MIT-IBM Watson AI Lab

2020–2021

Annual Report

Table of Contents

Message from the Directors	1
Lab by the Numbers	2
Selected Projects	4
People	14
Events	16
Students / Campus Engagement	18
Current Members	20
In the Media	22

Message from the Directors

Each fall, we look forward to the return to campus, but this year feels special. We'll have a chance to reunite with colleagues and friends in person after more than 1.5 years apart. This year, the MIT-IBM Watson AI Lab moves into its new home in Kendall Square. Built largely during the pandemic, the new Lab headquarters will facilitate collaboration between MIT and IBM researchers.

We'll also be welcoming our first class of interns into this new space. Starting this fall, the Lab will fund several students in MIT's Masters of Engineering program to work with us. With guidance from the Lab's researchers, they will address hand-picked projects that solve urgent business problems and have a strong chance of success. We're excited to see what fresh ideas the students come up with and to play a role in shaping their education.

This past year was difficult in so many ways, but we pulled through and accomplished quite a lot, from filing multiple patents with the U.S. Patent and Trademark Office to publishing outstanding papers in the top machine learning conferences. We hope you enjoy reading some of our highlights from the past year.

Aude Oliva, MIT and David Cox, IBM

Lab by the Numbers

160
researchers across MIT and IBM

$240M
10 year investment to found a joint lab

100+
Principal Investigators across 5 schools and 1 college

Over 380
project proposals submitted from MIT and IBM

7
member companies

100+
active projects

91 UROPs

Over 460
peer-reviewed papers

30
patent disclosures

Selected Projects

Active Fairness	Adversarial Learning	AI for Medical Imaging	AI for Quantum
Artificial Synapses	Audio Enhancement	Bootstrapping Fluid Intelligence	Brain-Inspired Vision
Contact Tracing	Creative AI for the Circular Economy	Cross-Species Data Translation	Debiased Object Recognition
Efficient Robustness	Explainable Financial Forecasting	Explainable Graph Learning	Exploiting Unlabeled Data
Learning from Small Data	Learning Priors	Lifelong Learning for Transfer	Lifelong RL for Zero-Shot Transfer
Model Fusion	Modeling Vaccine Manufacturing	Multi-Agent Planning	Neuro-Symbolic Integrations
Personalized Patient Models	Personalized Q&A	Planning in Complex Domains	Privacy, Fairness, & Causal Inference
Robust Probabilistic Programming	Sepsis Detection	Sketching AI Summaries	Stable ML Models

AI Information Theory	AI Mental Models	AI on IoT	AI Protein Design	AI Qubit Control	Anomaly Detection
Break the Curse of Dimensionality	Building Virtual Worlds	Causal Inference	Certifiably Safe AI	Clinical AI Translator	Collaborative Reinforcement Learning
Deep Phenotyping	Designing Silicate Materials	Dynamic Online Forecasting	Dynamic Resource Allocation	Dynamic Treatment Strategies	Economically Attractive AI
Generative Scene Decomposition	Graph Optimization	Intuitive AI	Inverse Design	Ising Machines	Learning 5000 languages
Machine Common Sense	Malware Detection	Memristor AI	Mixing Sight and Sound	ML and Cryptography	ML in Hilbert Space
Neuro-Symbolic Language	Neuromorphic Computing	Next-Gen Hearing Aids	Next-Gen Hearing Aids	Optimizing Drug Discovery	Patient Risk Stratification
Programs for Safe AI	Provably Trustworthy AI	Quantifying Robustness	Requirements for Trust	Robust Neural Networks	Robust Optimization
Supply Chain Optimization	Synthetic Drug Trials	Temporal Recommendation Algorithms	Tunneling Risk Management	UI for AutoML	Video Interpretation

AI is moving at an astonishing pace, making it hard to predict which ideas and technologies will land and grow into something bigger, and which will fade away. The MIT-IBM Watson AI Lab's strategy, as a result, has been to build a diverse portfolio and invest in projects with both near- and long-term payoffs.

We'll also be welcoming our first class of interns into this new space. Starting this fall, the Lab will fund several students in MIT's Masters of Engineering program to work with us. With guidance from the Lab's researchers, they will address hand-picked projects that solve urgent business problems and have a strong chance of success. We're excited to see what fresh ideas the students come up with and to play a role in shaping their education.

We are also trying to remove barriers to increase adoption of AI. To that end, we're working on methods for extracting feature representations from unlabeled data to give AI models more examples on which to train. We continue to experiment with embedding basic physics into AI models so that they can learn how objects relate and influence each other, which could make training more intuitive and less data-intensive.

The following is a snapshot of what we will be working on this year.

Incorporating linguistic structure into deep learning models

Jacob Andreas (MIT), Regina Barzilay (MIT), Tommi Jaakkola (MIT), Jim Glass (MIT), Roger Levy (MIT), Shiyu Chang (IBM), Yang Zhang (IBM)

Nearly 7,000 languages are spoken in the world today, but fewer than a dozen have the massive training data that machine learning models need to learn how to translate from one language to another, among other applications. Much of the world, as a result, is unable to quickly get answers to questions, retrieve information with a simple command or translate important documents not in their native tongue. In this project, researchers will leverage text from high-resource languages like English to develop better representations of low-resource languages like Quechua; weave grammar and other linguistic rules into AI models to better understand learned representations; and develop algorithms that can learn and transfer linguistic structures across languages.

Faster, cheaper AI with "neuromorphic" chips that work like neurons in the brain

Bilge Yildiz (MIT), Ju Li (MIT), Jesus del Alamo (MIT), John Rozen (IBM)

AI software has made stunning progress with the popularization of deep neural networks. Hardware innovation has lagged by comparison, with most chips still processing electrical signals digitally, in the binary logic of 1s and 0s, and off and on switches. The brain, by contrast, processes signals in an analog fashion, in bursts of varying intensity, which uses far less energy. The ability to recreate this process in a brain-like chip depends on the ability to precisely control electrical signals. In this project, researchers are exploring the use of ion intercalation to achieve multiple resistance states in the chip's channel material, tungsten oxide, while using as little energy as the brain.

Bringing deep learning to "internet of things (IoT)" devices

Song Han (MIT), John Cohn (IBM)

The branch of AI that curates your social media feed and serves up search results could soon check your vitals or set your thermostat. MIT's Song Han is working to bring deep neural networks to the tiny computer chips in wearable medical devices, household appliances, and the billions of other gadgets that make up the "internet of things." One recent system to come out of his lab, MCUNet, designs compact neural networks that allows AI applications to run smoothly on IoT devices despite their limited memory and processing power. The technology could facilitate the expansion of the IoT universe while saving energy and improving data security.

Can deep learning models be trusted?

Luca Daniel (MIT), Pin-Yu Chen (IBM), Sijia Liu (IBM)

As AI systems automate more tasks, the need to quantify their vulnerability and alert the public to possible failures has taken on new urgency, especially in safety-critical applications like self-driving cars and fairness-critical applications like hiring and lending. To address the problem, MIT-IBM researchers are developing a method that reports how much each individual input can be altered before the neural network makes a mistake, on their own or through a malicious attack. The team is now expanding the framework to larger, and more general neural networks, and developing tools to quantify their level of vulnerability based on many different ways of measuring input-alteration.

A better way to manage inventory

Georgia Perakis (MIT), Pavithra Harsha (IBM)

Companies that manage to stock just enough inventory to meet demand usually have higher profit margins. But accurately forecasting demand and supply at each step in a long supply chain is a tricky process, often because important data points are missing. Further, current methods to forecast supply chains often analyze one variable at a time, and don't capture the full story of influences. In this project, researchers are developing an interpretable, forecasting method that combines machine learning and inventory optimization, evaluating several drivers of supply chains over time, to make more accurate predictions. In this way, companies can make better decisions.

How do silica materials form?

Rafael Gomez-Bombarelli (MIT), Teodoro Laino (IBM), Leonid Kahle (IBM), Federico Zipoli (IBM)

Industrial silica is found in toothpaste, semiconductors, and even telescopes. Under different processing conditions, silicate minerals precipitate to form silica, creating materials with a wide range of properties. Though silica is everywhere, the structural changes responsible for creating its many properties are poorly understood. In this project, researchers will use neural networks to replace expensive quantum simulations with surrogate machine-learning models that can simulate and optimize chemical reactions to produce atom-by-atom structures of desired materials.

How do private, aggregated data influence decisions?

Dean Eckles (MIT), Naoise Holohan (IBM)

Personal data is increasingly aggregated and analyzed under a privacy model known as differential privacy which maintains individual privacy while allowing government, researchers, and companies to glean insights from large datasets. Despite the sudden increase in differential privacy methods, relatively little is known about the downstream effects, and whether resulting decisions and medical interventions are fair. In this project, researchers will explore how these methods serve as a series of randomized, accidental experiments, what can be learned about the effects of interventions, and what groups and places are more likely to be experimented on. They will also develop new methods for making use of privacy-preserving data.

Exploiting unlabeled data in supervised learning

Gregory Wornell (MIT), Prasanna Sattigeri (IBM)

One barrier to the wider adoption of AI is a shortage of images and other data with descriptive labels that allow machine learning models to find patterns and make predictions. In this project, researchers will develop tools for training models to extract feature representations from unlabeled data and transfer this information to supervised learning tasks. The research emphasizes application domains involving video, image, audio, and text data.

Toward an AI system that understands cause and effect

Joshua Tenenbaum (MIT), Antonio Torralba (MIT), Chuang Gan (IBM)

Deep learning models have delivered remarkable advances in AI, from image classification to speech recognition. But most of their success comes down to basic pattern matching. To push the capabilities of AI systems, researchers have created a physical world made up of colliding objects and pairs of questions and answers that can learn why things happen. The dataset, called CLEVRER, is aimed at testing AI models in their understanding of how objects relate to and influence one another. To beat the test, the researchers have developed a hybrid AI model that combines statistical deep learning with the interpretability of symbolic programs.

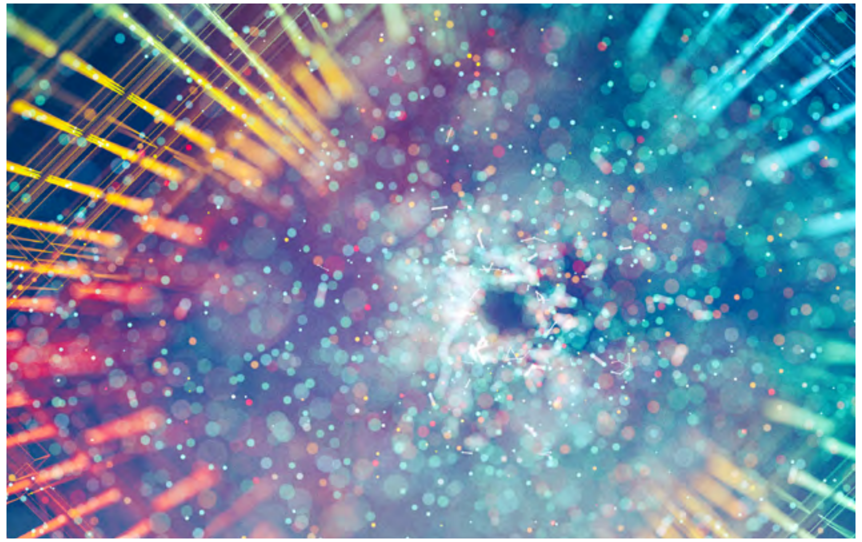

Learning and planning in hybrid domains

Leslie Pack Kaelbling (MIT), Michael Katz (IBM)

Policies and value functions are one way to train an intelligent agent, but such methods are data intensive, bad at explaining their decisions, and prone to failure in dynamic, complex environments. In this project, researchers will develop efficient, model-based planning methods for domains that require both discrete choices of what actions to take and numerical choices of precisely how to carry them out. Their proposed methods will need less data, be able to generalize across domains, acquire knowledge cumulatively, and provide explanations for their decisions.

Detecting changes, gaps and other pitfalls in time-series data

Devavrat Shah (MIT), Yang Zhang (IBM)

Working with time-series data to forecast, say, the closing price of a stock, can be tricky. Statistical errors can arise from missing or corrupted data, even without the complexity of analyzing continuously evolving events. The stock market, for example, behaved differently in the 1980s than it did before and during the pandemic, making accurate forecasting difficult. In this project, researchers will develop methods to improve the use of time-series data for forecasting in ever-shifting environments while also identifying anomalies in the data. Their methods will be incorporated into Time Series PredictDB, a database that stores time-series data and lets users query prediction tasks rather than build their own machine learning model.

Shrinking the environmental footprint of concrete

Stefanie Jegelka (MIT), Elsa Olivetti (MIT),
Richard Goodwin (IBM), Nghia Hoang (IBM)

Concrete is the most common building material on Earth, with current production methods contributing a major share of global carbon emissions. In this project, researchers will leverage AI to find more sustainable, high-performing and cost-competitive concrete mixtures. They will use text analysis tools to explore the scientific literature for concrete mixture designs and use generative modeling to test different approaches to achieve greater environmental performance. They will also investigate use of waste materials in alternative concrete binders to further reduce concrete's carbon footprint.

Merging deep nets with symbolic programs for robust decision-making

Armando Solar-Lezama (MIT), Nathan Fulton (IBM)

As AI systems evolve, more of their decisions will be made with less human supervision. Without strong and explicit safety guarantees, the expanded use of automated decision-making could have disastrous consequences for society. In this project, researchers will develop a family of code-writing algorithms that combine data-rich neural networks with traditional symbolic programs to ensure that AI systems update their assumptions about the world as the environment around them changes.

Advancing graph deep learning to analyze relationships at scale

Charles E. Leiserson (MIT), Jie Chen (IBM),
Toyotaro Suzumura (IBM)

The powerful recommendation algorithms guiding us to new products, friends and webpages are based on graph algorithms that analyze complex relationships among millions of data points. Combined with deep neural networks, which excel at picking out patterns in images and sequence data, graph algorithms have the potential to predict market behavior and much more. In this project, researchers will address the computational challenges of graph neural networks while making their decisions more explainable.

People

Aude Oliva

MIT Director, MIT-IBM Watson AI Lab

David Cox

IBM Director, MIT-IBM Watson AI Lab

Anantha Chandrakasan

MIT Chair, MIT-IBM Watson AI Lab; Dean, MIT School of Engineering

Dario Gil

IBM Chair, MIT-IBM Watson AI Lab; Senior Vice President and Director, IBM Research

Daniel Huttenlocher

MIT Co-Chair, MIT-IBM Watson AI Lab; Dean, MIT Stephen A. Schwarzman College of Computing

Events

What's next in AI

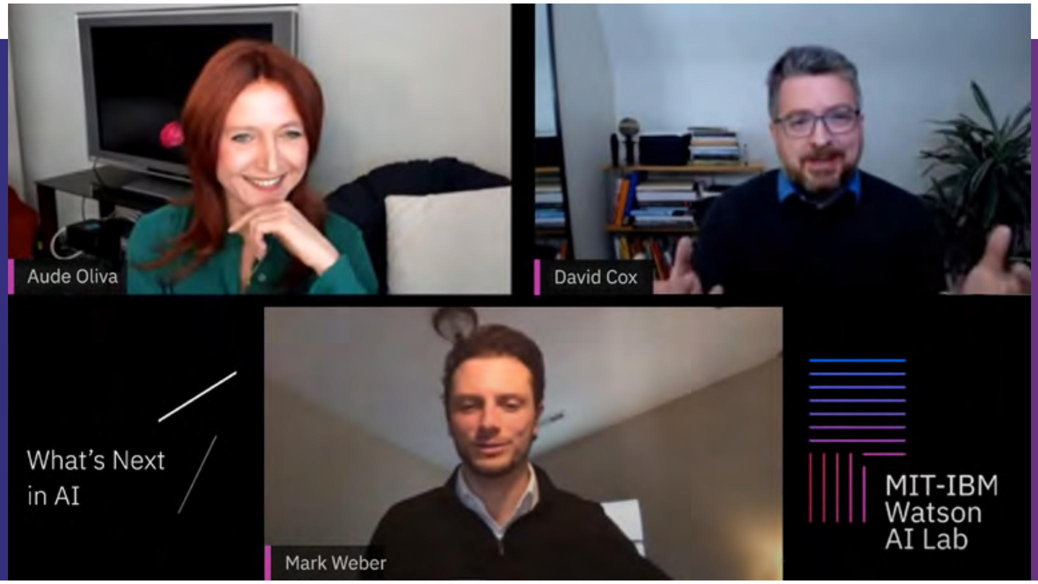

In this virtual event, scientists and business experts from the MIT-IBM Watson AI Lab came together to review three key barriers to adoption of AI – trust, scalability, and reasoning – and how we can solve these challenges through scientific advancement. These sessions took place over three weeks in November 2020.

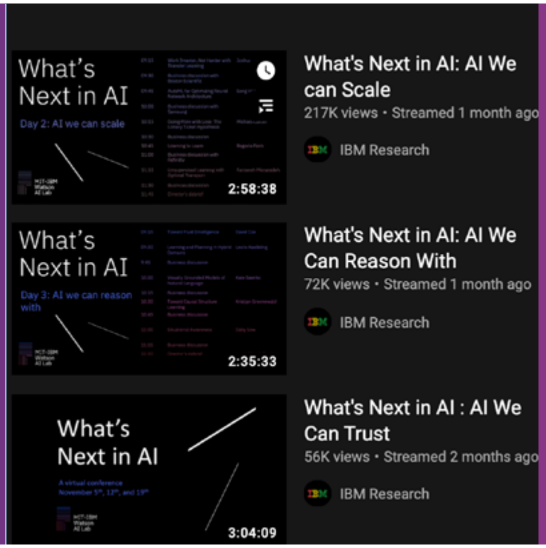

3 half-day sessions in November

AI we can trust
AI we can scale
AI we can reason with

YouTube Data

Views: 346,470
Hours watched 149,500
Average view time: 26 minutes
Unique viewers: 279,200
12.2 million impressions

Poster Session

During the virtual MIT-IBM Watson AI Lab open house, member companies and researchers were invited to learn about the Lab's current research portfolio, as well as various research projects proposed for next year. Additionally, there was a virtual, asynchronous poster session, where participants got the opportunity to speak with researchers and learn more about their projects.

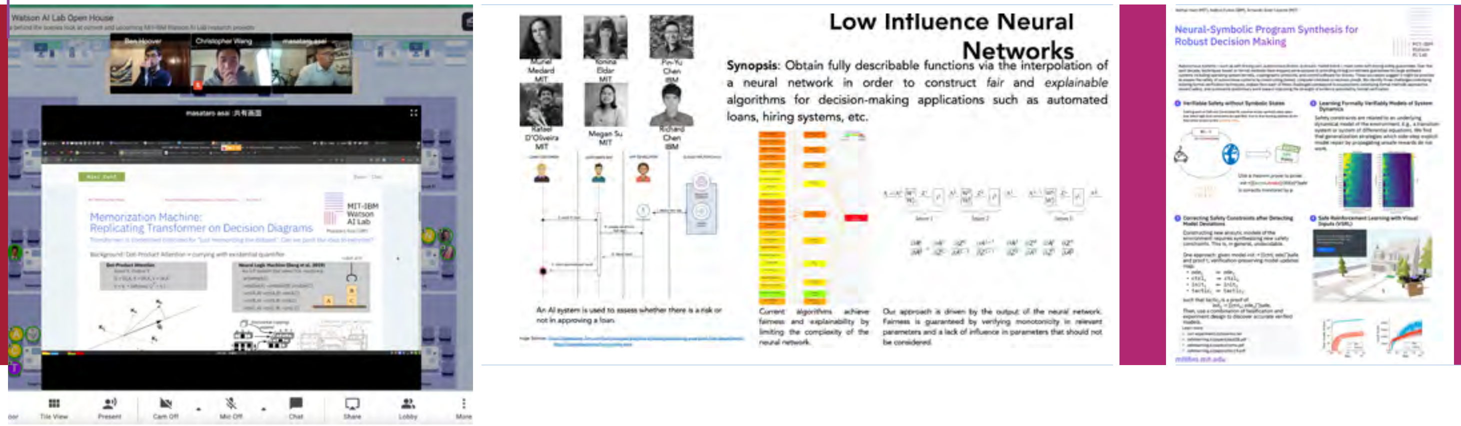

Technical Seminars

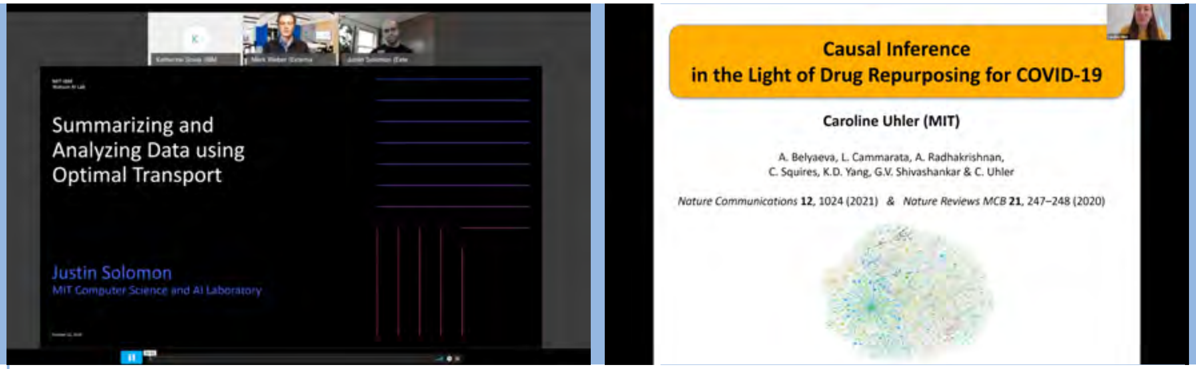

The Technical Seminar Series is a monthly offering for the MIT-IBM Watson AI Lab's corporate membership program hosted by leading MIT and IBM researchers in the field of AI. This series is designed to educate our members' employees on the latest advancements in the field of AI. While the seminar is open to any employee of our member companies, the talks are geared towards those with a technical or semi-technical background.

Students / Campus Engagement

The Lab funded the MIT Undergraduate Research Opportunities Program (UROP) and the SuperUROP program to support the research skills of MIT undergraduates. The research projects students work on push the frontiers of AI and its applications.

The Lab selected four Masters of Engineering students from the Electrical Engineering and Computer Science (Course 6A) program to conduct remote internships at IBM in Cambridge, MA. The 6A program allows students to pursue a cutting-edge industry project while simultaneously working towards the completion of their graduate degree.

Spring 2021 UROPs

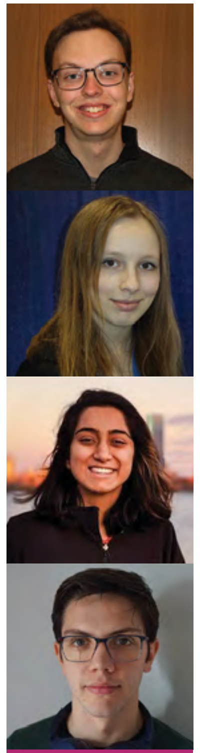

Fall 2020 UROPs

6A (IBM interns)

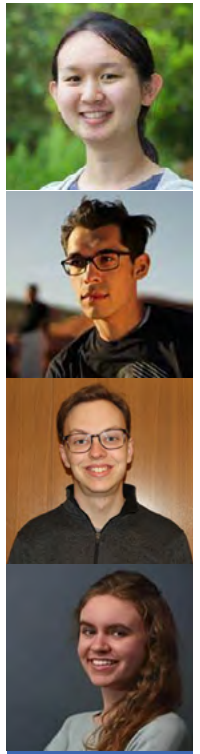

The MIT-IBM Watson AI Lab funded 10 MIT Summer Research Program (MSRP) undergraduate students. This program works to promote the value of graduate education, improve the research enterprise through increased diversity, and increase the number of underrepresented and underserved students in the research enterprise.

As part of the program, IBM hosted an Industry Night career panel for all MSRP students to discuss career paths beyond academia. Eighty MSRP students tuned in to hear from three field experts regarding their own personal experiences, advice, and areas of expertise.

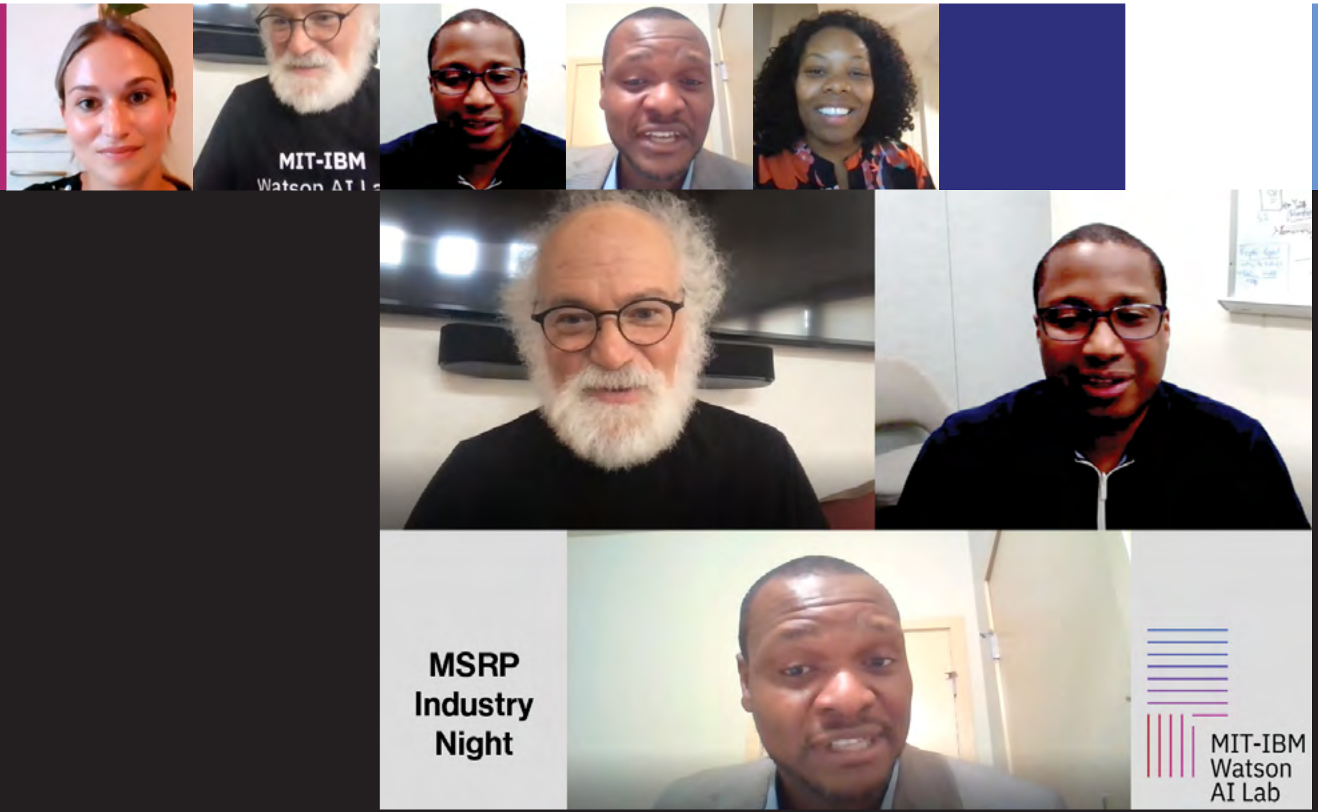

Current Members

The MIT-IBM Waston AI Lab's Membership Program, formally announced in 2019, invites companies to co-invest in the Lab in exchange for intellectual property and other Membership benefits. All of our members are market and technology leaders in their respective industries. They bring a breadth of perspective and strategic influence by providing real-world challenges where innovation is needed, helping us achieve our mission to push the envelope of what AI can do and to make a positive impact on business and society.

Samsung
A peer in AI research, the Samsung Advanced Institute of Technology is Samsung Group's R&D hub, established as the incubator for cutting-edge technologies under the founding philosophy of Boundless search for breakthroughs, and guided by the vision of Changing the World through Creative Research.

London Stock Exchange Group
LSEG is a leading global financial markets infrastructure and data provider. Our ambition is to shape the future of our industry by serving our customers and markets better. Through extensive experience, deep knowledge and worldwide presence across financial markets, we enable businesses and economies to fund innovation, manage risk and create more jobs.

Nexplore
Nexplore is the innovation arm of Grupo ACS, whose global portfolio of companies comprises the largest construction and industrial services company in the world. From the high-speed rail project in California to some of the world's most beautiful airports, Grupo ACS companies are building the world's most critical infrastructure.

We expected cutting edge innovation from the MIT-IBM Watson AI Lab. With their researchers and our financial services expertise we have been able to explore real market challenges.

– Hanna Helin, Global Head of Technology Innovation and Ecosystems, LSEG

The team at the MIT-IBM Lab made the effort to understand our business and our technology roadmap, and crafted together with the Woodside team an engagement that is aligned with where we want to go and achieve business impact.

– Shelly Kalms, CDO, Woodside

Following an extensive evaluation of options to invest in research programs, the combination of top-notch researchers, a broad project portfolio, and a flexible participation model led us to choose the MIT-IBM Watson AI Lab. Our experience as a member so far has been outstanding.

– Dr. Yves Gorat Stommel, Deputy Head of Function, Evonik Digital

Woodside
Woodside is the pioneer of the liquid-natural gas industry in Australia and a leading investor in renewable energy. A supporter of the Paris Agreement, Woodside aims to be net zero in its operations by 2050 and believes natural gas has a significant role to play in the world's energy transformation as the lowest emission hydrocarbon fuel.

Boston Scientific
Boston Scientific is dedicated to transforming lives through innovative medical solutions that improve the health of patients around the world. Boston Scientific's businesses include neuromodulation, interventional cardiology, cardiac rhythm management, electrophysiology, peripheral interventions, endoscopy, urology and pelvic health.

Wells Fargo
Founded in 1852 and headquartered in San Francisco, Wells Fargo provides banking, insurance, investments, mortgage, and consumer and commercial financial services through more than 8,400 locations and offices in 42 countries. With approximately 268,000 team members, Wells Fargo serves one in three households in the United States and employs about one in 600 working Americans.

Evonik
Evonik is a world leader in specialty chemicals. Evonik strives to provide the materials making tires fuel-efficient, mattresses more elastic, medications more effective, and animal feeds healthier. That's what specialty chemicals are all about. And when it comes to specialty chemicals, Evonik is among the best in the world. The company is active in more than 100 countries around the world and generated sales of €12.2 billion and an operating profit of €1.9 billion in 2020. Evonik goes far beyond chemistry to create innovative, profitable and sustainable solutions for customers.

The collaboration is great for Boston Scientific because it allows us to continue innovating with new technology and allows the AI Scholars, like myself, to learn and be up to date on what's happening in my field of interest.

– Umanga Poudel, Data Scientist, Boston Scientific

In the Media

The MIT-IBM Watson AI Lab has been cited or featured by media outlets including multiple mentions in *Academic Times, Forbes, Gizmodo, Medium, Morning Brew, Venture Beat*, and *Wired*.

The Academic Times

Machine learning could yield protein-based coatings that curb food waste

By Tara DiMaio
May 1, 2021

A new computational model combines artificial intelligence with biology to design and fold proteins in a method that could yield specialized proteins for different applications, such as an edible coating for crops that reduces food waste.

In a recent development in embodied AI, scientists developed a new challenge that will help assess the ability of AI agents in finding paths, interacting with objects, and planning tasks efficiently.

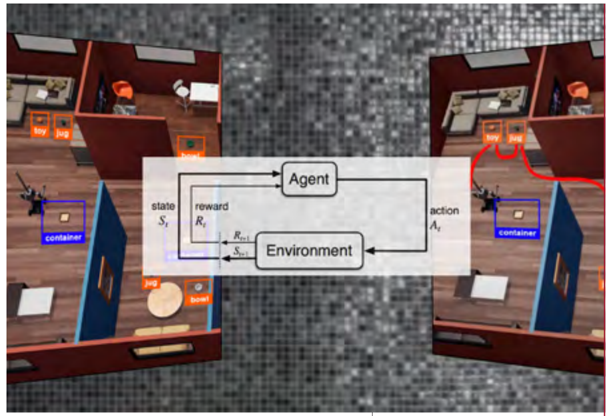

Reinforcement learning challenge to push boundaries of embodied AI

By Ben Dickson
April 26, 2021

TechTalks

Fruit Fly Brain Hacked For Language Processing

Discover

By the Phsics ArXiv Blog
Jan 28, 2021

This team has hacked the fruit fly brain network to perform other tasks, such as natural language processing. It's the first time a naturally occurring network has been commandeered in this way.

AI Algorithms Are Slimming Down to Fit in Your Fridge

By Will Knight
December 13, 2020

Wired

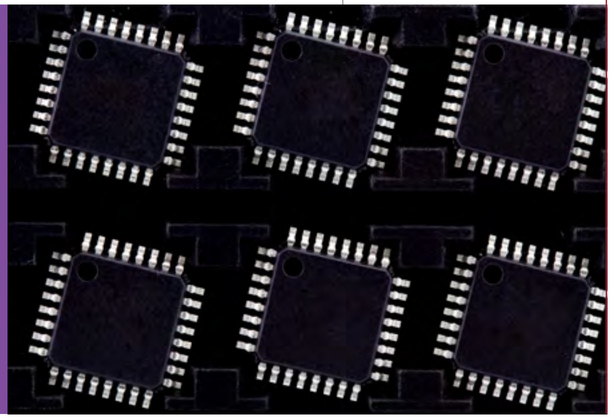

Researchers showed it is possible to squeeze a powerful AI vision algorithm onto a simple, low-power computer chip that can run for months on a battery. The trick could help bring more advanced AI capabilities, like image and voice recognition, to home appliances and wearable devices, along with medical gadgets and industrial sensors.

Scientists have developed a series of tests that will help evaluate the capacity of AI models to reason like children, by observing and making sense of the world.

Venture Beat

Can you teach AI common sense?

By Ben Dickson
July 27, 2021

MIT-IBM
Watson
AI Lab

MIT Schwarzman
College of Computing

MIT | School of Engineering

IBM Research