MIT-IBM Watson AI Lab

A New Leaf

Entering into its sixth year, the Lab begins a new chapter this fall. The changing of the seasons brings fresh opportunities to campus: new collaborators, thoughtful seminars and events, and exciting projects. Our portfolio will include new projects spanning topics such as financial decision-making, efficient computing, fairness and privacy, accelerated scientific discovery, and climate and sustainability.

Again, young minds will play a significant role in achieving project goals. Over the next year, the Lab will host numerous masters and graduate students for networking events and internships. Here, we'll provide them with opportunities to immerse themselves in AI/computing development, learn about translational research processes, and build out their market-critical skillsets, preparing them for the next step in their careers.

We look forward welcoming you into the Lab and tackling your persistent computing questions with new ideas and novel challenges with enthusiasm.

David Cox and Aude Oliva, directors of the MIT-IBM Watson AI Lab

In the Lab

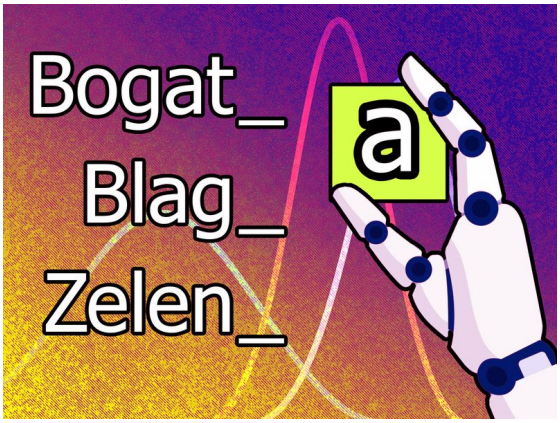

AI that can learn the patterns of human language

Machine-learning model discovers linguistic rules, often matching up with human-made ones.

Exploring the interaction of phonology and morphology, Lab researchers from MIT Professor Tenenbaum's team have developed a new machine-learning model that can learn the rules and patterns of human languages on its own, emulating reasoning and knowledge of linguists. The work can then be applied to many languages and could inform how infants acquire language.

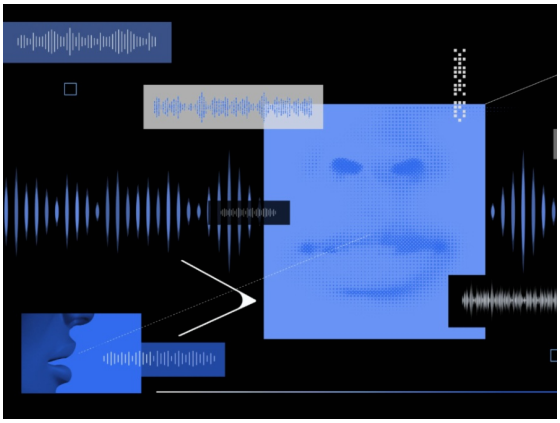

Speech recognition benefits from one voice

Converting several audio streams into one improves speech-processing

Spoken language comes with variations in accent, rhythm, and intonation—characteristics that can trip up speech recognition systems and obfuscate language nuances. To address this, Lab researchers added a voice-conversion algorithm to the cutting-edge speech-processing model, improving performance on key tasks.

How to create business value with AI

[Explore 12 real-world examples from organizations making it a reality.]

In a report, researchers from the Lab and the IBM Institute for Business Value interviewed global business and technology experts to reveal myths about AI implementation and show through data and real-world examples how companies are using AI to create real business value.

Boston-area undergraduates develop a taste for careers in tech

[Break Through Tech AI event promotes diversity in computing and career paths]

A program launched this summer in the MIT Stephen A. Schwarzman College of Computing aims to provide mentorship and computing opportunities for local, women and underrepresented groups. IBM researchers and managers supported this pipeline of talent development by sharing how they achieved success in the AI sector during a panel discussion.

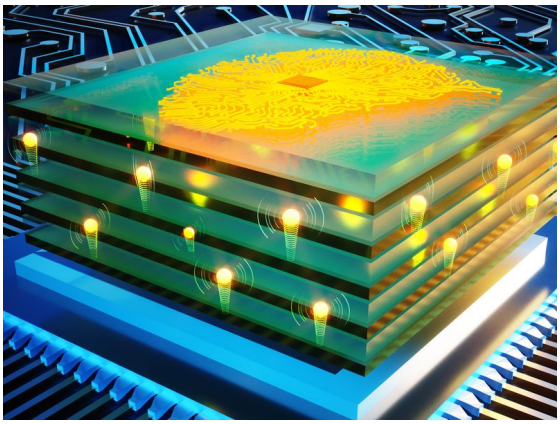

New hardware offers faster computation for AI

[Analog deep learning technique propels protons through solids at unprecedented speeds.]

A fabrication method from the Lab teams of Jesús del Alamo, Ju Li, and Bilge Yildiz allows a network of analog artificial “neurons” and “synapses” to run 1 million times faster than previous versions, which is also about 1 million times faster than the synapses in the human brain, and at a fraction of the energy cost.

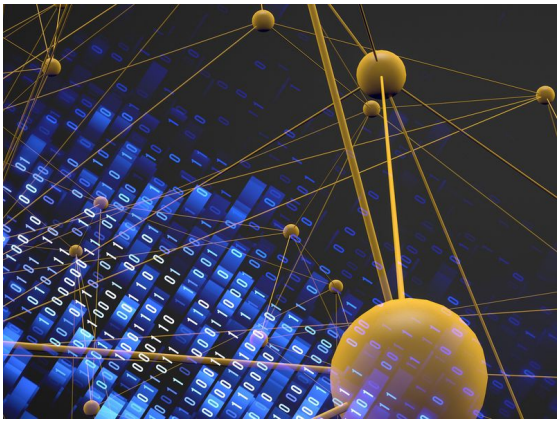

Explained: How to tell if AI is working the way we want it to

[Interpretability methods can be a double-edged sword, researchers say.]

How do Lab researchers check that a deep-learning model is working properly?
MIT's Marzyeh Ghassemi group explains three popular interpretability methods. While these can be helpful for debugging and quality assurance, the explanations should be taken with a grain of salt.

Giuseppe Romano

[Modeling energy materials and devices for a brighter, more efficient future]

MIT and Lab researcher Giuseppe Romano uses multiscale models to simulate engineered nanomaterials to discover ones with highly-efficient thermoelectric properties, with applications for renewable energy. “We’ve had to hack nature and create semiconducting nanostructures that combine the properties we need in terms of electrical and thermal conductivity,” says Romano.

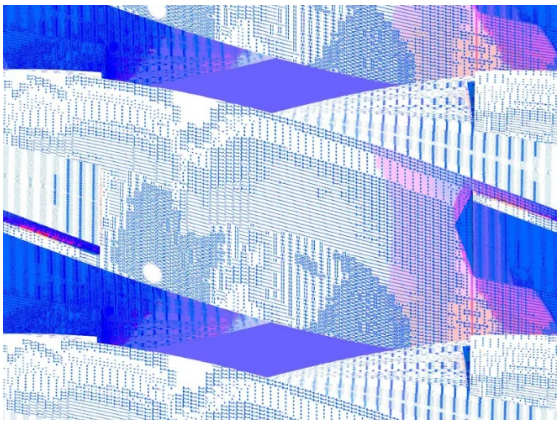

Explained: What is confidential computing?

[Confidentiality, integrity, and attestation are key aspects to security.]

With the boom in cloud computing, new types of security threats have emerged. Confidential computing is a solution to the added IT security issues of working with the cloud. IBM Research breaks down the components that that go into cloud computing and how researchers are addressing open questions.

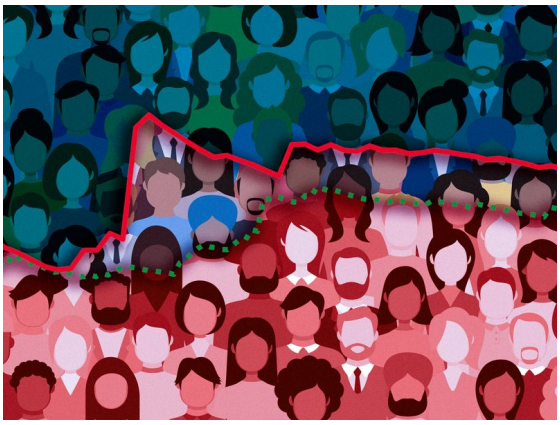

A technique to improve both fairness and accuracy in AI

[Technique to mitigate disparities among minority subgroups in machine learning models]

Selective regression can improve the overall predictions of a machine learning model, but can do the opposite for underrepresented subgroups. A technique from Lab researchers, working with the Wornell group, fixes the flaw, improving fairness and accuracy.

Teaching AI to ask clinical questions

[More efficiently finding information in a patient’s health record]

Patient electronic health records (EHR) contain a plethora of information, but it's not easily accessible. Lab researchers working with the Peter Szolovits group trained a machine-learning model to generate high-quality, authentic clinical questions and are working toward building a model that can automatically answer physicians’ questions in an EHR.

Explained: What is federated learning?

[Processing data at their source]

Federated learning is a way to train AI models without anyone seeing or touching your data, offering a way to unlock information to feed new AI applications, for instance, in healthcare and finance. However, there are calculated trade-offs for this privacy: accuracy, efficiency, transparency, and incentives for good behavior.

Marzyeh Ghassemi

[Updating the state of the art]

Lab researcher and MIT Professor Marzyeh Ghassemi leads the Healthy ML Group at MIT, where she creates and applies machine learning to understand and improve health and healthcare in ways that are novel, robust, private, and fair.

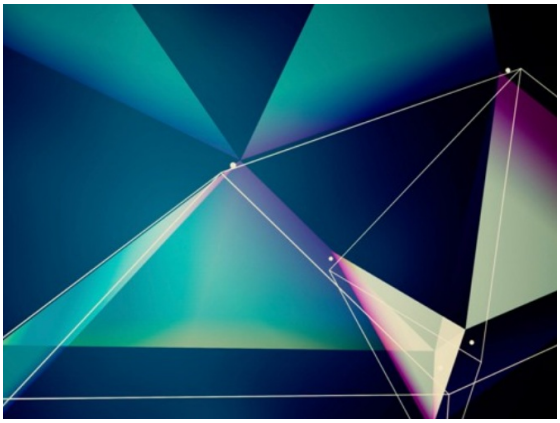

Exploring the Applications of Geometric Data

[Justin Solomon on extensions to medical vision, 3D animation and autonomous vehicles]

Justin Solomon's group uses geometric approaches for analyzing data and processing geometric data for work in computer vision, autonomous vehicles, 3D shapes, high-dimensional data, and computer graphics. He says, he's developed useful and intuitive tools that don't require detailed knowledge of mathematics.

Manipulating stock prices with an adversarial tweet

[Simple adversarial attacks can trick prediction models.]

Predictive models can give off an air of trustworthiness, but their robustness might be unfounded. Work from IBM Research demonstrates how, for instance, stock price models can be fooled through targeted tweets, and Lab research shows how a patterned t-shirt can cloak someone in a people-detecting model, so it's important to understand vulnerabilities.

In the Media

"Artificial synapse" could make neural networks more like brains

Researchers from the Lab have created an "artificial synapse" device, specifically, a network of nanoscale resistors, that can transmit protons across terminals at speeds 10,000 times faster than biological synapses, reports *New Scientist.* Further, this analog machine learning technique requires little energy and is efficient possibly "offer[ing] advantages over digital machine learning."

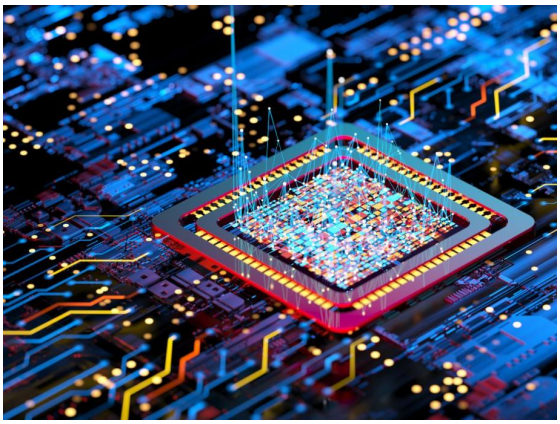

Perceptron: Analog AI chips

MIT researchers from the Lab have developed a new hardware that offers faster computation for AI with less energy, reports *TechCrunch*. “The researchers’ processor uses ‘protonic programmable resistors’ arranged in an array to ‘learn’ skills.” The study's lead author MIT postdoc Murat Onen says that analog processors can provide a leg-up over digital ones. “You will be training networks with unprecedented complexities that no one else can afford to, and therefore vastly outperform them all. In other words, this is not a faster car, this is a spacecraft.”

Computational method to create synthetic 3D spider web structures

Lab researcher Markus Beuchler and a colleague have developed a way to model complex 3D spider webs, which are structurally strong, but not well understood, reports *AIP Scilight.* Their work combines "advanced modeling techniques to relate spider web graph microstructures to effective mechanical properties, with a focus on strength and toughness. They used deep neural networks for web property prediction, which are trained on graph-structured web data and simulated mechanical properties."

Selective risk could improve AI fairness and accuracy

The Lab's Wornell team has developed a technique, monotonic selective risk, that can help correct model biases and reduce error rates for underrepresented groups, reports *RT Insights.* To improve AI model fairness and accuracy, "instead of abstaining, one model includes sensitive attributes such as race and sex, and the other does not. In tandem, both models make decisions, and the model without the sensitive data is used as calibration for biases in the dataset."

How transformers seem to mimic parts of the brain

Transformers, which operate on a mechanism called self-attention, are being used to model the brain. Hopfield networks retrieve memories in a similar way. When trying to mimick memory retrieval, Lab researcher Dmitry Krotov and his colleague "proved that a transformer-based Hopfield network was biologically plausible," reports *Quanta Magazine.*

Beyond deep learning

Grappling with real-world environments that can be confusing to AI, so models are now incorporating two modes of thought, fast/automatic and slow/conscious, reports *PNAS*. For instance, complex scene perception, Lab researcher Dan Gutfreund and colleagues developed a neuro-symbolic system, 3DP3, that uses this technique to handle variation. "Perhaps most importantly, 3DP3's novel combination of existing ideas — symbols, networks, and Bayesian inference — offers much more scope for incorporating capabilities that go beyond scene recognition."

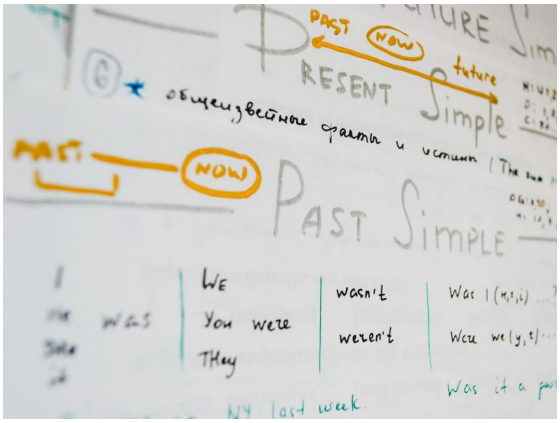

AI learns and recognizes language norms and patterns

The Lab's Tenebaum team recently used an Bayesian Program Learning approach to create a model capable of learning grammar, *Dataconomy* reports. "This model can also learn higher-level linguistic patterns that can be applied to several languages, allowing it to achieve greater outcomes."

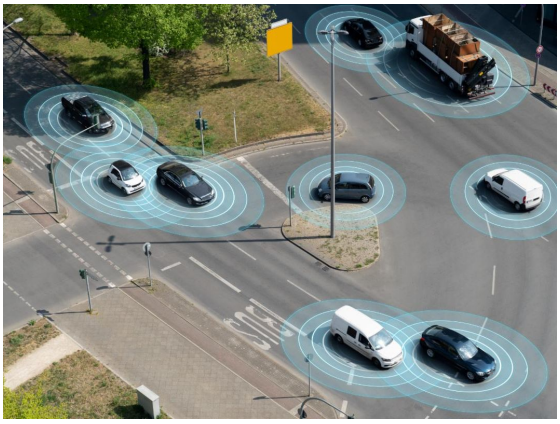

Green Driving

Loh Down on Science host Sandra Tsing Loh spotlights work from the group of Lab researcher Cathy Wu to develop a new method for self-driving vehicles that would help minimize idling at red lights. “In their method, self-driving can be taught to minimize stops at red lights. To make this work, traffic lights and self-driving cars would have sensors. This would let them check in with each other on their surroundings,” says Loh.

Upcoming Events

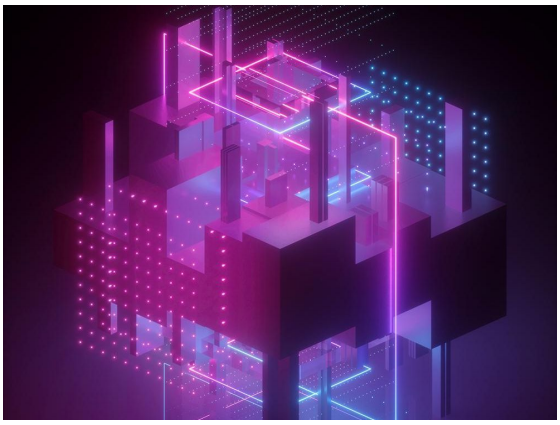

AI Transformation in Analog Systems

September 28, 12-1p.m. ET
RLE 36-462 and 36-428

In this in-person talk, MIT MTL and the AI Hardware Program will host Wenjie Lu of Analog Devices Inc., who will discuss how his team leverages AI-for-Hardware to tackle challenges in complex analog systems beyond the reach of traditional algorithms and how AI provides new possibilities in design, optimization, and testing of future analog systems. Recording available after the talk.

AI Policy Forum Summit

September 28, 9a.m.-5p.m. ET
Live webcast

Hosted by the MIT AI Policy Forum — an initiative of the MIT Schwarzman College of Computing, join leaders in government, business, and academia for the AI Policy Forum Summit. The program will focus on the global policy challenges surrounding deployment of AI in many societal areas and consider possible future developments and concrete guidance on implementing AI-related policies. Register here.

Lab Highlights

MIT professor Justin Solomon wins the SIGGRAPH Significant New Researcher Award, recognizing an individual who has made recent, significant contributions to the field of computer graphics and is new to the field.

MIT's Faez Ahmed is a recipient of the 2022 3M Non-Tenured Faculty Award, recognizing his research on the development of deep generative models for engineering design problems.

Markus Buehler of MIT receives 2022 IACM Fellows Award for his extraordinary research accomplishments and publications in computational mechanics.

Lab-supported "Distributed Adversarial Training to Robustify Deep Neural Networks at Scale" received Best Paper Runner-Up Award at the UAI 2022 conference.

MIT professor Georgia Perakis appointed one of the new associate deans of the Social and Ethical Responsibilities of Computing (SERC), a cross-cutting initiative in the MIT Stephen A. Schwarzman College of Computing.

MIT professors Justin Solomon, Kevin O'Brien, Phillip Isola, Isaac Chuang receive EECS awards in teaching, innovation, and inclusion.

Online Learning

Unsupervised Machine Learning: Unlocking the Potential of Data, a joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course, begins November 16.

Making AI Work: Machine Intelligence for Business and Society, a joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course, begins October 19.