MIT-IBM Watson AI Lab

In the Lab

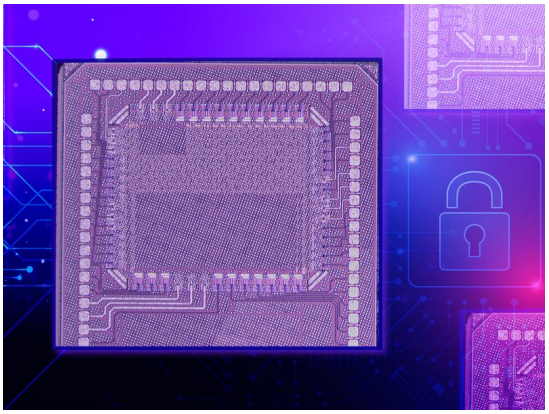

Tiny chip can safeguard user data while enabling efficient computing

A security solution for power-hungry AI models that offers protection against two common attacks

Research from the Lab groups of Anantha Chandrakasan — Lab MIT chair and MIT's chief innovation and strategy officer — John Cohn, and Xin Zhang has produced a machine-learning accelerator (digital in-memory compute) that can help keep sensitive data private, while enabling huge AI models to run efficiently on devices, like smartphones.

3 Questions: Enhancing last-mile logistics with machine learning

Using AI to make vehicle routing more efficient and adaptable for unexpected events

Lab researcher Matthias Winkenbach shares how traditional operations research falls short when it comes to optimizing how packages are delivered on a large scale, and how his machine-learning research with the Lab, leveraging models from natural language processing could unlock several advantages.

A crossroads for computing at MIT

[The MIT Schwarzman College of Computing building will form a new cluster of connectivity]

“The building is designed to be the computing crossroads of the campus," says Daniel Huttenlocher, Lab MIT co-chair, dean of the MIT Schwarzman College of Computing, and the Henry Ellis Warren Professor of Electrical Engineering and Computer Science. "It’s a place to bring a mix of people together to connect, engage, and catalyze collaborations in computing, and a home to a related set of computing research groups from multiple departments and labs,” including the Lab.

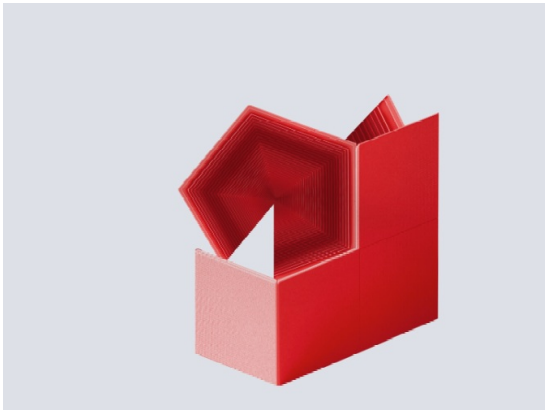

What is red teaming for generative AI?

[A way of interactively testing AI models to protect against harmful behavior, bias, and data leaks]

Bad actors frequently probe systems like foundation models for vulnerabilities to steal data or disrupt service. Generative AI poses additional risks, making adversarial AI-testing and re-alignment crucial for safe, secure, and trustworthy AI. IBM researchers have developed datasets and tools, including work from the Lab groups of Aldo Pareja, James Glass, Akash Srivastava, and Pulkit Agrawal that identifies uncommon prompts that could produce harmful results.

Faster, better way to prevent an AI chatbot from giving toxic responses

[A curious machine-learning model finds a wider variety of prompts for training a chatbot.]

Work from the Lab groups of Aldo Pareja, James Glass, Akash Srivastava, and Pulkit Agrawal developed a better way to red team. The team created a reinforcement learning approach using curiosity-driven exploration which rewards the system for generating more novel prompts that elicit toxic chatbot responses.

In the Media

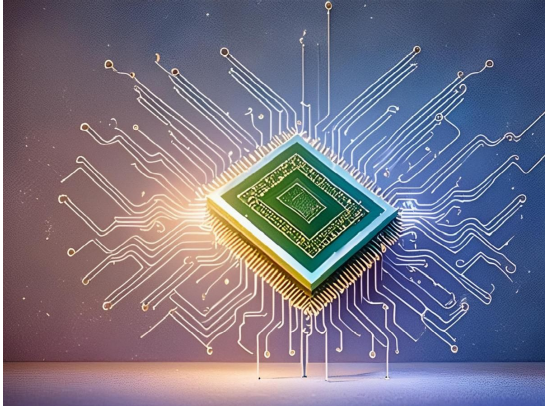

MIT's new powerful chip thwarts millions of data theft attacks

"This new chip will adopt a three-part approach to protect AI devices from data attacks," reports Interesting Engineering. "As security has become a critical issue in the design of edge devices, there is a need to develop a complete system stack focusing on secure operation," said Anantha Chandrakasan, MIT chief innovation and strategy officer and Lab MIT chair, about his team's new digital in-memory compute chip.

How AI will step off the screen and into the real world

Lab researcher Daniela Rus shares on theTED stage how AI and robotics are converging in her lab. Rus highlights her group's "liquid networks" (bio-inspired AI), and text-to-robot design. This Lab research, DiffuseBot, generates soft robots and considers physical constraints, before iterating on their design in a simulated physical environment.

Upcoming Events

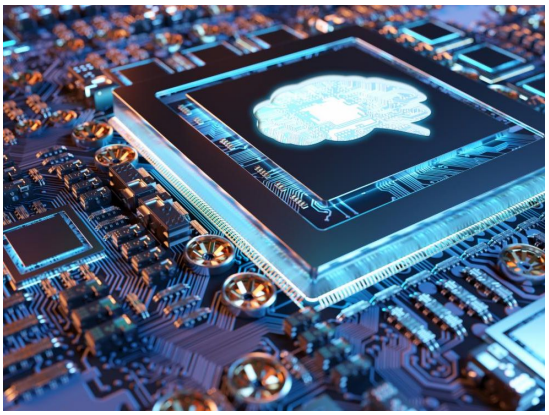

MIT AI Hardware Program

2024 Annual Symposium
May 1, 9:30 a.m.-3 p.m. ET

In a hybrid event, MIT AI Hardware Program researchers will provide project reviews of the current portfolio as well as provide exposure to new projects. The event will be hosted by Lab co-director Aude Oliva and Lab researcher Jesús del Alamo, both of whom are the MIT AI Hardware Program leads. The talks will be followed by interactive demos and posters. Registration required.

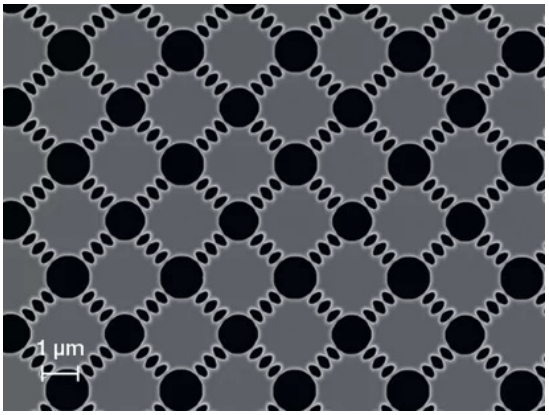

Accelerated Discovery with Differentiable Programming: From Nanomaterials to Chiplets Design

Nano Explorations
May 14, 11-11:45 a.m. ET

In a virtual event, Lab researcher Giuseppe Romano will cover Automatic Differentiation (AD) software. He will report on recent materials and systems optimization efforts, along with the implemented open-source software. One application pertains to chiplets floorplan design, where, in collaboration with the Lab, the researchers design a framework for minimizing the maximum temperature during operation. More information and registration is required.

Lab Highlights

Lab researcher William Oliver was has been elected to the2023 class of the American Association for the Advancement of Science (AAAS) Fellows, recognizing scientifically and socially distinguished achievements.

Lab researchers Antonio Torralba and Phillip Isola published a MIT Press textbook"Foundations of Computer Vision."

Online Learning

Artificial Intelligence: Implications for Business Strategy
A joint MIT CSAIL and MIT Sloan School of Management Course begins
May 29.

Machine Learning in Business
A joint MIT CSAIL and MIT Sloan School of Management Course begins
June 5.

Making AI Work: Machine Intelligence for Business and Society
A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
June 5.

Unsupervised Machine Learning: Unlocking the Potential of Data
A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
June 12.

AI in Robotics: Learning Algorithms, Design and Safety
A Professional Education Course begins
July 10.

Reinforcement Learning
A Professional Education Course begins
July 29.

Advanced Reinforcement Learning
A Professional Education Course begins
August 1.