MIT-IBM Watson AI Lab

In the Lab

The future of AI is open

[The age of value creation with AI]

At IBM's [Think], Darío Gil — Lab IBM chair, IBM SVP and director of Research — presented IBM workflows and innovations that allow businesses to effectively and securely deploy AI. One from from Lab researchers Shivchander Sudalairaj, Abhishek Bhandwaldar, Aldo Pareja, Kai Xu, Akash Srivastava, and Lab co-director David D. Cox, called InstructLab, allows users to tune a model with new skills and knowledge specific to a need without losing generality in its capabilities quickly.

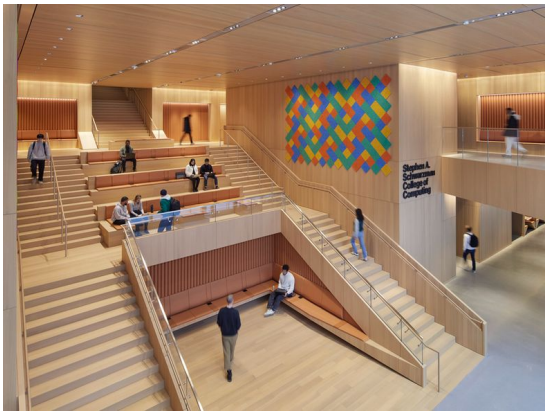

New faculty at the MIT Schwarzman College of Computing

[Fulfilling its commitment to create 50 new faculty positions]

Lab researchers Sara Beery, Yoon Kim, Sherrie Wang, and Connor Coley are some of the more recent hires, whose work spans computer vision to monitor biodiversity to streamlining discovery in the chemical sciences.

Better cardiovascular care through AI

[Collin Stultz wants to help heart patients by applying machine-learning techniques to cardiovascular medicine.]

Lab researcher Collin Stultz is working to leverage health data from wearable devices to ECGs along with machine learning to pick up on subtle physiological changes that can impact and predict heart health.

Biometrics in the age of artificial intelligence

[Lessons from the past can inform the ethical use of government-sponsored identification systems.]

Michelle Spektor PhD ’23, the MIT-IBM Postdoctoral Fellow in Computing and Society in the MIT Schwarzman College of Computing’s Social and Ethical Responsibilities of Computing (SERC), examines how biometric identification and national surveillance has changed over the years globally and how it can impact populations and communities.

Using ideas from game theory to improve the reliability of language models

[A new “consensus game” elevates AI’s text comprehension and generation skills.]

The Lab teams of Jacob Andreas, Gabriele Farina, and Yikang Shen have developed an equilibrium-ranking algorithm that harmonizes generative and discriminative querying to enhance prediction accuracy across various tasks, outperforming larger models.

Creating bespoke programming languages for efficient visual AI systems

[Jonathan Ragan-Kelley optimizes how computer graphics and images are processed.]

Lab researcher Jonathan Ragan-Kelley develops new programming languages that make it easier to write programs that run efficiently on increasingly complex hardware of today and tomorrow.

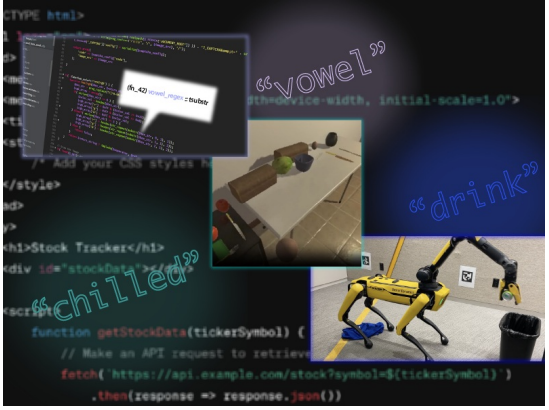

Natural language boosts LLM performance in coding, planning, and robotics

[Neurosymbolic methods help models find better abstractions within natural language.]

Work from the Lab groups of Jacob Andreas and Joshua Tenenbaum has produced frameworks, dubbed LILO and Ada, that can synthesize, compress, and document code; and explore sequential decision-making for AI agents, respectively.

Vincent Sitzmann

[Next steps for AI: Creating 3D understanding from 2D images]

Computer vision and Lab researcher Vincent Sitzmann develops neural networks that can interpret and represent 2D images and video in 3D with the goal of understanding real-world physics. To do this he explores and applies differential rendering, uncertainty analysis in 3D reconstruction, and unsupervised algorithms that understand 3D geometry of objects in a scene from video in order to create models that can perceive.

In the Media

A new AI discovery sure looks Like the dawn of true machine reasoning

Abstracting information out of natural language can help improve the reasoning of LLMs to better mimic that of humans, and as *Popular Mechanics* reports, techniques like LILO and Ada from the teams of Joshua Tenenbaum and Jacob Andreas demonstrate ways models can gather more context for better decision-making.

QoQ and QServe: A new frontier in model quantization

Research from the Lab groups of Song Han and Chuang Gan has created a "progressive group quantization" technique that reduces accuracy losses that would result from standard quantization methods used for large language models, reports *Marktechpost*.

Game theory can make AI more correct and efficient

How can a reward system and a large language model pitted against itself improve its accuracy and consistency? *Quanta Magazine* reports on work from Lab intern Athul Paul Jacob and researchers Gabriele Farina, Jacob Andreas, and Yikang Shen, who are developing a consensus game based on Diplomacy to do just that.

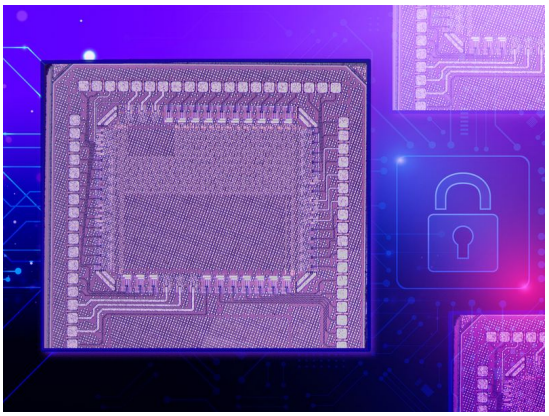

New powerful chip thwarts millions of data theft attacks in tests

Research from the Lab groups of Dean Anantha Chandrakasan — Lab MIT chair and MIT's chief innovation and strategy officer — John Cohn, and Xin Zhang has produced a machine-learning accelerator chip to make health-monitoring apps more secure, reports *Interesting Engineering*. "The researchers subjected this new chip to intensive testing, simulating real-world hacking attempts, and the results were impressive."

Event Recordings

SERC Workshop: Eric Horvitz, Chief Scientific Officer of Microsoft
In a recorded fireside chat, Dan Huttenlocher, dean of the MIT Schwarzman College of Computing and MIT Lab co-chair, spoke with Eric Horvitz, Chief Scientific Officer of Microsoft, on the multifaceted dimensions of responsible AI, as part of a regular series of workshops presented by the Social and Ethical Responsibilities of Computing (SERC).

Lab Highlights

Lab researcher Pulkit Agrawal received an IEEE 2024 Early Academic Career Award in Robotics and Automation "for pioneering contributions to self-supervised robot learning and advancing sensorimotor control for contact-rich, dynamic, and dexterous tasks."

Lab researcher Jesús del Alamo and his colleague receive Intel's 2023 Outstanding Researcher Award Outstanding Researcher Award for work on "Exploring the Limits of Vertical-Nanowire Tunnel Field-Effect Transistors in the Nanoscale."

Darío Gil — SVP, director of IBM Research, and the Lab's IBM chair — has been elected as chair of the National Science Board. He is the NSB's first industry chair in over 30 years.

Lab researcher Anette Peko Hosoi received an Australian Academy of Science Selby Fellowship.

Former Lab intern Junhong Lin is a winner of the 2024 J. Francis Reintjes Excellence in 6A Industrial Practice Award presented for "outstanding performance in a 6A work assignment and/or exceptional quality in an M.Eng. Thesis performed at a 6A company."

Lab researcher Wojciech Matusik received a Humboldt Research Award

Lab researchers Rameswar Panda, Akash Srivastava, Yikang Shen, and Lab co-director David D. Cox led the development of IBM's open source Granite code models and InstructLab Project to bring open innovation to AI.

Lab researchers Daniela Rus and Piotr Indyk were elected to the National Academy of Sciences. Rus was honored for her ongoing research achievements in artificial intelligence and soft robotics, while Indyk was inducted for his contributions to algorithms addressing geometry and massive data problems.

Lab researcher Julie Shah named head of the MIT Department of Aeronautics and Astronautics.

Lab researcher Julian Shun promotioned to associate professor with tenure.

Yu Wang, a student in Lab research Justin Solomon's group, received an honorable mention for SIGGRAPH's Outstanding Doctoral Dissertation Award or his dissertation “Geometric Computing beyond the Laplacian.”

Lab researcher Yury Polyanskiy was named a 2024 IEEE Fellow "for contributions to information measures and finite-blocklength information theory."

Lab researchers had 21 papers accepted to the International Conference on Machine Learning (ICML), a leading international academic conference in machine learning.

Online Learning

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.