




In the Lab



Five takeaways from this year's IBM AI Hardware Forum

[The challenges and opportunities in designing infrastructure to support future AI workloads](#)

The 2023 IBM AI Hardware Forum highlighted advancements in AI-specific chips and infrastructure, focusing on energy-efficient technologies, the importance of open-source AI, and the need for skilled AI and semiconductor workers.



New method uses crowdsourced feedback to help train robots

[Human Guided Exploration \(HuGE\) enables AI agents to learn quickly with help from humans.](#)

Researchers from the Lab's Pulkit Agrawal group and elsewhere have developed a new reinforcement learning approach that doesn't rely on an expertly designed reward function. Instead, it leverages crowdsourced feedback, gathered from many nonexpert users, to guide the agent as it learns to reach its goal.



Merging science and systems thinking to make materials more sustainable

[Elsa Olivetti is passionate about materials science "from the atom to the system."](#)

For Lab researcher and Professor Elsa Olivetti, tackling a problem as large and complex as climate change requires not only lab research but also understanding the systems of production that power the global economy. With a holistic view, she innovates to improve the life cycle of materials.



Students pitch transformative ideas in generative AI at MIT Ignite competition

[Twelve teams presented innovative startup ideas with potential for real-world impact.](#)


During an event co-sponsored by the Lab, teams of students and postdocs put forth entrepreneurial ideas that employ generative AI, judged to be the most feasible and provide significant potential for real-world impact.



Technique enables AI on edge devices to keep learning over time

[On-device training is significantly sped up without any dip in accuracy.](#)

The PockEngine training method from the Lab groups of Song Han and Chuang Gan enables deep-learning models to efficiently adapt to new sensor data directly on an edge device.



The latest AI safety method is a throwback to our maritime past

[The origins of 'AI governance' and addressing generative AI's potential risks and harms](#)

IBM's Kush Varshney discusses core issues of fairness, trust, and transparency and tools to address them, like those embedded in watsonx, IBM's platform for enterprise AI.



Explained: Generative AI


[How do these systems work, and what makes them different from other types of AI?](#)

MIT News explains that the technology "can be thought of as a machine-learning model that is trained to create new data, rather than making a prediction about a specific dataset. The potential applications are immense, for instance at MIT, Lab researchers are using it to create synthetic data and designing new proteins or materials.

What is AI alignment?

[Encoding human values and goals into large language models](#)


IBM Research explains that AI alignment, which frequently is a two-step process, involves programming AI to understand and adhere to human values, goals, and ethics, ensuring its actions and decisions are beneficial and aligned with human intentions.



Language gives robots a better grasp of an open-ended world

[Blending 2D images with foundation models to build 3D feature fields](#)

New research from the Lab groups of Leslie Pack Kaelbling and Phillip Isola enables robots to interpret open-ended text prompts using natural language, helping the machines manipulate unfamiliar objects. The system's 3D feature fields could be helpful in environments that contain thousands of objects, such as warehouses.




Quantum's impact on security

[A Conversation with MIT CSAIL and MIT Center for Quantum Engineering](#)

Experts and Lab researchers Daniela Rus, Peter Shor, William Oliver, and Vinod Vaikuntanathan provide some points about what the current field of quantum computing research looks like, its impact on security, and how the work happening today might impact society in the future at a recent MIT talk.



In the Media



Google, IBM make strides toward quantum computers that may revolutionize problem solving

Dario Gil —SVP, director of IBM Research, and the Lab's IBM chair — speaks with Scott Pelley on [60 Minutes](#) about the capabilities of this technology. "We are now in a stage where we can do certain calculations with these systems that would take the biggest supercomputers in the world to be able to do some similar calculation," says Gil. Further, a million or a billion of these supercomputers will not be able to keep up with the future of these machines.



AI tipping point

AI experts, including Lab researcher Yoon Kim, weigh in on the incredible benefits and potential dangers of AI, particularly generative AI, with [Curiosity Stream](#). "These large language models are trained to predict the next word, given context. They take all the sentences that you see out on the web now. Given a sentence of, let's say five words, this results in five different prediction tasks. Fundamentally, it's actually not that surprising that accurate prediction of a next word results in models that seem intelligent. I think what surprised people was that this actually happened."



President Biden signs executive order to require guidelines for AI

President Biden recently passed an AI order requiring development standards, labels and watermarks, and safety test results to be shared. [NewsCenter 5](#) spoke with an Lab researcher Regina Barzilay about the challenges of regulating AI.



IBM's Gil says unrealistic to call for global AI body

Dario Gil —SVP, director of IBM Research, and the Lab's IBM chair — tells Bloomberg TV that, instead "existing enforcement agencies should be empowered to ring-fence AI's risks." Further, players in the industry should incorporate AI governance and specific use cases to ensure accountability against risks.



IBM's AI Academy

IBM is providing education to help get companies and users on track to leverage AI through its [AI Academy](#). Dario Gil —SVP, director of IBM Research, and the Lab's IBM chair — and others provide background to help jumpstart a workforce, becoming a value creator with generative AI, learning elements of enterprise AI and seeing how it could be used in business operations.

Event Recordings

[Analog In-Memory Computing For Deep Learning Inference](#)

Abu Sebastian of IBM Zurich speaks about analog in-memory computing for deep learning, detailing a prototype chip with phase-change memory and concluding with future research trajectories in the field.

[Generative AI Week](#)

The MIT community recently came together to "come to grips with the tectonic forces of generative AI – to understand its potential, contain its risks, and harness its power for good." Experts, including a large lineup of Lab researchers, provided insights and analysis on the ever-evolving AI landscape.

[SERC Fireside Chat](#)

Dan Huttenlocher, dean of the MIT Schwarzman College of Computing and MIT Lab co-chair, spoke with Mira Murati, chief technology officer of OpenAI, about her experience leading the teams behind ChatGPT and Dall-E, and shared her opinions on the explosive growth of generative AI and the social and ethical issues in computing technologies. Their conversation is part of a series hosted by the Social and Ethical Responsibilities of Computing (SERC).

Lab Highlights

Lab researcher Bilge Yildiz was selected as a [2023 Fellow of the Electrochemical Society](#).

Lab researcher Dava Newman in the Department of Aeronautics and Astronautics received an [Honorary Doctorate from the Royal College of Art](#), London, U.K.

Lab researcher Josh Tenenbaum was named a [Schmidt Futures AI2050 Senior Fellow](#).

Online Learning

[Artificial Intelligence: Implications for Business Strategy](#)

A joint MIT CSAIL and MIT Sloan School of Management Course begins
January 31, 2024.

[Machine Learning in Business](#)

A joint MIT CSAIL and the MIT Sloan School of Management Course begins
February 7, 2024.

[Making AI Work: Machine Intelligence for Business and Society](#)

A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
March 13, 2024.

[Unsupervised Machine Learning: Unlocking the Potential of Data](#)

A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
March 20, 2024.