

In the Lab



Creating space for the evolution of generative and trustworthy AI

[Industry Showcase](#)

The Lab and its member companies delved into the history and future of explainable AI and trustworthy foundation models, its intersection with different parts of society and business, and academic-industry partnership's role in its growth at this year's Industry Showcase.



Building a foundation for the future of AI models

[Dario Gil on the future of generative AI in enterprise at THINK 2023](#)

Dario Gil, SVP, director of IBM Research, and the Lab's IBM chair, delivers a keynote, encouraging everyone to be "value creators" in this era. Here, he demystifies the technology — including behind watsonx, which leverages some Lab-supported research — and shares a set of principles to guide generative AI business strategy.



3 Questions: Jacob Andreas on large language models

[Creating state-of-the-art models and investigating how language can enhance other types of AI](#)


Lab researcher Jacob Andreas examines the mechanics, implications, and future prospects of the technology at hand, delving into issues like "in-context learning," hallucinations, inaccuracies in its assertions, and how this work will likely progress in the near term.



Introducing the technology behind watsonx.ai

[IBM's AI and data platform for enterprise](#)


Lab co-director and IBM VP David Cox describes IBM's work to build a suite of enterprise-grade, IBM-trained foundation models, including their approach to data and model architectures. IBM's new platform and tooling includes watsonx.ai studio, which uses Lab work, for instance, in fairness and model "growth" from smaller models.



A tool for accurately simulating complex systems

[Eliminating a source of bias in simulations](#)


Work from the team of Lab researcher Devavrat Shah draws on principles of causality and enables unbiased trace-driven simulations. The new technique could help researchers design better algorithms for a variety of applications, including improving video quality on the internet and increasing the performance of data processing systems.



Open-source platform simulates wildlife for soft robotics designers

[SoftZoo, a soft robot co-design platform, can test for optimal performance in different environments.](#)


A Lab team led by Joshua Tenenbaum, Daniela Rus, and Chuang Gan has developed "SoftZoo," a way for researchers to investigate soft robot morphology, control, and behaviors within a variety of feature-rich, simulated environments.



Machine learning model finds genetic factors for heart disease

[The model can predict heart-related traits and drive genetic discovery in heart disease.](#)


Multimodal work from the Lab's Caroline Uhler, Kenney Ng and others investigates cardiac physiology in health and diseases. Here, they use a machine-learning approach that leverages electrocardiograms and MRIs to learn patterns and identify characteristics of a patient's heart. This tool could potentially be used to detect and diagnose heart conditions.



How IBM is tailoring generative AI for enterprises

[Developing generative foundation models that are trustworthy, energy efficient, and portable](#)


AI, machine learning and deep learning are transformative technologies. However, history of some applications has shown that they do not come without risks. In order for generative models to be useful and create value for businesses and industry applications, IBM, with research contributions from the Lab, is developing them in a thoughtful and sustainable manner.



Generating novel proteins that meet structural design targets

[Tunable proteins could be used to create new materials with specific mechanical properties.](#)

A machine-learning system, which is based on an attention-based diffusion model and from the team of Lab researcher Markus Buehler, can generate protein designs with certain structural features, and which do not exist in nature. These could be utilized to make materials that have similar mechanical properties to existing materials, like polymers, but with a smaller carbon footprint.



Explainer: What is generative AI?

[Deep-learning models that can generate high-quality text, images, and other content.](#)

Generative AI has gone through many cycles of extensive growth and hype, from breakthroughs in computer vision to the rise of deep generative models, particularly for language, with different architectures including the pivotal transformer. In this era, Lab researchers David Cox, Akash Srivastava, and others are considering the power of supervised learning, prompt engineering, and compact models to achieve better output.



MIT researchers discuss frontiers of generative AI

[AI-generated code, language, and images as well as its capabilities, limitations, and future impact](#)

In a recent panel discussion, Lab researchers Phillip Isola, Daniela Rus, Armando Solar-Lezama, and Jacob Andreas touched on topics like using images to remove noise, text to image capabilities, word embeddings and transformer models, and the power of "attention." Together, they looked at model benchmarking and academia's role in the future of this technology.

In the Media




IBM intros a slew of new AI services, including generative models

At Think 2023, IBM announced watsonx, giving creators access to a "toolset, infrastructure and consulting resources they need to create their own AI models or fine-tune and adapt available AI models on their own data," reports [TechCrunch](#). "You still need a very large organization and team to be able to bring [AI] innovation in a way that enterprises can consume," Dario Gil, SVP, director of IBM Research, and the Lab's IBM chair said.




Will artificial intelligence help — or hurt — medicine?

Doctors, particularly in the U.S., are inundated with paperwork and natural language AI models propose a potential solution, reports [NPR](#), but not without bias, a serious concern. "When you take state-of-the-art machine-learning methods and systems and then evaluate them on different patient groups, they do not perform equally," the Lab's Marzyeh Ghassemi shares.



Using smaller models to train large language models

Linear Growth Operator (LiGO), a method from the Lab researchers Rameswar Panda, Leonid Karlinsky, Rogerio Feris, David Cox, and Yoon Kim, allows for "creators of these LLMs to integrate old models into new development," reports [RTInsights](#). The technique enables researchers to train these models with a lower computational and time cost.



Some glimpse AGI in ChatGPT. others call it a mirage

With the launch of powerful large language models like GPT-4, which is behind ChatGPT, some researchers and the general public wonder if we're seeing the makings of artificial general intelligence (AGI). Lab researcher Josh Tenebaum says that, "GPT-4 is remarkable but quite different from human intelligence in a number of ways," [WIRED](#) reports.



Quadrupeds are learning to dribble, catch, and balance

Dribblebot, from the Lab group of Pulkit Agrawal, can dribble a soccer ball on different landscapes, "using only onboard sensing and computation for this task, and it was first trained extensively through reinforcement learning in simulation," reports [IEEE Spectrum](#). Looking further, "the real problem that's being solved here is legged locomotion while manipulating an occasionally adversarial object in the real world. This obviously opens up other potential applications."

Event Recordings

Lab researchers Jim DiCarlo and Pulkit Agrawal present at TEDxMIT, "[Neuroscience+AI can unlock hidden visual interface for the emotional brain](#)" and "[Why it's harder for AI to open doors than play chess.](#)" respectively.

Lab member program lead Kate Soule shares [how to build enterprise-ready foundation models](#).

Lab Highlights

Lab researcher Dina Katabi was elected as a member of the [National Academy of Sciences](#), which aims to "encourage education and research, recognize outstanding contributions to knowledge, and increase public understanding in matters of science, engineering, and medicine."

Lab researchers James DiCarlo and Piotr Indyk were elected to the [American Academy of Arts and Sciences](#).

Lab researcher Chuchu Fan received a [National Science Foundation Career Award](#).

Lab researcher Justin Solomon named [Edgerton Award](#) winner, which recognizes exceptional distinction in teaching, research, and service at MIT.

Lab researchers Jim DiCarlo and Josh McDermott received [MIT Department of Brain and Cognitive Sciences awards](#) for Excellence in Graduate Teaching, and DEIJ, respectively.

The MIT Case Studies in Social and Ethical Responsibilities of Computing (SERC) aims to advance new efforts within and beyond the MIT Stephen A. Schwartzman College of Computing and regularly shares [findings in its issues](#).

Online Learning

[MIT Introduction to Deep Learning](#)

A Lab-sponsored course teaches the fundamentals of deep learning.

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

A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
June 7.

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

A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
June 14.

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

A joint MIT CSAIL and MIT Sloan School of Management Course begins
June 28.

[Designing Efficient Deep Learning Systems](#)

An MIT Professional Course begins
July 17.

[Machine Learning in Business](#)

A joint MIT CSAIL and the MIT Sloan School of Management Course begins
July 26.