MIT-IBM Watson AI Lab

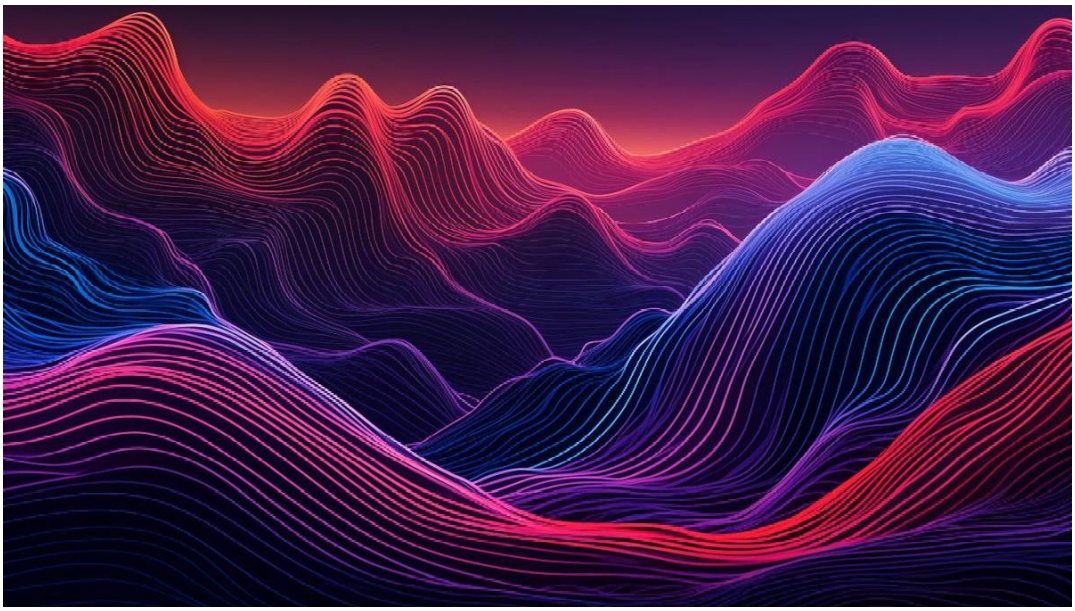

In the Lab

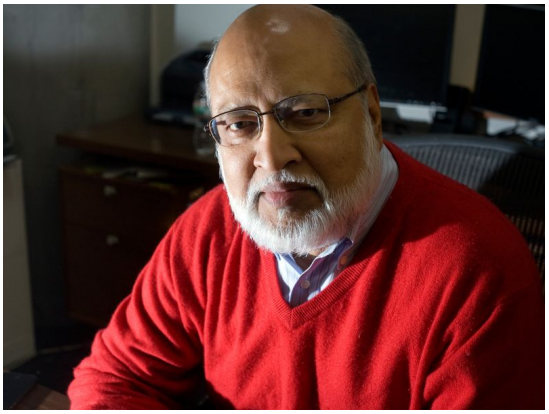

Arvind, MIT professor and prolific computer scientist, dies at 77

A dedicated teacher, academic leader, and Lab researcher

"Arvind was both a tremendous scholar in the fields of computer architecture and programming languages and a dedicated teacher, who brought systems-level thinking to our students. He was also an exceptional academic leader, often leading changes in curriculum and contributing to the Engineering Council in meaningful and impactful ways. I will greatly miss his sage advice and wisdom," says Anantha Chandrakasan, MIT Lab chair, chief innovation and strategy officer, and dean of engineering.

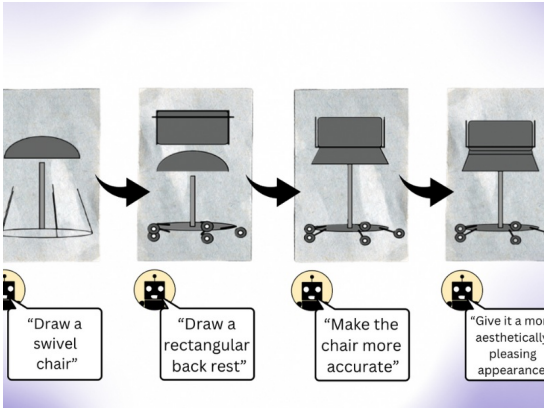

Understanding the visual knowledge of language models

Text-based large language models can be prompted to code better illustrations.

The teams of Lab researchers Phillip Isola and Antonio Torralba tested the ability of LLMs to draw and self-edit images, based on their internal knowledge. Further, training a computer vision model, that can recognize objects within real photos, with a synthetic dataset of these pictures

outperformed a model trained with real photos.

Researchers use large language models to help robots navigate

A new method uses language-based inputs to direct a robot through multistep navigation tasks.

Completing manual tasks for an AI agent requires several machine-learning models and expensive visual data. The Lab teams of Rameswar Panda, Rogerio Feris, Lab co-director Aude Oliva, Phillip Isola, and Yoon Kim instead developed a technique to uses visual text captions to train a model and then perform these functions.

QS ranks MIT the world's No. 1 university for 2024-25

Ranking at the top for the 13th year in a row

The Institute received a No. 1 ranking in the following QS subject areas: Chemical Engineering; Civil and Structural Engineering; Computer Science and Information Systems; Data Science and Artificial Intelligence; Electrical and Electronic Engineering; Linguistics; Materials Science; Mechanical, Aeronautical, and Manufacturing Engineering; Mathematics; Physics and Astronomy; and Statistics and Operational Research.

The world at MIT: Daniela Rus

A journey into robotics

In a short video, Lab researcher and CSAIL director Daniela Rus shares how a foundation in machining and math led her to a career in computation and robotics for practical applications and the advancement of technology.

Looking for a specific action in a video? This method can find it

Approach could streamline virtual training processes or aid clinicians in reviewing diagnostic videos.

A new, more efficient method from the teams of Lab researchers Samuel Thomas, Rogerio Feris, James Glass, and Hilde Kuehne uses only videos and their automatically generated transcripts to perform spatio-temporal grounding and identify particular actions in a long video.

School of Engineering welcomes new faculty

[Fifteen new faculty members join six of the school's academic departments.]

Many of these new faculty also report into the MIT Stephen A. Schwarzman College of Computing and specialize in research that intersects with multiple fields. "I am particularly struck by the interdisciplinary approach many of these new faculty take in their research. They are working in areas that are poised to have tremendous impact," says Anantha Chandrakasan, MIT Lab chair, chief innovation and strategy officer, and dean of engineering"

In the Media

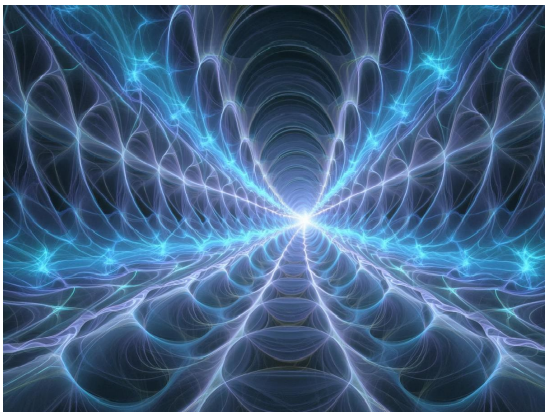

How physics can improve image-generating AI

A team of Lab researchers, led by Tommi Jakkola, is looking to improve image generation using the physics of electric fields to outperform diffusion models, reports The Economist. "Poisson flow generative models" (PFGMs) describe "the electric field created by static electrical charges. Judged by industry standards, PFGMs generate images of equal or better quality than state-of-the-art diffusion models, while being less error-prone and requiring between ten and 20 times fewer computational steps."

Researchers turn visual data into language to help robots navigate

Directing AI agents to perform tasks requires knowledge of the task as well as the environment. AI Business reports that work from the groups of Rameswar Panda, Rogerio Feris, Lab co-director Aude Oliva, Phillip Isola, and Yoon Kim has found a way to offer the agent visual information, using text descriptions, eliminating the need to train on expensive visual data. "The language-based approach even worked well in low-data settings, where only a few expert navigation examples were available for training."

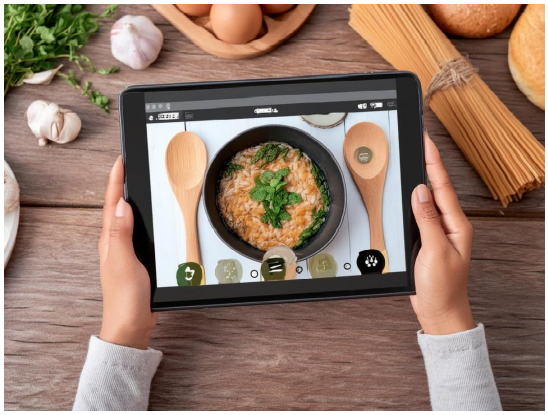

See what you mean

Searching for particular content within an unlabeled video can be tedious. The Lab groups of Samuel Thomas, Rogerio Feris, James Glass, and Hilde Kuehne developed a self-supervised technique to address this, reports Hackster.io. "During the evaluation, they found that their new approach was generally much more accurate in identifying specific actions in videos than existing methods. It also proved to be much better at identifying human-object interactions, which are crucial in identifying a great many actions of interest."

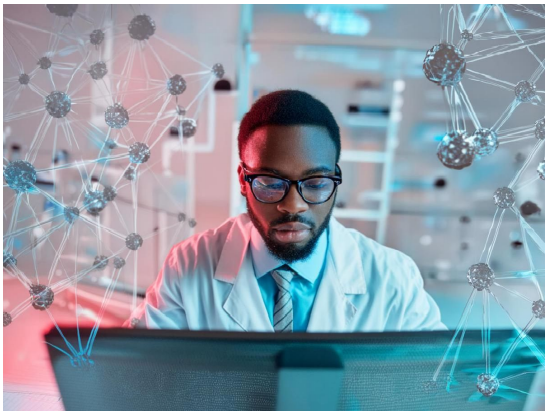

This AI paper introduces the 'scientific generative agent'

A new computational technique from the Lab can assist in the scientific process and help accelerate discovery, reports Marktechpost. Research from the groups of Joshua Tenenbaum, Daniela Rus, Chuang Gan, and Wojciech Matusik leverages LLMs and simulations to "transcend specific domains and offer a unified method for physical science. The framework combines the knowledge-driven, abstract reasoning abilities of LLMs with the computational strengths of simulations, providing a more comprehensive approach to scientific inquiry."

Event Recordings

Imagination in Action
Watch as distinguished speakers from MIT's CSAIL and EECS probe AI frontiers and implications. At the event, guests shared their insights and expertise, including Lab researchers Ramesh Raskar, Andrei Barbu, Armando Solar Lezama, Arvind Satyanarayan, Una-May O'Reilly, Daniela Rus, Yoon Kim, and Neil Thompson. Read more at Forbes.

Making sense of training large AI models
MIT graduate student Kwangjun Ahn, working with Lab researcher Ali Jadbabaie, defends his PhD in two parts: First, he investigates how training transformer-based models can lead to remarkable properties, such as in-context learning, and he also aims to understand the main challenges associated with transformer training. The second part of the thesis focuses on understanding practical optimization algorithms.

Lab Highlights

The MIT School of Engineering has awarded Lab researcher Leslie Kaelbling the 2024 Ruth and Joel Spira Award for Excellence in Teaching. The award is given annually to faculty members that embody the tradition of high-quality engineering education at MIT.

Lab co-director Aude Oliva is the recipient of The Justine and Yves Sergent Award 2024, honoring "a female researcher who has developed an international reputation in the field of cognitive neuroscience, in particular in connection with research on the hemispheric specialization or the interface between cognitive neuropsychology and functional brain imaging."

Lab researchers Caroline Uhler and Tamara Broderick named Fellows of the Institute of Mathematical Statistics (IMS). Uhler received the award "for interdisciplinary excellence and for merging mathematical statistics and computational biology in innovative and impactful ways." Broderick was acknowledged "for significant contributions to theoretical modeling and computational methodology at the intersection of Bayesian Statistical Machine Learning and Bayesian nonparametric theory and applications."

Lab researcher Nancy Kanwisher named a 2024 Kavli Prize Laureate in neuroscience "for the discovery of a highly localized and specialized system for representation of faces in human and non-human primate neocortex."

27 Lab-supported papers were accepted to premier conferences: 10 to the Conference on Computer Vision and Pattern Recognition (CVPR) and 17 to the International Conference on Learning Representations (ICLR).

Online Learning

Artificial Intelligence: Implications for Business Strategy
A joint MIT CSAIL and MIT Sloan School of Management Course begins
July 10.

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.

Machine Learning in Business
A joint MIT CSAIL and MIT Sloan School of Management Course begins
August 7.

Making AI Work: Machine Intelligence for Business and Society
A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
August 14.

Unsupervised Machine Learning: Unlocking the Potential of Data
A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
August 21.