MIT-IBM Watson AI Lab

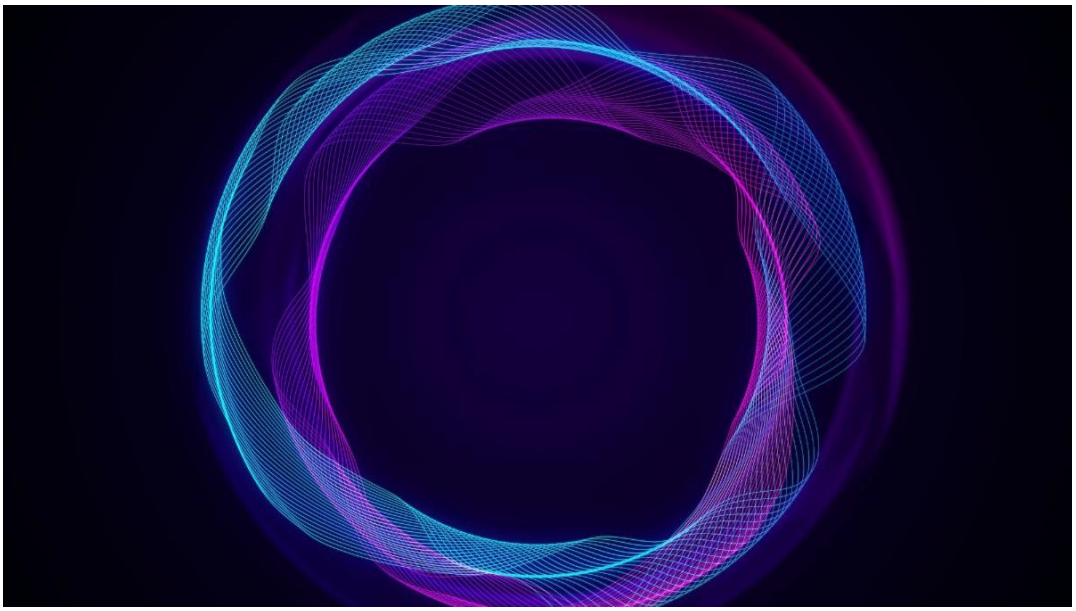

In the Lab

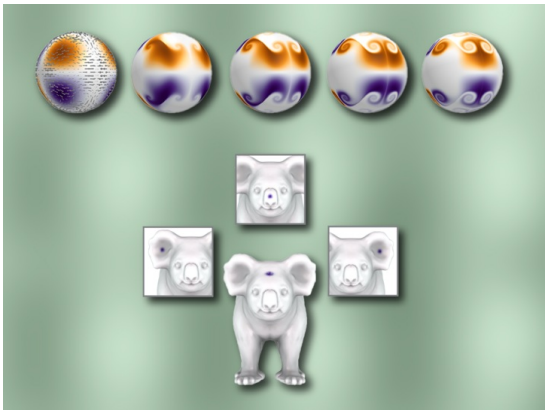

A framework for solving parabolic partial differential equations

[A new algorithm breaks them down into simpler problems.]

Generating phenomena like fire and flames with computer graphics and geometry processing requires solving complicated physics. A new framework from the team led by Lab researcher Justin Solomon achieves this for different nonlinear parabolic PDEs on triangle meshes by splitting them into three simpler equations.

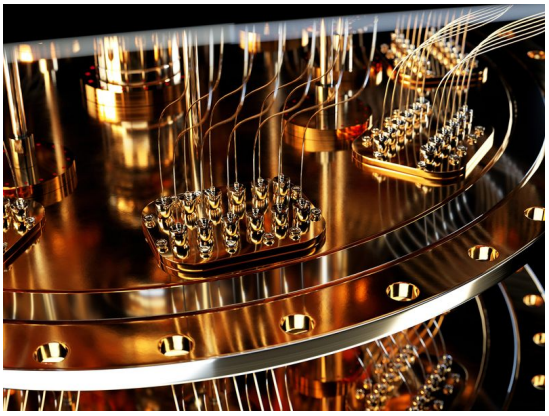

Toward a code-breaking quantum computer

[Making a smaller and more noise-tolerant quantum factoring circuit for cryptography]

Research out of the Lab group of Vinod Vaikuntanathan has proposed a new fast algorithm that requires fewer qubits and has a higher tolerance to quantum noise, which could make it more feasible to implement in practice.

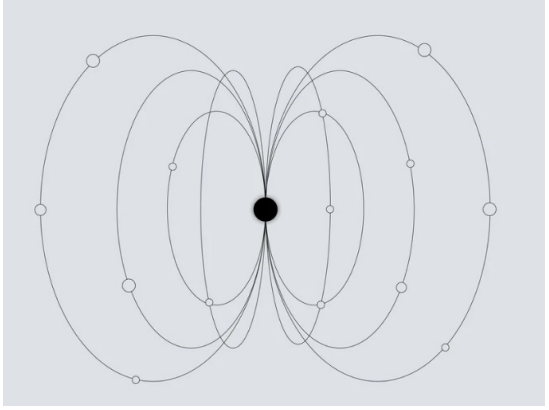

Teaching AI models to improve themselves

[How self-specialization and deductive closure training can improve the accuracy of language models without weighing them down.]

Large language models are skillful at many tasks but fact-checking themselves is not one of them. Addressing this, Lab researchers Yada Zhu, James Glass, Rogerio Feris, Leonid Karlinsky, and co-director David Cox, along with their colleagues developed a method that leverages in-context learning and synthetic data to enhance expertise.

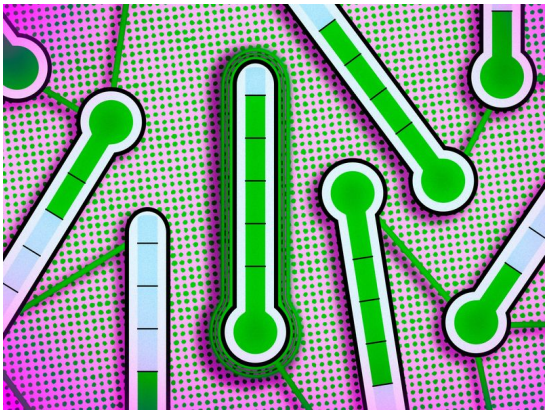

Method prevents an AI model from being overconfident about wrong answers

["Thermometer" technique could help someone know when they should trust an LLM.]

Since LLMs are used for a variety of tasks, the Lab groups of Subhro Das, Kristjan Greenewald, Prasanna Sattigeri, Gregory Wornell, Soumya Ghosh have developed an auxiliary model that runs on top of an LLM and helps to calibrate it.

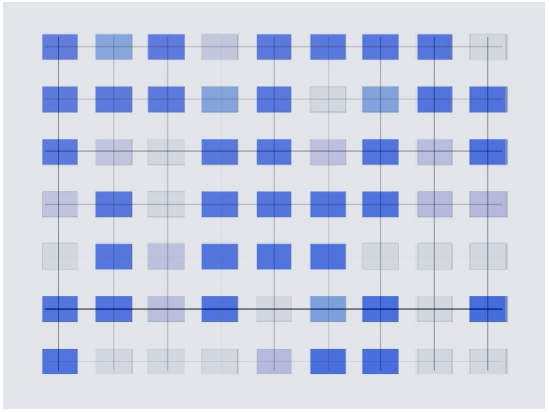

What's an LLM context window and why is it getting larger?

[Providing more background can lead to more coherent and relevant answers.]

Many strategies have been employed to extend context windows, leading to better LLM performance, including from the Lab research groups of Rameswar Panda and Yoon Kim, who used 500 million tokens and long form documents for pretraining. Additionally, Lab researchers Subhro Das, Dmitry Krotov, Leonid Karlinsky, Rogerio Feris, and co-director Aude Oliva developed a method that uses synthetic longform instruction data and compression.

Researchers advance automated interpretability in AI models

[Multimodal agent, MAIA, can iteratively design experiments to better understand AI components.]

Lab teams led by Jacob Andreas and Antonio Torralba developed a pre-trained vision-language model with a set of tools that can interrogate the subcomponents of other models to help explain their behavior.

AI model identifies certain breast tumor stages likely to progress to invasive cancer

[The model could help clinicians with assessment and ultimately help in reducing overtreatment.]

Research out of the Lab group of Caroline Uhler has created a machine-learning model can leverage chromatin staining to identify the type and stage of ductal carcinoma in situ (DCIS) — a task that's difficult for physicians.

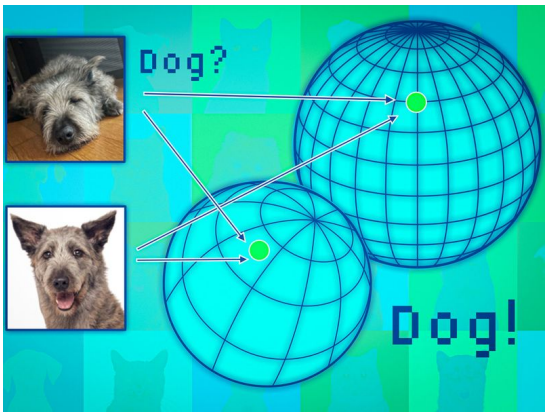

How to assess a general-purpose AI model's reliability before it's deployed

[Enabling users to compare several large models and choose the one that works best for their task.]

Not all foundation models are created equal, so a Lab team led by Navid Azizan and Hao Wang developed a method to quantify uncertainty by assessing the consistency of representations across an ensemble of similar models and discerning how each model would perform on on a variety of downstream classification tasks.

IBM reaffirms its commitment to the Rome Call for AI ethics

[IBM, other leaders and representatives came together to discuss developing AI ethically.]

Over the summer, leaders from major religions and industry organizations gathered in Hiroshima to commit to and reaffirm the use and development of AI for the benefit of humanity.

Foundations of Computer Vision

[Antonio Torralba, Phillip Isola, and William Freeman take a closer look at a changing field]

MIT Professor William Freeman and Lab researchers Antonio Torralba and Phillip Isola discuss their latest textbook, which explores the technology through the years, as well as its connection to machine learning and themes of ethics and fairness in the field.

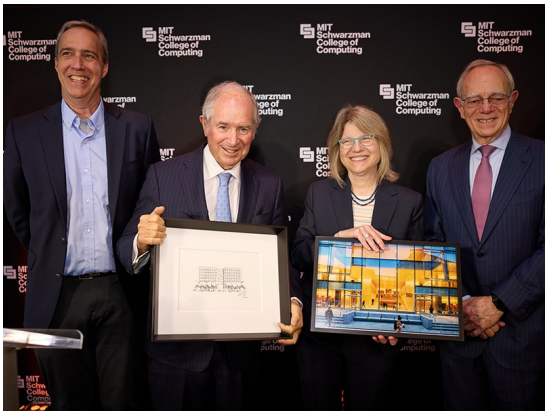

Marking a milestone

[Dedication ceremony celebrates the new MIT Schwarzman College of Computing building]

The MIT Stephen A. Schwarzman College of Computing celebrated the completion and inauguration of its new building with a dedication ceremony. “It fills me with immense satisfaction and pride to see the vibrant activity of the MIT students, researchers, faculty, and staff who spend time in this building,” said Dean Dan Huttenlocher, who is also the Lab's MIT co-chair. “It’s really amazing to see this building come to life and become a resource for so many across the MIT campus and beyond.”

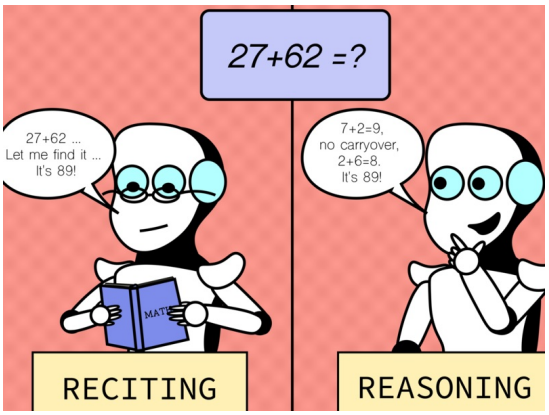

Reasoning skills of large language models are often overestimated

[How LLMs excel in familiar scenarios but struggle in novel ones]

On account of their build, it's difficult to examine the inner workings of LLMs, when it comes to discerning their memorization versus reasoning abilities. A Lab team led by Jacob Andreas and Yoon Kim probed this through an investigation around common tasks and unfamiliar counterfactual scenarios.
7/11

In the Media

A Greek-Indian friendship driven by innovation

Anantha Chandrakasan, MIT Lab chair, chief innovation and strategy officer, and dean of engineering, and Pavlos-Petros Sotiriadis PhD '02 sit down with Tassoula Eptakili of Kathimerini to discuss MIT’s unique approach to entrepreneurship, the future of AI and the importance of mentorship. MIT’s approach to education “focuses both on foundational knowledge and its practical application. This prepares [students] to tackle even the most challenging problems,” says Chandrakasan.

IBM wants to teach AI the language of your business

“Open innovation is really the story of human progress,” Lab co-director David Cox said at VB Transform 2024, as reported by VentureBeat, making the case for openness in innovation for enterprise generative AI.

Lab Highlights

Fotini Christia named director of the MIT Institute for Data, Systems, and Society (IDSS).

Lab researcher Hilde Kuehne has joined the faculty of the Universität Tübingen in Germany.

Lab co-director Aude Oliva received the 2023 INNS Donald O. Hebb Award from the International Neural Network Society, recognizing outstanding contributions to research in biological learning.

Lab researcher Armando Solar-Lezama received the Robin Milner Young Researcher Award for his research, which has "reignited" program synthesis approaches. "Bridging also into machine learning, Armando has brought program synthesis to enable neuro-symbolic computation with better interpretability, verifiability, generalization, counterfactual reasoning, and structuring computation search spaces for more effective inference."

Lab researcher Manish Raghavan received a Google Research Scholar Award, awarded to professors based on merit to support their cutting-edge research.

Lab researcher Alex Shalek named director of the Institute for Medical Engineering and Science (IMES) at MIT.

Lab researcher Julian Shun received the 2023 Association for Computing Machinery Paris Kanellakis Theory and Practice Award, which honors specific theoretical accomplishments that have had a significant and demonstrable effect on the practice of computing.

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

The MIT Schwarzman College of Computing and the MIT Washington Office are leading an effort to produce policy briefs with recommendations on the governance of AI. The goal of these briefs is to help shape a technically informed discussion of how to govern AI in a way that will make it safe while enabling AI to thrive. Dan Huttenlocher, MIT Lab co-chair and dean of the college, along with Lab researchers Yoon Kim and Jacob Andreas have contributed insights.

Over 100 students attended the Lab's Open House this past month, showcasing projects and demos and facilitating recruitment for summer 2025 internship positions and 6A MEng opportunities.

Online Learning

Artificial Intelligence: Implications for Business Strategy
A joint MIT CSAIL and MIT Sloan School of Management Course begins
October 9.

Driving Innovation with Generative AI
An MIT xPRO Course begins
October 15.

Machine Learning in Business
A joint MIT CSAIL and MIT Sloan School of Management Course begins
October 16.

Unsupervised Machine Learning: Unlocking the Potential of Data
A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
November 13.

Making AI Work: Machine Intelligence for Business and Society
A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins
November 20.