Computer Science and Artificial Intelligence Laboratory

CSAIL pioneers research in computing that improves the way people work, play, and learn. CSAIL serves the MIT community, the country, and society at large by creating a positive future enhanced by computer science through contributions of ideas, artifacts, and people.

CSAIL leadership spearheaded several new initiatives in FY2023:

- AI Accelerator training program for Dept. of Defense and national security leaders, Spring 2022
- Negotiation Skills for Women in AI symposium, April 2023
- TEDxMIT Superpowers, April 2023
- Weekly CSAIL Generative AI discussions: established in May, 2023 discussions about generative AI and research directions, exploring how we can use generative AI to address issues such as climate change.
- CSAIL-LIDS Joint Machine Learning Advances Symposium, May 2023
- CSAIL Ten Year Impact Report was completed and slated for publication.
- CSAIL celebrated 60th anniversary with a reunion event and a one-day Computing the Future symposium marking Project MAC's 60th anniversary, the CSAIL's 20th anniversary, and key role in the advancement of computer science.
- MIT CSAIL + Imagination in Action: AI Frontiers & Implications – 12-hour program with media highlights by Forbes, featured 60 talks featuring CSAIL researchers, 40 startups, 1000+ attendees.

Growth

FY2023 total combined research volume (Primary and Secondary funds) was $96,522,419. Research volume was 65% federal and 35% non-federal. CSAIL managed 633 active research accounts and over 112 PIs with appointments across 11 MIT departments. Through 2022-2023 academic year we had 615 graduate students with RA appointments, and 317 UROP students.

New funding sources include USGA (US government sponsor) and the following US and international non-federal sponsors:

- Foreign Federal Government
  - Swiss National Science Foundation
- Foreign Foundation
  - Hasso Plattner Foundation
- Foreign Private Profit
  - Ericsson
  - Hyundai Motor Company R&D Center
  - Hyundai NGV Company Ltd

- Private Non-Profit
  - Eric and Wendy Schmidt Fund for Strategic Innovation
- Private Profit
  - Boston Dynamics AI Institute
  - Sanofi US.

Other organizations sponsoring research include:

- Institution of Higher Education
  - University of Illinois-Urbana Champaign
- Private Non-Profit
  - The Research Foundation - Stony Brook University
- Private Profit
  - Aurora Flight Sciences RDC
  - Lawrence Livermore National Security, LLC.

Research Initiatives

Capital One Research Program

Established to support seeding machine learning and data-centric research with potential applications to the financial service industry such as: exploring challenges related to systems that support ML at scale, data and privacy, model training and efficiency, and/or explainability of ML models. Six funded projects received funds to support the work of 1 PhD student for 1 year.

METEOR Program

Five Postdocs: Amy Fox (Advisor: Satyanarayan), Luana Ruiz (Advisor: Jegelka), Elizabeth Bondi-Kelly (Advisors: Sontag, Ghassemi), Kwesi Rutledge (Advisors: Roy, Fan), and Vasha DuTell (Advisors: Rosenholtz, Freeman) Mackenzie Leake is going to Adobe and has an offer from Brown. We received 22 applicants for year three. Some may be shared with FODSI, SoE, SERC. Funding is provided by industrial sponsors with the focus on scientific excellence. Eligibility criteria: US citizens, green card holders or DACA recipients, broad eligibility otherwise (race, ethnicity, gender, 1st generation, significant hardship).

Research Highlights

- Daniela Rus's lab created "VISTA 2.0," a data-driven simulation engine where vehicles can learn to drive in the real world and recover from near-crash scenarios. All of the code is being open-sourced to the public.
- Pulkit Agrawal's lab ensures AI works with the right dose of curiosity. Researchers make headway in solving a longstanding problem of balancing curious "exploration" versus "exploitation" of known pathways in reinforcement learning.

- Hari Balakrishnan has made roads safer for drivers.
- Stefanie Jegelka seeks to understand how machine-learning models behave, to help researchers build more robust models for applications in biology, computer vision, optimization, and more. Unpacking the "black box" to build better AI models.
- Manolis Kellis' lab charts how exercise affects the body. A new study maps the genes and cellular pathways that contribute to exercise-induced weight loss.
- Aleksander Madry urged lawmakers to ask rigorous questions about how AI tools are being used by corporations: "We are at an inflection point" with AI"
- Ted Adelson's group developed a robotic hand that uses high-resolution touch sensing to accurately identify an object after grasping it just one time.

Department of the Air Force-MIT AI Accelerator

Since 2019, the projects of the Department of the Air Force (DAF)-MIT AI Accelerator (aia.mit.edu) have advanced AI research in a broad range of areas, including weather modeling and visualization, optimization of training schedules, and enhancement of autonomy for augmenting and amplifying human decision-making. The AI Accelerator engages more than 150 faculty, researchers, and students, who are affiliated with more than 20 different organizational units across MIT Campus and MIT Lincoln Laboratory, including the School of Architecture and Planning, the School of Engineering, the School of Science, the School of Humanities, Arts, and Social Sciences, and the Stephen A. Schwarzman College of Computing. The interdisciplinary AI Accelerator project teams also include DAF personnel, who are embedded in the research teams and serve as liaisons between the projects and Department of Defense stakeholders. There are currently 15 active AI Accelerator projects that involve 35 MIT faculty and also include recently launched and expanded projects, namely, Automation in Space Domain Awareness, Better Networks via AI Enabled Hierarchical Connection Science, Few-Shot and Continual Learning, and Trustworthy AI.

Defense Science and Technology Agency Singapore

Research collaboration with DSTA continues its third year. Projects started in 2021:

- SYNTHBOX: Establishing Real-World Model Robustness and Explainability Using Synthetic Environments (Aleksander Madry).
- Next Generation NLP Technologies for Low Resource Tasks (Regina Barzilay, Tommi Jaakkola)
- Trustworthy, Deployable 3D Scene Perception via Neuro-symbolic Probabilistic Programs (Vikash Mansinghka, Joshua Tenenbaum)
- Computationally-Supported Roleplaying for Social Perspective Taking (D. Fox Harrell)
- Analytics-Guided Communications to Counteract Filter Bubbles and Echo Chambers (Deb Roy)
- Decentralized Learning with Diverse Data (Costis Daskalakis, Asu Ozdaglar, Russ Tedrake)

- Data-driven Optimization under Categorical Uncertainty, and Applications to Smart City Operations (Alexandre Jacquillat)
- Provably Robust Reinforcement Learning (Ankur Moitra)
- Building Dependable Autonomous Systems Through Learning Certified Decisions and Control (Chuchu Fan)
- Online Learning and Decision Making under Uncertainty in Complex Environments (Patrick Jaillet)
- New Representations for Vision (William Freeman, Joshua Tenenbaum)
- Improving Situational Awareness for Collaborative Human-Machine First Responder Teams (Nicholas Roy) (concluded in 2022)
- Low Resource Multilingual Speech Recognition (James Glass)
- Sparse Data Methods for Speech Recognition (James Glass)

One additional project was launched in January, 2023:

- 3D Fusion from Multiple Images (Sertac Karaman) This project is focused on developing new algorithms that allow for the construction of 3D views from multiple viewpoints. Specifically, the project aims to investigate the reconstruction of multiple viewpoints through few images or videos from very different angles and scales.

Gwangju Institute of Science and Technology

CSAIL initiated the following research projects, in collaboration with the Gwangju Institute of Science and Technology (GIST) in South Korea. The projects vary in length from 24 to 27 months, and continued from 2021:

- AI for Human Computer Interaction in Education (Wojciech Matusik, Daniela Rus)
- Artificial Compound Eye with Artificial Intelligence (ACE.A.I) for Enhanced Sensing (Fredo Durand, William T Freeman)
- Extending Contrastive Learning to New Data Modalities and Resource-Limited Devices (Dina Katabi, Piotr Indyk)
- AI for Energy: Designing high-performance catalysts and electrodes for efficient hydrogen production (Tommi Jaakkola, Regina Barzilay)
- AI-Driven Soft Robot Skin for Recognition, Modeling, and Exploration (Stefanie Mueller, Daniela Rus)
- AI-driven discovery of co-evolutions of host-microbiome interactions for the application of microbiome-based therapeutics (Yoon Kim, Marzyeh Ghassemi)

Research Collabs Program - Google

In this joint research program, a Googler, MIT faculty member and MIT PhD student join in a “triad” to collaborate on a one-year joint research project through a sponsored research agreement. Current projects include:

Discovery of super ion conductors for clean energy, enabled by machine learning and targeted experiments (Bilge Yildiz)

MIT collaborator will discover inorganic solid materials that enable very fast conduction of small ions, such as proton and lithium ion, based on physical principles. MIT collaborators will leverage machine learning techniques and targeted validation experiments. The expected results of this research are important for enabling safe and high energy density batteries for energy storage, as well as fast and energy efficient analog computing devices.

Improving Conditional Computation in Neural Networks using (Combinatorial) Optimization (Rahul Mazumder)

This project will determine robustness properties via large-scale experiments that will extend the scope of MIT collaborator's approach to many applications like distillation, increasing robustness of existing pre-trained models that are too expensive to retrain, meta learning and few-shot learning etc. MIT collaborators will research using combinatorial optimization to improve sparsity for tree-based gate to solve expert selection problem to (near) optimality.

Extracting 3D Knowledge from Large Generative Models (Antonio Torralba)

MIT Collaborators will train generative models at scale using uncurated image collections that are easily acquired from the Internet. MIT Collaborators will extract 3D information from large generative models, and incorporate generative training to improve 3D inference and inverse rendering. Google will provide feedback to the MIT Collaborators and MIT Collaborators will develop code and perform experiments to verify the models. Googler, Faculty, and MIT Collaborators will jointly analyze results, and produce joint publications. Code from this research will be open-sourced.

Symbolic Distillation of Large Language Models for Trustworthy NL (Yoon Kim)

MIT collaborators will extract a set of task-specific rules from large language models via structured query inputs, and will use the extracted rules as part of an interpretable symbolic model (e.g., based on probabilistic graphical models). Spurious correlations are eliminated by inspecting and refining extracted rules and further develop a more trustworthy system that is better aligned with expert behavior. MIT collaborators will achieve a level of controllability currently unavailable when directly working with the underlying language model.

Project that Concluded June 2023

- Indoor Localization using one-sided round-trip time (RTT) (Berthold Horn)

Projects that Concluded End 2022

- Understanding and Improving Attention Mechanisms in Transformer-based Deep Networks (Philip Isola)
- Improving Conditional Computation in Neural Networks using (Combinatorial) Optimization (Rahul Mazumder) (Note: The initial 2022 project)
- Non-Gaussian Data Assimilation (Yousef Marzouk)

Quanta - CSAIL Research Collaboration

CSAIL is continuing its five-year collaboration. The primary focus of the collaboration is largely defined by the troika of patient, hospital, and doctors; other aspects include privacy/security, and nutrition and self-care. Current projects include:

- Using Machine Learning to Curb Infectious Disease-John Guttag
- Learning to Assess Breast Cancer and Lung Cancer Risk to Enable Early Detection and Prevention-Regina Barzilay
- Revolutionizing the Care of Patients with Cardiovascular Disease- Collin Stultz Deep
- Metric Learning to Uncover Diabetic Patient Subtypes and Personalize Treatments - Marzyeh Ghassemi

Toyota-CSAIL Joint Research Center

CSAIL continues its projects with Toyota Research Institute in University 2.0, a long-term collaboration with other universities, in its 3rd year (funded through April 2024). Projects include:

- Task Driven Development of Nimble, Reactive, Rugged Hands (PIs: Ted Adelson, Sangbae Kim, Alberto Rodriguez, Wojciech Matusik, Pulkit Agrawal)
- How can we create superintelligent human-computer groups? (PIs: Tom Malone, Daniela Rus, Abdullah Almaatouq)
- Learning Interactive and Responsive Driving (PIs: Daniela Rus, Sertac Karaman)
- Mini City Road Challenge (PIs: Daniela Rus, Evangelos Theodorou [GA Tech], James Rehg [GA Tech], Sertac Karaman)
- Scalable Self-Supervised Learning for 3D Scene Understanding (PIs: Justin Solomon, Greg Shakhnarovich [TTIC])
- Physical and Functional Inductive Biases for Visual Representation Learning (PIs: Josh Tenenbaum, Jiajun Wu [Stanford], Fredo Durand)

Wistron-CSAIL Research Collaboration

Current projects include:

- Multi-Modal Sensing Systems for Smart Manipulation and Human-Environment Interactions (PIs: Antonio Torralba and Wojciech Matusik)
- A Hardware Accelerator to Enable Computing on Encrypted Data (PIs: Daniel Sanchez and Srini Devadas)
- Using Test Data Augmentation to Improve Classification of Data (PI: John Guttag)
- Multimodal AI: Learning Joint Representations of Heterogeneous Data (PIs: Polina Golland and Peter Szolovits)

Industrial Outreach

Alliances

The CSAIL Alliance Program (CAP) supports the mission of CSAIL by connecting researchers, students and technological advances to industry and organizations across the globe in three levels: (1) Student Engagement – focused on connecting with students and post docs for career opportunities (2) Affiliate, which provides lab visits, access to the annual meeting, recruiting assistance, research briefings and professional education discounts and (3) Partner, which includes expanded options with added access to research initiative meetings, custom faculty-led seminars and expanded recruiting options. Currently there are over 100 member companies, including CSAIL start-ups and global brands such as Apple, Cisco, Dell EMC, Google, Adobe, Citibank, and Microsoft.

CSAIL Alliances also produces and manages online professional development courses in partnership with MIT's Professional Education, MIT's Office of Digital Learning (ODL), Get Smarter, Simplilearn and Emeritus. We have repurposed our content to reach global audiences through partners in Japan (CyberU) and India (Simplilearn) as well as connecting courses as a single offering in US markets (Harvard Extension School). Previous courses in partnership with edX ran 2014-2018 and enrollments totaled 18,978 for total enrollment in excess of 46,000 online learners.

In addition, CSAIL Alliances also runs several research focused initiatives. Initiatives are pre-competitive, are formed with a small group of companies around a theme, seed projects aligned with the theme and include all the benefits of the Affiliate level. Presently we have the following active initiatives:

- FinTech@CSAIL's final remaining founding members made 2 awards PIs in the Spring of 2023. There is interest from both industry and faculty to restructure this initiative and the process for restructuring has begun.
- MachineLearningApplications@CSAIL focuses on applications of the latest machine learning (ML) technologies and research on the resolution of current challenges limiting the abilities of ML and professional development that will help prepare a company's workforce for this digital transformation. This initiative is led by Daniela Rus. Current members include: Arrow Electronics, BT, Cisco, CapitalOne, SAP and Ernst & Young. In FY23, the initiative provided discretionary funding for 7 proposals and supported 7 CSAIL PIs.
- Future of Data, Trust and Privacy launched April 2021 in collaboration with MIT Internet Policy Research Initiative (IPRI), co-led by Daniel Weitzner and Professor Srini Devadas. Weitzner's research pioneered the design of accountable systems as a new approach to privacy and investigated the interaction between cryptographic technology and surveillance law, while Devadas' research areas are computer architecture, computer security, and applied cryptography. Current members include: Capital One, Mass Mutual Insurance, Fidelity, and Visa. The initiative provided discretionary funding for 6 proposals and supported 7 CSAIL PIs. The initiatives also produced a white paper on accountability and data traceability https://www.csail.mit.edu/news/mit-future-data-publishes-new-white-paper-accountability-and-traceability.

Internet Policy Research Initiative

The mission of the Internet Policy Research Initiative (IPRI), an institute-wide initiative, is to work with policymakers and technologists to bridge the gap between technologists and policymakers and increase the trustworthiness and effectiveness of interconnected digital systems, like the Internet. We accomplish this via a three-pronged approach: targeted engineering and public policy research, educational programs for students and policymakers, and outreach programs to build policy communities that facilitate communication. IPRI's research efforts cover six categories:

Cybersecurity

IPRI's cybersecurity research focuses on the technical and policy aspects of cybersecurity issues as they relate to the communication networks and software systems affecting the global society and economy. This multidisciplinary research area encompasses encryption policy, accountability, cryptography, data sharing, securing core economic and social infrastructure, measuring cyber risk, and more. Major projects include an interdisciplinary project called SCRAM (Secure Cyber Risk Aggregation and Measurement) that uses cryptographic computation tools to safely aggregate sensitive security data and develop cyber risk metrics and models. We have ongoing research projects with the US Federal Reserve System (including an annual meeting on measuring cyber risk), municipalities, and security communities such as ISACs and ISAOs.

The MIT Future of Data Initiative

This is a leading multi-disciplinary research agenda to design and stimulate the deployment of accountable systems to provide trusted, traceable uses of personal data on an ecosystem-wide scale. The Initiative gathers computer science and Internet policy researchers with leading commercial enterprises in financial services, payment technology, cloud platforms, insurance and other sectors to discuss current challenges and opportunities in privacy and data governance. Modern privacy laws place appropriately high expectations on organizations processing personal data. At the same time, users report declining trust in those who handle their personal data and regulators around the world struggle with the scale of the enforcement challenge. A key challenge addressed by today's roundtable is to identify and put into service technical infrastructure for enterprises seeking to handle personal data in a trustworthy and lawful manner, and with guardrails to enable the scalable use of that data.

Privacy

IPRI's work also focuses on privacy policy and its critical role in trustworthiness. The Privacy group has published work on topics like privacy and security in home assistants, exposure minimization, and the data-sharing practices of smartphone apps. Current projects include the development of privacy-aware databases.

Networks

The Advanced Network Architecture (ANA) group is organized around five themes: Internet architecture, Internet security, Internet economics, Internet policy, and network management.

AI Policy

Current research areas include studying the role of AI in financial decision making, increasing access to new training data sets with policy, working with stakeholders on AI principles, and shaping global Internet policymaking via policymaker engagement and informing the public debate.

The Decentralized Information Group

Focuses on data and systems governance (primarily on the Web) and explores both policy and technical issues. Current projects include a decentralized privacy-preserving platform for clinical research, evaluating the trustworthiness of autonomous systems, studying the relationship between privacy and machine learning, developing explanations for complex machines and models, securely aggregating distributed data, and developing smart contracts for data sharing.

MIT App Inventor

Empowers young people to develop useful apps that serve as novel digital solutions to problems they face in their lives, communities, and world.

Research Highlights

Numerous individual and multi-investigator projects are under way with major research discoveries across all areas of computing. Our work opens many opportunities to propel science, create new businesses, protect the planet, understand life, improve our cities and enhance our well-being and quality of life. A sampling of the work is highlighted below.

Acquiring multiple complex robot skills for challenging conditions (Pulkit Agrawal)

A goal of robotics is to develop robots that can reliably perform tasks in diverse circumstances that are unknown apriori. Some examples of where such capability is critical are robots deployed in disaster relief, search or rescue missions, or assisting the elderly with household chores. However, most of the robots that exist today are designed to perform a single task in a narrow range of environmental conditions and therefore cannot be deployed in general scenarios. The primary reason for their deficiency is: (a) errors in estimating task-relevant system parameters (such as friction coefficient or object mass) from sensory data and (b) the laborious search for task-specific heuristics required to construct controllers. Moreover, these heuristics only operate in limited conditions, requiring humans to analyze failure modes and propose new heuristics for new scenarios. Consequently, building robot systems that can perform even a single task in diverse environmental conditions requires substantial human intuition, effort and thereby remains a formidable challenge.

We use machine learning methods to build a scalable way for acquiring multiple complex robot skills that operate in challenging conditions. The main idea is to leverage large amounts of simulated data to learn high-reward action sequences (or skills) valid in diverse conditions directly from sensory observations. These skills are then transferred to the real world. The developed approach only requires the human to specify "what" task the robot should do and therefore mitigates human effort in defining "how" to do the task (i.e., coming up with task-specific heuristics or estimating task-relevant system parameters). Because the robot can automatically learn skills given only the task specification, this framework provides a scalable way of learning a large number of skills.

We applied these ideas to the well-known challenging problem of re-orienting objects using a human-like robotic hand, an essential step in robotic tool use. This work demonstrated the ability to re-orient many previously unseen objects for the first time and won the Best Paper Award at the Conference on Robot Learning (CoRL). The same approach was applied to enable the MIT Mini Cheetah (a four-legged robot) to jump over gaps and run across diverse terrains at speeds faster than prior methods tested on this robot. The robot was robust to unexpected changes such as motor breakdown during deployment. The results from this work were widely covered in the press.

Leveraging Multi-Material 3D printing and Computational Design to Enable Novel Product Applications (Stefanie Muller)

Most of today's objects that product designers create have only a single appearance, i.e. one static surface texture. To assist designers in creating visually dynamic objects, such as a kettle bell that can provide the viewer with instructions to raise or lower it based on the current exercise position, we developed Lenticular Objects. Published in the premier Human-Computer Interaction conference ACM UIST 2021, Lenticular Objects use 3D printed lenticular lenses to create objects that display different patterns from different viewpoints. For this project, we collaborated with MIT Architecture and the Harvard Graduate School of Design and received the Red Dot Design Concept Award and IF Design Award for the design contributions of this project. Creating objects with viewpoint-dependent design requires expert knowledge in optics, computer graphics and 3D modeling. To make this process more accessible, the design tool we developed supports designers in generating a viewpoint-dependent design by taking as input the desired surface patterns for each of the different viewpoints. With our design tool, the designer can ideate by trying out different viewpoints and visual textures and evaluate the result by previewing it with the simulation tool. When the design is ready, our system exports the fabrication file, which can be directly used for 3D printing the object in one pass.

Private Web Search (Henry Corrigan-Gibbs)

To use a web-search engine today, a client must send its search query directly to the search provider. This is a privacy risk: a client's search queries contain sensitive data, including its location ("Boston weather"), health conditions ("lung cancer"), and personal beliefs ("recipes for Passover seder"). Even a well-intending search engine may disclose users' sensitive query data through a data breach or seizure (e.g., in an authoritarian state).

We have built the first web-search engine that allows a client to search the web without revealing _any_ information about its query to the search provider. To perform a search, the client first uses a special type of encryption scheme and sends the encrypted query up to the search provider's servers. The provider can transform the client's encrypted query into an encrypted set of search results without ever seeing the client's query in unencrypted form. To make this possible, we developed a new high-speed encryption scheme that supports performing computations on encrypted data.

We call our new search engine "Tiptoe," and have built a prototype of it that allows clients to search over 364 million webpages in under three seconds. The search quality cannot match Google yet, but our scheme is competitive with traditional text-search algorithms. With this project, we hope to give everyday users access to the breadth of the web while protecting them from some of its most severe privacy risks.

This project was done in collaboration with Alexandra Henzinger, Emma Dauterman (UC Berkeley), and Nickolai Zeldovich.

Addressing Challenges of Designing and Evaluating Systems for Health Deployments (Marzyeh Ghassemi)

Good model performance is no guarantee of correct model use, and there are many decision-making dynamics that should be considered early on in the model development process. Model bias and subgroup underperformance are not currently regulated by the FDA. To explore what impact biased models might have in deployed health settings, the Healthy ML Lab led recent work to evaluate biased advice in a mental health crisis setting [1]. Specifically, we presented advice from intentionally-biased LLMs to subjects who are deciding whether to call for medical or police help. We instructed subjects to call for police help if, and only if, there was a risk of violence, and found that subjects at baseline were not more likely to call police on Black and Muslim patients. When our biased LLMs gave prescriptive advice to subjects, e.g., "our model thinks you should call the police", clinicians and non-clinicians tended to follow the advice and call the police on minority patients. However, we also found that the biased LLM advice was mitigated by delivery method, specifically by presenting the advice descriptively, e.g., "our model flagged a risk of violence," rather than prescriptively. While both delivery methods present the same biased predictions, there is a significant difference in the human outcome.

We also investigated the impact of standard data collection methods being used in benchmark machine learning systems. Importantly, my lab has demonstrated that human labelers label data differently depending on whether they are in a factual or normative setting [2]. E.g., subjects label the presence of fried food differently when asked a) whether a meal has fried food, versus b) whether a meal has fried food as part of a meal policy that prohibits fried food. Importantly, we demonstrated that training a model on factual data creates machine learning systems that do not reproduce human judgements in normative systems. This is an important warning for a field where datasets are often used without close examination of labeling practices, and underscores the need for caution in automated decision systems, particularly in contexts where compliance with societal rules is essential.

Laboratory Sponsored Activities

Media Outreach

Consistently higher media attention and output than competitors.

- MIT CSAIL has 1574 press mentions this year (CMU SCS received 411 and Stanford HAI received 130).
- Combined following of 543,451 users across Twitter, Instagram, Facebook, LinkedIn, and YouTube (31% growth from last year).
- Consistent media coverage in various top-tier outlets: Forbes, WIRED, TechCrunch, IEEE Spectrum, Boston Globe, Popular Science, and Washington Post.

CSAIL's Twitter following has grown over 2000% since 2015, well above UCB (roughly 1500%), Stanford Engineering (over 250%), MIT Engineering (roughly 1000%), MIT EECS (about 500%), and CMU CS (under 250%).

YouTube Channel

Channel	Views	Videos	Average
MIT CSAIL	168,400	23	7,322
Harvard SEAS	141,198	20	7,060
Stanford HAI	86,398	58	1,490
CMU SCS	26,496	31	855

Highlights

- Drones navigate unseen environments with liquid neural networks: 42 media mentions
- AI models fail to reproduce human judgements about rule violations: 39 media mentions
- Robotic hand can identify objects with just one grasp: 31 media mentions

CSAIL Hosted Lecture Series

Dertouzos Distinguished Lectures have been a tradition since 1976, featuring some of the most influential thinkers in computer science. Speakers featured during 2022-2023:

- Prof. Fei-Fei Li, What we see and what we value: AI with a human perspective, November 2, 2022
- Prof. Pat Hanrahan, Shading Languages and the Emergence of Programmable Graphics Systems, December 7, 2022
- Prof. Kate Crawford, Generative Machines and Ground Truth, February 15, 2023
- Prof. Dorit Aharanov, The Search for Evidence of Quantum Advantage, March 15, 2023

Hot Topics in Computing was launched in 2017, convening experts to discuss emergent potential, perception, and problems associated with the proliferation of computation and machines. Topics featured during 2022-2023:

- Dr. Patrick Dykstra, Patrick and the Whale, December 6, 2022
- Dr. Sebastien Bubeck, First Contact, March 22, 2023
- CSAIL PI Panel discussion on Generative AI, March 23, 2023

Organizational Changes

- Hired 7 fiscal staff, 5 HR staff, and 4 AAs
- Shepherded the successful W3C spinout as of December, 31, 2022
- Developed and implemented CSAIL research scientists review plan

New faculty started during 2022–2023 include:

Sara Beery, Assistant Professor.

Christina Delimitrou, Assistant Professor

Mohsen Ghaffari, Assistant Professor

Mina Konakovic Lukovic, Assistant Professor

Vincent Sitzmann, Assistant Professor

Faculty taking leave during 2022–2023 included:

Costis Daskalakis - Sabbatical

Erik Demaine - Sabbatical

John Guttag - Sabbatical

Piotr Indyk - Sabbatical

Leslie Kaelbling - Sabbatical

Ruth Rosenholtz - Leave (NVIDIA)

Howard Shrobe - DARPA IPA

Nicholas Roy - Sabbatical

Julie Shah - Sabbatical

David Sontag – Sabbatical

Jack Wisdom - Sabbatical

Awards and Honors

- Alizadeh, Mohammad, 2023 Association for Computing Machinery (ACM) Grace Hopper Award
- Andreas, Jacob, 2023 Junior Bose Award for Excellence in Teaching
- Balakrishnan, Hari, 2023 Marconi Prize
- Barzilay, Regina, 2023 inducted into National Academy of Engineering (NAE)
- Brooks, Rodney, 2023 IEEE Founders Medal
- Daskalakis, Costis, 2023 ACM Fellow
- Ghobadi, Manya, 2023 ACM-W Rising Star Award
- Goldwasser, Shafi, elected 2023 Foreign Member of the Royal Society, UK National Academy of Sciences
- Hadfield-Menell, Dylan, 2023 Schmidt Futures AI2050 Early Career Fellow
- Indyk, Piotr, 2023 elected to American Academy of Arts & Sciences
- Kalai, Yael, 2022 ACM Prize in Computing
- Katabi, Dina, 2023 elected to National Academy of Sciences (NAS)
- Leighton, Tom, 2023 IEEE John von Neumann Medal
- Liskov, Barbara, 2023 Benjamin Franklin Medal in Computer and Cognitive Science
- Metcalfe, Robert, 2022 ACM A.M. Turing Award
- Mueller, Stefanie, 2022 MIT Technology Review's "Innovators Under 35"
- Ragan-Kelley, Jonathan, 2023 Sloan Foundation Sloan Research Fellow

- Rubinfeld, Ronitt, 2023 Guggenheim Fellow
- Rus, Daniela, 2022 IEEE Robotics and Automation Technical Field Award
- Rus, Daniela, 2023 Boston Globe Tech Power Players
- Shor, Peter, 2022 Breakthrough Prize in Fundamental Physics
- Solomon, Justin, 2023 Harold E. Edgerton Faculty Achievement Award
- Tedrake, Russ, 2023 MIT SoE Teaching with Digital Technology Award
- Yan, Mengjia, 2022 Intel Rising Star Faculty Award

CSAIL members recognized for excellence in service and research contributions:

- Sponsored by CSAIL faculty lead Martin Rinard since 1997, MIT student team won the 2023 World Finals of the International Collegiate Programming Contest (ICPC)
- Sponsored by CSAIL, TEDxMIT won the 2023 MIT Student Awards for Outstanding Event
- Infinite Mile Awards: Jackie Cen, Senior Financial Officer – unsung hero; and Mary McDavitt, Administrative Assistant – Excellence
- Collier Medal, Martin Nisser, graduate student
- 31 Spot Awards achieved by CSAIL staff
- 27 Gratitude Book Club Awards to faculty and staff to recognize exceptional service

Key Statistics for Academic Year 2022–2023

Headcount	Women	Men	Women %
Faculty	25	102	19.69%
Postdoc Assoc/Fellow	24	73	27.84%
Principal Research Scientist	4	3	57.14%
Research Staff	15	64	18.99%
Senior Research Scientist	2	5	28.57%
Administration, technical, support staff	63	81	60.58%
Graduate Students	177	438	28.78%
UROP	139	178	43.85%
Visitors	12	43	21.80%
TOTAL PERSONNE: 1,497	**482**	**1,015**	**31.84%**

*Please note, 17 or 1.12% of headcount have unknown genders.

Daniela Rus
Director