# MIT Quest for Intelligence

The nature of intelligence—how the brain produces intelligent behavior and how it can be replicated in machines—is one of the fundamental problems in science and technology. Progress in this area could have enormous societal impact, making it easier to solve challenges in many other disciplines. We believe the key to progress is to integrate the science and engineering of intelligence. Understanding how the human mind works in engineering terms will lead to transformative advances in artificial intelligence (AI), allowing us to create machines that exhibit human-level intelligence in how they reason, see, communicate, and learn. An understanding of intelligence will also advance the fields of cognitive science and neuroscience and help us explain how the brain functions in both health and disease. We hope to translate this knowledge into a broad spectrum of societal benefits tied to health care, information security, transportation, global communication, education, and more.

The MIT Quest for Intelligence (MIT Quest) seeks to drive the development of transformative tools and technologies to benefit society. MIT Quest brings together researchers from across disciplines to discover the foundations of intelligence and to address the ethical and societal impacts of automated decision making.

## Goals, Objectives, and Priorities

MIT Quest is composed of two parts: The Core, intended to advance the science and engineering of intelligence, and The Bridge, intended to develop tools and technologies with positive societal impact.

In the past year, The Core has formulated a set of team-driven "moon shots" to tackle the most enduring questions around intelligence. Moon shots are long-term research projects focused on issues such as visual and linguistic intelligence (a commonsense understanding of events and language), the development of intelligence (how a human infant learns), embodied intelligence (allowing robots to behave autonomously and with a sense of morality), and collective intelligence (allowing humans and machines to work collaboratively to achieve what neither can do alone).

The Bridge is developing a set of "missions" aimed at democratizing AI in education, business, health care, and across society more broadly. Project-based missions are designed to produce training and technologies that address specific business and societal challenges. Our ultimate aim is to move from projects that need a dedicated AI specialist and domain expert to an era in which domain experts can intelligently access and deploy tools to solve their problem.

MIT Quest is committed to translating its research into tangible, real-world applications. In our first year, we formed several corporate and philanthropic collaborations to maximize our impact. These collaborations are described below.

### MIT-IBM Watson AI Lab

The MIT-IBM Watson AI Lab aims to advance AI hardware, software, and algorithms related to deep learning and reasoning; increase AI's impact on industries such as health care and cybersecurity; and explore the economic and ethical implications of artificial intelligence for society. Co-led by Antonio Torralba (inaugural director of MIT Quest) and David Cox (IBM Research) and chaired by Anantha Chandrakasan (dean of the MIT School of Engineering) and Dario Gil (vice president of AI and IBM Q at IBM Research), the lab will invest $240 million in AI research and education over the next 10 years. In two years, the lab has received 253 submission proposals and awarded 83 grants to 23 departments, labs, and centers across campus. The grants range from a single year of seed money to multi-year funding. The lab currently includes more than 75 MIT principal investigators and 40 IBM researchers. Funded projects are answering questions such as:

- How can advanced algorithms expand capabilities in machine learning and reasoning?
- How can quantum computing optimize machine-learning algorithms and other AI applications?
- How can AI ensure the security and privacy of medical data?

### MIT-SenseTime Alliance on Artificial Intelligence

The MIT-SenseTime Alliance on Artificial Intelligence opens new avenues of discovery across MIT in areas such as computer vision, brain-inspired algorithms, medical imaging, and robotics. SenseTime, founded by MIT alumnus Xiao'ou Tang PhD '96, specializes in computer vision and deep learning technologies. In its first year, the MIT-SenseTime Alliance funded 27 projects involving 50 MIT principal investigators across all five schools. Funded projects will answer questions such as:

- How can linguistic theory transform machine-learning algorithms to better approximate how people converse?
- How can artificial systems such as robots "learn" commonsense knowledge?
- How can product design and systems architecture tools capitalize on strategies that combine human and machine intelligence?

### Abdul Latif Jameel Clinic for Machine Learning in Health

The Abdul Latif Jameel Clinic for Machine Learning in Health (J-Clinic) harnesses the power of machine learning for preventative medicine, clinical diagnostics, and drug discovery and development. MIT professors Regina Barzilay and James Collins are the faculty leads, Anantha Chandrakasan is the chair, and MIT professor Phil Sharp is chair of the advisory board.

The J-Clinic draws on MIT's expertise in cellular and medical biology, computer science, engineering, and the social sciences as it focuses on developing machine-learning technologies to revolutionize the prevention, detection, and treatment of disease. It

concentrates on creating and commercializing high-precision, affordable, and scalable machine-learning technologies in areas of health care ranging from diagnostics to pharmaceuticals, with three main areas of focus: (1) preventative medicine methods and technologies with the potential to stop non-infectious disease in its tracks, (2) cost-effective diagnostic tests that may be able to both detect and alleviate health problems, and (3) drug discovery and development to enable faster and cheaper discovery, development, and manufacture of new pharmaceuticals, particularly those targeted for individually customized therapies. During its first year, the J-Clinic issued one call for proposals and funded 18 principal investigators across MIT. Funded projects will answer questions such as:

- How can personalized machine learning be used to monitor depression?
- How can machine learning and electronic health records be harnessed to repurpose existing drugs?
- How can deep learning help identify biomarkers linked to Alzheimer's disease for earlier detection?

### MIT-Liberty Mutual Insurance Collaboration

The MIT-Liberty Mutual Insurance Collaboration represents a $25 million, five-year commitment to advance artificial intelligence research in computer vision, natural language processing, machine-learning fairness, data privacy and security, and risk-aware decision making, among other areas. Research topics under discussion include efforts to make decision-making algorithms transparent to customers and regulators, to use computer vision to reduce crashes by identifying dangerous driving conditions and roadways, to further protect the anonymity and security of personal data, to use natural language processing to analyze insurance claims and speed processing and compensation, and to structure investment portfolios. The collaboration sent out a request for proposals in May 2019 and will start funding its first set of projects this fall.

### Siegel Family Initiative: Scaling AI the Human Way

Supported by a 10-year, $10 million gift from the Siegel Family Endowment and in collaboration with researchers at Harvard and Stanford Universities, MIT professor Josh Tenenbaum and The Core have kicked off an ambitious research program motivated by the oldest dream in AI: building a machine that grows into intelligence the way a human does—starting like a baby and learning like a child. Success would give us machine learning that involves true learning and artificial intelligence that is truly intelligent. In contrast to modern AI systems that require large data sets and are incapable of learning or thinking for themselves, we would have machines that we could teach, talk to, and trust. We would have AI that grasps new concepts from a few examples, as children do; that builds upon layers of concepts to reach a genuine understanding of the world; and that has the flexibility to solve a range of problems.

### Schmidt Family Funding for AI Education

The generous support of Eric Schmidt, technical adviser to Alphabet Inc., and his wife, Wendy, has helped fund a range of AI education projects through The Bridge.

## Accomplishments and Activities

### Crowdsourcing Bold Research Ideas

In a series of meetings with faculty last fall, we developed a set of moon shot and mission ideas for advancing human and machine intelligence research and for developing tools to democratize AI. We have used the moon shot/mission proposals in our pitches to corporate and philanthropic prospects. We have also held or are planning to hold a series of workshops to further develop our portfolio of ideas.

Our four moon shot proposals aim to teach robots common sense by reverse-engineering the intuitive physics, psychology, and language that underpin a child's ability to learn quickly and flexibly with a limited amount of data; develop computational models of how humans learn language to guide the development of leaner, more flexible natural language processing models; optimize combinations of people and computers to harness their collective intelligence and accomplish what neither can do alone; and develop computational models of the brain to help create new computer architectures.

Our three mission proposals aim to digitize MIT's collection of eight million books, videos, and other reference material to develop machine-learning models that can perform cutting-edge visual, text, and speech processing tasks; develop an AI curriculum for K–12 students that includes the use of robot mentors to personalize learning and gather data to improve teaching; and develop a platform of tools and services that any researcher at MIT can use to solve domain-specific problems.

### Current Research

Our research falls within two main categories: work to advance research in human and machine intelligence and work to develop tools and technologies to make AI more broadly accessible. The projects summarized below show the wide range of ideas MIT Quest has funded.

#### *Debugging Neural Networks*

MIT-IBM researchers Antonio Torralba, Stefanie Jegelka, and Hendrik Strobelt are developing tools to visualize what deep learning models are "thinking," allowing software developers to find and fix mistakes and ward off malicious attacks. The tools will enable developers to root out bugs in neural network nodes much as they do now in lines of code. For example, if the network confuses a construction scene with a street bazaar, the tools pinpoint the set of nodes that produced the mistake. The nodes are then retrained to fix the error.

#### *An App to Track Declining Brain Function*

MIT professors Thomas Heldt and Vivienne Sze are developing low-cost tools to identify and track Alzheimer's and other neurodegenerative diseases using a simple mobile phone app. As patients play an eye-tracking game on their phone, the camera records how quickly and accurately their eyes respond to prompts on the screen. The resulting data can tell researchers how well the patient's brain is functioning. The app, and the software being developed to crunch the data, could provide a way to track disease

progression in Alzheimer's patients and serve as an adjunct to clinical drug trials by making it easier to track improvements over time.

### *Fighting the Opioid Epidemic*

More than 115 people in the United States die each day after overdosing on opioids. The type of opioid, how much was prescribed, and for how long are all factors in who succumbs to addiction. Public health officials now hope to use this information to change how painkillers are prescribed. MIT-IBM researchers David Sontag, Dennis Wei, and Kush Varshney are applying machine learning tools to medical insurance claim records to understand what kinds of medical histories and prescription practices raise red flags. Their goal is to develop a model that can help doctors tailor prescriptions to individual patients to reduce addiction risk.

### *Understanding Real-World Actions as They Unfold*

The brain has a remarkable ability to size up a scene and quickly understand what's going on. MIT-IBM researchers Aude Oliva and Daniel Gutfreund are training machine-learning models to do something similar with a data set of one million short video clips called Moments in Time. The models learn to recognize what's happening in any particular frame, whether that's pandas playing or robots dancing or a poodle jumping for joy. If AI systems can learn to extract the gist of dynamic scenes, researchers hope to transfer the knowledge to other domains.

### *Designing a Robot with Common Sense*

A robot that can break down high-level tasks and run for weeks without getting stuck is still a long way from being built. But MIT researchers Leslie Kaelbling, Tomas Lozano-Perez, and Josh Tenenbaum hope to crack the problem by applying lessons from computing, cognitive science, and neuroscience. They are currently building an experimental infrastructure that will allow computer simulators—and, one day, real robots—to perceive and interact with the world around them, eventually achieving a semblance of common sense.

### *A Model to Learn All of the World's Languages*

Swahili may sound nothing like Quechua, but research suggests they, like most languages, have much in common. Their similar properties may explain why humans learn language so easily, acquiring new words and concepts from context, while deep learning models require mountains of training data. The need for so much data leaves voice recognition and translation software beyond reach for thousands of languages that are spoken globally but are not yet in machine-readable form. Researchers are developing a machine-learning framework to reveal the biases that let children learn language so quickly and to improve and extend language-learning models to thousands of data-scarce languages in the world.

### *Rebooting Jibo the Home Robot as a Personal Wellness Coach*

Depression and other mood disorders are still diagnosed and tracked with information patients give their doctors. To provide more personalized therapy, MIT researchers

Rosalind Picard and Cynthia Breazeal want to see if the individualized attention and support provided by a home robot like Breazeal's creation, Jibo, can make a difference. Subjects recruited to the study will grade their coach's ability to provide timely and effective advice. The study will also examine whether robot coaches offer better emotional support than state-of-the-art mobile apps already in use.

### Research Milestones

In two years, the MIT-IBM Watson AI Lab has published more than 50 peer-reviewed papers co-authored by MIT-IBM researchers. The lab's portfolio has expanded to 60 active projects with the launch of 40 new projects in fall 2019.

One project, led by MIT professor Roger Levy and presented at the 2019 annual conference of the North American Chapter of the Association for Computational Linguistics, is investigating what deep learning language models know about grammar. A second project, led by Josh Tenenbaum and presented at the 2019 International Conference on Learning Representations (ICLR), combines deep learning and symbolic programming to speed up visual recognition tasks. A third project, led by Antonio Torralba and presented at the Association for Computing Machinery's 2019 SIGGRAPH conference, lets users add and remove features from photos. This tool could help flag images that have been manipulated to mislead the public.

Funded by the MIT-SenseTime Alliance on Artificial Intelligence, a team led by MIT professor Marin Soljačić developed a form of AI that can read scientific papers and produce a plain-English summary in a sentence or two. A project led by MIT professor Song Han and presented at the 2019 ICLR describes an algorithm for optimizing machine-learning models, allowing them to run up to 200 times faster than traditional methods.

The Bridge developed a set of Jupyter notebook prototypes to help researchers identify causal relationships in gene regulation, classify stress fractures in shattered materials, generate synthetic data to supplement small data sets, and track athletes in the field from their movements to model performance. The Bridge also helped integrate GANpaint, a tool developed by the MIT-IBM Watson AI Lab, into the coding platform Scratch to teach kids about AI-generated art.

The Bridge provided $1.8 million in cloud-computing credits for AI projects on campus: $1 million in credits from Google and $800,000 from IBM. One project made possible by the IBM cloud credits led to the discovery that a deep learning model could get by with 90% fewer connections if the right subnetwork could be trained at the outset. The "lottery ticket hypothesis" proposed by MIT professor Michael Carbin and graduate student Jonathan Frankle won a best paper award at the 2019 ICLR and has already inspired several spin-off papers. The Bridge also helped facilitate IBM's donation of a nearly $12 million supercomputer to MIT to ease the demand for computing power on campus.

## Events

We organized three major tutorials and workshops that explored some of the top technical and ethical challenges facing intelligent systems. We also co-sponsored several lectures that brought familiar names in computing and AI to campus: Eric Schmidt of

Alphabet, Demis Hassabis of DeepMind, Amnon Shashua of Mobil Eye, and Alyosha Efros of the University of California, Berkeley. We organized a College of Computing poster session that highlighted computing-related research across MIT's five schools. We also co-sponsored an AI policy forum with the MIT Internet Policy Research Initiative that drew senior policymakers from the Organisation for Economic Co-operation and Development and resulted in coverage in the *New York Times*. Two of the three workshops—on robustness and interpretability and AI hardware—convened MIT leaders in the field to further develop the moon shot ideas earlier proposed by faculty.

### Workshop on Robust, Interpretable Deep Learning Systems

Organized by Professors Aleksander Madry and Antonio Torralba, this workshop explored methods for attacking and defending deep neural networks and visualizing their behavior, structure, sensitivities, and biases. As more decision making becomes automated, the emerging field of adversarial machine learning has focused on making deep learning systems more reliable, secure, and transparent. The workshop featured talks by MIT leaders in the field and drew about 200 participants, including approximately 40 graduate students who presented posters on their research.

### Workshop on Intelligent Hardware Technologies: 10 Years Out

Organized by MIT professor Jesus del Alamo, this workshop examined the various approaches researchers at MIT and elsewhere are taking to develop next-generation computers, from new material systems to circuits and architectures. The workshop featured talks by MIT leaders in AI hardware, including those looking to the brain for inspiration on how to build faster, more energy-efficient machines. The event drew about 100 participants, including about 20 students who presented posters.

### GANocracy: Workshop on Theory, Practice and Artistry of Deep Generative Modeling

Organized by Professor Phillip Isola and Research Scientist Aude Oliva, this tutorial and workshop explored the art and science of generative adversarial networks, or GANs, a machine-learning technique that harnesses a pair of neural networks to create hyper-realistic images and sounds. The morning tutorial took attendees through the technical elements of GANs, while the afternoon workshop featured artists and researchers discussing emerging applications, including the use of GANs to create synthetic data to fill in gaps in time-series data and to test AI models for bias. The workshop also touched on the problem of deep fakes, or images manipulated to deceive people, and the AI tools that might be developed to combat them.

The tutorial and workshop together drew about 250 people, with 50 students presenting posters. The tutorial remains online and available to anyone who wants to learn more about GANs. Two months after the event, Isola and Torralba were contacted by Facebook to join a global competition to develop AI tools to detect deep fakes. Announced by Facebook in September, the competition is intended in part to address the spread of misinformation before the next presidential election.

## Education Activities

### Undergraduate Research

MIT Quest has helped to train and mentor undergraduates interested in AI-related applications through the Undergraduate Research Opportunities Program (UROP). We sponsored 47 fall UROP and 62 spring UROP projects. We also sponsored 47 projects over Independent Activities Period (IAP) and 24 projects this summer.

In one ongoing project, several UROP students are working with Bridge engineers to build a computer interface called Monkey that will allow students to send AI models and training data to the cloud, put projects in a queue, train the models, and send projects back when finished. Monkey will also track individual cloud-credit usage. In another project, students worked with Bridge engineers to adapt GANpaint, a program developed by the MIT-IBM Watson AI Lab, to Scratch, the popular coding platform for kids. The work involved training a GAN on new images and integrating GANpaint with Scratch. A related project will use Jupyter notebooks to teach people to think critically about generative models.

Other UROP projects have focused on applying AI tools to new domains. Professor Rafael Gomez-Bombarelli is using machine-learning tools to narrow the search for promising drug candidates by predicting which molecules are most likely to bind with a target protein in the body. A sophomore worked with him this year to build a database of small molecules and proteins detailing their chemical structures and binding properties, also working on a deep learning framework aimed at predicting which molecule-protein pairs have the strongest binding affinity and are thus more promising drug candidates. Professor James Glass has developed a voice-controlled nutrition app called Coco Nutritionist to let users log their meals by talking into their phone instead of typing. A third-year student helped to develop the user interface and build a new feature on the back end for adding recipes and homemade meals. So far, the app has reached 8,500 downloads in the Apple Store.

### In the Classroom and Beyond

MIT Quest has forged a partnership with the undergraduate-led MIT Machine Intelligence Community (MIC) to advance our common cause of making AI tools accessible to all. MIC hosts regular AI talks and tutorials for undergraduates, and its members are working with Quest staff and faculty on several UROP projects, including the Monkey interface (described above), to expand and promote AI computing on campus. The MIT-IBM Watson AI Lab will help fund MIC's fall symposium on September 7, 2019, at Boston University, which will feature MIT and IBM speakers.

The MIT-IBM Watson AI Lab made available cloud computing services and a unique insurance-claims data set to students taking David Sontag and Peter Szolovits's 6.S897/HST.956 Machine Learning for Healthcare class. The lab also helped support a student-led introduction to deep learning class during IAP week that drew a record number of students. Similar collaborations are expected in the year ahead.

## Leadership

MIT Quest leadership is composed of world experts in computer vision, robotics, the human brain, and human behavior.

Antonio Torralba is the director of MIT Quest and the MIT director of the MIT-IBM Watson AI Lab, a professor of electrical engineering and computer science, and an investigator at the Computer Science and Artificial Intelligence Laboratory.

Aude Oliva is the executive director of MIT Quest and the MIT-IBM Watson AI Lab and a principal research scientist at the Computer Science and Artificial Intelligence Laboratory.

James DiCarlo is the director of The Core for MIT Quest, the head of the Department of Brain and Cognitive Sciences, the Peter de Florez Professor of Neuroscience, and an investigator at the McGovern Institute for Brain Research.

Nicholas Roy is the director of The Bridge for MIT Quest, the Bisplinghoff Professor of Aeronautics and Astronautics, and an investigator at the Computer Science and Artificial Intelligence Laboratory.

Daniela Rus is the associate director of The Core for MIT Quest, the director of the Computer Science and Artificial Intelligence Laboratory, and the Andrew and Erna Viterbi Professor of Electrical Engineering and Computer Science.

Cynthia Breazeal is the associate director of strategic initiatives for The Bridge for MIT Quest, an associate professor of media arts and sciences, and the director of the Personal Robots Group at the MIT Media Lab.

Joshua Tenenbaum is a scientific director of The Core for MIT Quest; a professor of computational cognitive science; a researcher at the Center for Brains, Minds and Machines; and an investigator at the Computer Science and Artificial Intelligence Laboratory.

Leslie Kaelbling is a scientific director of The Core for MIT Quest, the Panasonic Professor of Computer Science and Engineering, and an investigator at the Computer Science and Artificial Intelligence Laboratory.

Tomaso Poggio is the founding scientific advisor of The Core for MIT Quest; the director of the Center for Brains, Minds and Machines; the Eugene McDermott Professor in the Department of Brain and Cognitive Sciences; and an investigator at the Computer Science and Artificial Intelligence Laboratory.

Erik Vogan is the director of corporate engagement for MIT Quest and a program director in the Office of Corporate Relations.

Rachel Donahue is the director of development for MIT Quest and the Department of Brain and Cognitive Sciences.

## Communications

We have steadily grown our following on Twitter and Facebook (@MIT_Quest and @MITQuest4Intelligence) by posting general AI news and news items featuring MIT AI researchers, including those affiliated with MIT Quest. We have built out our website to feature events, staff, selected research projects, and messaging for prospective corporate and philanthropic donors. We have also established a presence on YouTube and MIT Events. We put out our first MIT Quest newsletter this summer, highlighting MIT Quest events and research news from the last year. We plan to send out a similar newsletter on behalf of the MIT-IBM Watson AI Lab.

## Summary

By bringing together researchers across MIT's five schools, MIT Quest aims to go beyond business as usual. In the past year, we have made great strides in advancing our ambitious research goals in the science and engineering of intelligence and providing custom-built AI tools to researchers across disciplines. We have formulated a series of moon shots and missions and are now actively seeking funding to set these projects in motion. We have also established several multi-year collaborations with industry, nonprofits, and individual philanthropists. Through their generosity, we have expanded our research portfolio to more than 110 projects and created 180 UROP projects for undergraduates to get involved. Research topics range from fundamental problems in the science and engineering of intelligence to applications with shorter term and more immediate benefits. MIT Quest remains committed to its holistic vision for shaping the future of the science and engineering of intelligence, and we look forward to joining the new MIT Stephen A. Schwarzman College of Computing in our common cause.

**Antonio Torralba**
**Director**

**Aude Oliva**
**Executive Director**