The MIT Quest for Intelligence Report to the President 2022–2023

Understanding human intelligence is one of the greatest human endeavors of all time – right alongside existential questions such as our world's position in the universe and the origin of life. The MIT Quest for Intelligence (The Quest) is the only organizational unit at MIT aimed directly at this question. The past year has been a good one for the Quest as it continues to have significant influence on research and applications at the interface of Natural Intelligence (NI) and Artificial Intelligence (AI). Adjustments and improvements to the organization's focus, administration, and staffing have resulted in positive outcomes. Our efforts to fund Missions — interdisciplinary teams of researchers, each spanning science and engineering, and each focused on a specific domain of intelligence — have been successful, and we have launched a Language Mission and an effort in Scaling Inference. After several successful hiring searches, our expanded staff are working to support these Missions, our interface with the MIT community, and ongoing research efforts.

Leadership and Affiliated Researchers

James DiCarlo, Peter de Florez Professor of Systems and Computational Neuroscience (BCS), is the Quest Director; Nicholas Roy, Professor of Aeronautics and Astronautics, is the Director of MIT Quest Systems Engineering; Professor Joshua Tenenbaum (BCS) is the Director of Science; Leslie Kaelbling, Panasonic Professor in the Department of Electrical Engineering and Computer Science (EECS) is the Director of Research; Vikash Mansinghka, Principal Research Scientist (BCS) is the Director of Modeling and Inference; and Erik M. Vogan is Executive Director. CBMM is led by Tomaso Poggio, the Eugene McDermott Professor (BCS).

Researchers representing labs, centers, and academic departments across the Institute are involved in Quest-sponsored research:

- Schwarzman College of Computing (SCC), Electrical Engineering and Computer Science Department (EECS): Associate Professor Vivienne Sze; Pulkit Agrawal, Steven G. (1968) and Renee Finn Career Development Professor; Associate Professor Jacob Andreas; Jesus del Alamo, Donner Professor of Science; Joel Emer, Professor of the Practice; William Freeman, Thomas and Gerd Perkins Professor of Electrical Engineering and Computer Science; Harry Lee, Advanced Television and Signal Processing Professor of Electrical Engineering; Tomas Lozano-Perez, School of Engineering Professor of Teaching Excellence; Professor Martin Rinard; Russ Tedrake, Toyota Professor.
- Computer Science and Artificial Intelligence Laboratory (CSAIL): Daniela Rus, Director.
- Brain and Cognitive Sciences Department (BCS): Ev Fedorenko, Middleton Career Development Professor of Neuroscience; Professor Ila Fiete; Nancy Kanwisher, Walter A. Rosenblith Professor; Rebecca Saxe, John W. Jarve (1978) Professor, Associate Dean of the School of Science; Laura Schulz, Professor of Cognitive Science; Professor Michale Fee, Glen V. and Phyllis F. Dorflinger Professor and head of BCS; Associate Professor Steven Flavell; John Gabrieli, Grover M. Hermann Professor; Associate

Professor Mehrdad Jazayeri; Professor Roger Levy; Professor Alexander Rakhlin; Mriganka Sur, Newton Professor; Assistant Professor Guangyu Robert Yang.

- Sloan School of Management: Abdullah Almaatouq, Douglas Drane Career Development Professor in Information Technology; John Horton, Associate Professor of Information Technologies; Thomas Malone, Patrick J. McGovern (1959) Professor of Management; Manish Raghavan, Drew Houston (2005) Career Development Professor; Professor David Rand.
- Departments of Nuclear Science and Engineering and Materials Science and Engineering: Bilge Yildiz, Breene M. Kerr (1951) Professor; Ju Li, Battelle Energy Alliance Professor in Nuclear Engineering.
- Department of Mechanical Engineering: Faez Ahmed, Brit (1961) and Alex (1949) d'Arbeloff Career Development Professor in Engineering Design.
- Department of Civil and Environmental Engineering (CEE): Ali Jadbabaie, J.R. East Professor and head of CEE.
- Media Arts and Sciences: Professor Cynthia Breazeal, Dean of Digital Learning; Alex (Sandy) Pentland, Toshiba Professor of Media Arts & Science.
- Political Science: Adam Berinsky, Mitsui Professor of Political Science.

Research

Researchers in the Quest aim to understand intelligence — how brains produce it and how it can be replicated in artificial systems. We approach this as a single grand challenge requiring the organized, collaborative efforts of science, engineering, the humanities and beyond. MIT's Center for Brains, Minds, and Machines (CBMM), an NSF-funded science and technology center focused on the interdisciplinary study of intelligence and how it can be replicated in machines, is administratively housed in the Quest and is currently in its tenth year of funding.

To execute on its vision, the Quest has established "*Missions*," long-term collaborative projects rooted in foundational questions in and centered around a single domain of intelligence, and "*Platforms*," software systems that enable Missions research in new directions and benchmarking and testing interfaces that use data from the Missions to help the researchers refine and expand their work. This past year, the Quest launched a Language Mission and the Scaling Inference Platform and continued to support the original four Missions, Developing Intelligence, Embodied Intelligence, Collective Intelligence and Hardware for AI. The Quest provides institutional support, guidance, and engineering support for each Mission.

Center for Brains, Minds, and Machines (CBMM)

CBMM is a multi-institutional, NSF-funded Science and Technology Center headquartered at MIT. For ten years, CBMM's research, education, and outreach programs have become an important part of the MIT environment, exploring the ways that the natural science of the brain and mind plays a key role advancing machine intelligence. As its NSF funding ends this year, CBMM leadership has focused energy on preserving its legacy of educational and outreach programs, and has worked closely with the Quest to continue to nurture and grow the community of intelligence researchers that CBMM nucleated a decade ago.

CBMM hosts regular internal research meetings, providing researchers a chance to present and discuss current work with peers. CBMM continues to strengthen its relationships with industry and international partners.

In January, the CBMM Diversity Program hosted the annual Quantitative Methods Workshop, an intensive seven-day program on computational and cognitive neuroscience methods. This summer it is hosting ten-week Undergraduate Summer Research Internships in Neuroscience.

In August 2022, CBMM held the ninth annual Brains, Minds, and Machines summer course, an influential three-week, multidisciplinary course on the science of intelligence. This course is cultivating a community of leaders knowledgeable in neuroscience, cognitive science, and computer science who will lead the development of true biologically-inspired AI. The course returned to its pre-pandemic enrollment level of 35 students attracted from top universities around the world.

Engineering Team

The Quest Engineering Team develops and maintains software systems that accelerate and integrate the work of the Quest Missions. A key tool in this endeavor is the Brain-Score platform, which enables direct comparison between neural representations and behavior of humans/non-human primates and the artificial neural representation and behavior of computational models. These comparisons help researchers understand how closely aligned a particular computational model is with data collected from biological agents, allowing them to explore areas of similarity or difference. While Brain-Score originally focused on models of the ventral (visual) stream, its comparison framework is generalizable across cognitive systems. In Q3 & Q4 of 2022, in service of the Quest Language Mission, the Engineering Team extended Brain-Score to include the domain of language, and development is currently underway to accommodate embodied agents for the Embodied Intelligence Mission. In Q1 2023, the Team focused on a project requested by the Developing Intelligence Mission: improving the accuracy and usability of iCatcher+, an AI-based tool for automatically annotating video data collected during Early Childhood Development studies. In addition to a nearly 10% gain in overall accuracy (more than 90% of frames per video are correctly labeled), the team constructed a graphical user interface that allows researchers to easily review the annotated data and modify as needed. These contributions significantly reduce the time and resources required to process collected data, thereby removing a major obstacle to conducting and scaling research involving infants and pre-verbal children.

Quest Platforms

Scaling Inference

Lead PIs: *Mansinghka and Tenenbaum*
Collaborators: *Rinard, Kaelbling, Roy, Andreas, Freeman, Schulz, Flavell, DiCarlo*
A great deal of enthusiasm, in both AI and in brain and cognitive sciences, is focused on building large neural network models. This team is pursuing an alternate scaling route for AI systems and for NI models, based on inference in probabilistic programs. Their AI-facing goal is to show that end-to-end explainable AI systems built using probabilistic programming can match and exceed the speed, robustness, and flexibility of human intelligence, using 100x-1,000x less computation than deep learning. Their NI-facing goal is to leverage new techniques for neural mapping of

probabilistic programs to build and test these AI systems as computational models of perception and cognition.

The AI-facing goals both draw on and contribute to an open-source platform:

- ChiSight: Real-time 3D-scene perception that learns to perceive new objects and scenes on one GPU in real time, aiming to be more robust than transformers trained offline using 1,000+ GPU-hours.
- ChiExpertise: Trustworthy conversational AI that gives grounded, auditable answers, and aims to be more accurate than GPT4 in data-driven domains, using models built and fine-tuned on one GPU.
- OpenGen: The first probabilistic programming stack aiming to be as widely adoptable as TensorFlow v1.

The group is collaborating with the Embodied Intelligence, Developing Intelligence, and Language Missions to test ChiSight and ChiExpertise as models of NI, as well as enabling adoption of the OpenGen platform more broadly across MIT labs.

Accomplishments include:

- A 3D-scene perception system that is more robust than deep learning and that provided one of just two "bases of confidence" for a new DARPA program in Assured Autonomy
- The first generalization of automatic differentiation that correctly estimates gradients of the expected values of probabilistic processes, receiving a SIGPLAN Distinguished Paper Award
- OpenGen was used to build the first whole-brain probabilistic models linking neural activity and behavior in *C. elegans*, in press at the journal *Cell.*

Quest Missions

Embodied Intelligence Mission (EI)

Lead PIs: *Kanwisher and Kaelbling*
Collaborators: *DiCarlo, Jazayeri, Fiete, Lozano-Perez, Roy, Tedrake, Tenenbaum*
This Mission's goals are to improve both our understanding of intelligent behavior in animals and humans and our ability to construct artificial systems that interact (via simulation or in reality) with the physical world. The team studies intelligent behavior at two scales: a "table-top" scale and a "house" scale.

At the "table-top scale", the PIs have formulated a set of manipulation problems under partial observability (called "find the grape") that they believe can be solved by monkeys, humans in virtual reality, simulated robots, and real robots. Progress includes:

- Using human VR studies to test the richness and precision of people's understanding of the geometry and physics of tabletop scenes.
- Designing and fabricating apparatus for monkey experiments
- Improving and implementing robot infrastructure, and developing novel planning algorithms for efficiently handling uncertainty and reasoning about unobserved space.

The "house scale" is in earlier stages with a focus on robot software architecture and implementation, and the PIs are working to integrate capabilities that have traditionally been studied independently, including mapping, navigation, object recognition, long-term memory, and planning. They aim to enable a robot to explore a house-sized environment, noting the locations of salient objects but not necessarily making a complete inventory, and subsequently being able to retrieve an object (perhaps one it has already seen or one it will still have to look for) or put away new objects, respecting the "organization" of the house. In the coming semester, they will focus on outlining concrete connections to natural intelligence systems (possibly rats and humans in VR) and on pushing the implementation effort on simulated and real robots.

Developing Intelligence: Scaling AI the Human Way

Lead PI: *Tenenbaum*

Collaborators: *Schulz, Saxe, Mansinghka, Kaelbling, Tedrake*

This Mission broadly aims to understand how human learners grasp new concepts from very few examples, and how children build upon layers of concepts to reach an understanding of the world with the flexibility to solve an unbounded range of problems. Is it possible to build AI that starts like a baby and learns like a child? Aims include building computational models of the core common-sense knowledge that represents the "start-up software" of the brain, the perception algorithms that allow infants to grasp the state of the physical world and other agents' goals in terms of these common-sense representations, and the learning algorithms used by babies over the first 18 months to grow, enrich, and ultimately move beyond their initial mental models. Progress towards these aims will have many technological and societal payoffs, including robots that can more flexibly adapt to new situations and robustly perceive their environments, and a better understanding of how children learn for the purposes of early childhood education and developmental interventions. Over the last year, our team has made multiple advances, including:

- Establishing a new integrated online platform for scalable, at home developmental experiments with infants and young children, which has already been used by dozens of other labs around the country, and which merges our previous systems "Children Helping Science" and "Lookit.mit.edu"
- Developing and benchmarking software tools for automating running and analysis of infant experiments
- Development of probabilistic programming AI models for infant-inspired 3D common-sense scene understanding that qualitatively and quantitatively beat state-of-the-art baselines in robotic perception, while requiring far less training data and generalizing more robustly to atypical but still quite frequent situations
- Computational models of theory of mind that represent the first comprehensive account of inferring beliefs, goals, and preferences, in both infants and adults, including the first solution of the "Baby Intuitions Benchmark"

Collective Intelligence (CI)

Lead PIs: *Malone, Almaatouq, Rand, Rus*

Collaborators: *Ahmed, Berinsky, Horton, Jadbabaie, Pentland, Raghavan, Yang*

This Mission studies one of the most important types of intelligence in the world: the collective intelligence that arises in groups of individuals, whether those individuals are people, computers, animals, or combinations of these and other entities. One important goal for the Mission is to identify similarities in how intelligence can emerge in these different kinds of groups. To do this, the team is developing a body of theory called an "ontology of collective intelligence" that identifies a wide range of possible processes and performance prediction models for tasks such as group decision-making.

The team is also developing and testing a variety of innovative examples of how humans and AI systems can work together. For instance, they have developed a generative AI-based tool, called the "Supermind Ideator," to help design innovative group processes that combine people and AI in solving business and societal problems. They are also developing AI-based tools to help humans write software, design visual objects, and write short stories. Together, they expect this work to help better understand both the science of collective intelligence and its application to important practical problems.

Brain-Guided Intelligence Hardware

Lead PIs: *Yildiz, del Alamo*

Collaborators: *Fee, Li, Sze, Lee, Emer*

Memristive synaptic devices constitute the key engineered element that links neuroscience with future machine intelligence hardware. The programmable non-volatile analog memory characteristics of memristors resemble the functionality of biological synapses. This Mission's research is developing a framework for the analysis and design of new machine intelligence systems based on novel memristive devices. It spans the entire abstraction stack, from new material systems and device principles, to AI system architecture design. Permeating the entire effort is a keen focus on energy efficiency and bio-plausibility. The team is developing novel electrochemical ionic synapses that target the specs required to implement analog neural networks with inference accuracy that matches digital neural networks while enhancing energy efficiency by 100X. They are also constructing a framework to efficiently explore the vast design space of possible system architectures to optimize system energy consumption, computational throughput and accuracy. Due to its rich ion dynamics, the synapses mimic the temporal dynamics of brain synapses, capturing millisecond-processes, such as synaptic potentiation, as well as evolving in the time scale of years associated with long-term memory. Based on this, they are developing a novel ionic hardware architecture that emulates the biological learning of the complex vocal behavior of a songbird.

Language

Lead PIs: *Fedorenko, Andreas, Levy*

Large language models are fundamental building blocks in many modern AI systems — for language processing, as well as robotics, computer vision, software engineering, and more. For models trained on text to be useful for general AI and scientific applications, they must understand not just the structure of language, but the structure of the world; moreover, their language, reasoning, and world knowledge capabilities must align with those in humans. This research aims to provide a robust theoretically motivated and empirically grounded framework for studying and improving world knowledge and reasoning capabilities in large language models, and using our understanding of human cognition to make models better. The team

- has proposed a framework for dissociating language and thought in large language models;
- has built a proof-of-concept system that improves "thinking" in language models by translating language into code, then running it with a probabilistic inference engine;
- has identified substantial limitations in the ability of these models to reason flexibly;

- is designing a comprehensive language modeling benchmark targeting key aspects of world knowledge including knowledge of both physical and social concepts;
- is constructing benchmarks for learning word meanings, using these benchmarks to evaluate language models trained with ecologically plausible data, and identifying why children acquire words in different patterns compared to models.

Seed funding

In addition to the Platforms and Missions, the Quest also funds four early-stage projects.

Building and evaluating multi-system functional brain models

Lead PI: *Yang*
Integrative understanding of the brain is offered through multi-system neural network models of brain functions incorporating multiple modules for diverse neural systems. Using a high-throughput pipeline that was developed recently to compare brain models with neural data recorded from animals, the researchers have made a wide range of discoveries. They found that neural networks with more biological structures ended up having neural activity more closely aligned with the brain. Meanwhile, "low-rank" neural networks with attractive mathematical properties do not match the brain as well, despite the theoretical appeal.

These results have been presented at major conferences in the field (COSYNE, CCN).

Reimagining Reinforcement Learning via Memory-Based Biological Learning

Lead PI: *Agrawal*
Co-PIs: *Rakhlin, Sur, Fiete*
This team's significant progress in efforts to understand the role of memory in animal cognition and how it enables rapid multi-task learning has led to computational models that may explain animal behavior in rich naturalistic settings, offering tools to study and mitigate memory loss and its consequences.

- Advances in Exploration–Exploitation: Rakhlin's lab is developing a theoretical foundation for sequential decision-making. Agrawal lab developed an algorithm that trades off exploration bonus against task reward. Agrawal and Rakhlin are measuring "novelty" in high-dimensional spaces. Preliminary results in an Agrawal/Sur collaboration show that developed models can explain glia activity and their control of exploration-exploitation in rodents.
- Exploration via Inductive Extrapolation: Fiete's lab developed approaches for fast exploration of new spaces by identifying regularities in a given environment. In collaboration with Agrawal lab, they developed a computational model for exploration that leverages long- and short-term memory for quick learning; publications are submitted and posted to arXiv.
- On both the computational and biological side, they are pursuing interdisciplinary directions to inform machine learning models and to use ML models to explain brain data better.

A group led by Agrawal received MURI funding to study memory-based computation in biology and build ML systems capable of multimodal and compositional generalization.

Modeling Working Memory using Assemblies of Neurons

PI: *Poggio*

Working Memory (WM) is the online maintenance of information available for processing and executing higher cognitive functions. Transferring WM across hemispheres involves novel ensembles of neurons which converge to the generic ensembles recruited in the original hemisphere. This project replicates these observations using a computational model of neural ensembles. While random synaptic connectivity explains the recruitment of novel ensembles in the new hemisphere, synaptic plasticity is required to explain the convergence of novel ensembles in the new hemisphere to the ensembles in the original hemisphere. Ablation of recurrent synaptic connections confirms the role of synaptic plasticity in WM transfer.

Understanding Social Mechanisms of Learning in Humans and Machines with Application to Early Childhood Education

PIs: *Breazeal and Gabrieli*

This project studies how to design AI technologies responsibly to support young children's learning with parents. The team designs, develops, and evaluates social robots that are capable of fostering long-term parent–child social and affective interaction at home. In their most recent real-world robot deployment, more than 70 families engaged in co-reading sessions with social robots in their households over 2–3 months. Data were gathered through surveys, interviews, and audiovisual recordings. Employing a multimodal, mixed-method analysis approach, they explored the influence of robot interactions on parent-child conversational dynamics, adaptive robot behavior, parents' self-reported experiences and perceptions, and their imagination for future robot design. Their findings highlight the potential of social robots as a social catalyst and present valuable insights for the responsible robot design in the context of parent-child interaction.

Research Affiliates and Industry Collaborations

The Quest's engagement program offers companies a variety of ways to advance their strategic goals. Through this program, we host the MIT-Liberty Mutual Insurance Collaboration.

In 2022–2023 our total spent fund volume was $4.4 million, and the MIT-Liberty Mutual Insurance Collaboration had $1.4 million in total secondary research volume.

Community Outreach and Activities

Education

The Quest trains and mentors undergraduates interested in neuroscience, psychology, and software engineering through UROP. This year, we supported nine fall UROPs, four IAP UROPS, twelve spring UROPs (plus two volunteer UROPS), and we are currently supporting three summer UROPS. A past UROP, Khaled Shehada, continued to work with the Engineering Team and will complete his MEng this summer.

The Quest will continue to support 50% of a post-doc fellow working in SERC in SCC, Miriam Boulicault, through the end of her appointment in August 2023, and will support an incoming post-doc for the next fiscal year.

Events

The Quest provided financial and organizational support to events and student activities to benefit the AI community at MIT and beyond. In November, the Quest and CBMM held "Advances in the quest to understand intelligence," the first in a series of planned bi-annual workshops to share the latest progress on understanding natural intelligence and how that scientific progress will drive AI and other impact areas. In March, we co-sponsored a panel discussion with the student group Artificial Intelligence at MIT (AI@MIT) which included Quest PIs and members of the Engineering Team. We also hosted two AI@MIT lunches that allowed students to have conversations with PIs in an informal setting. The Quest and CBMM co-sponsored seminar series hosted nine speakers, representing industry and academia in neuroscience, cognitive science, and computer science.

Communications

The Senior Communications Officer is working closely with the Senior Development Officer and Events Administrator to create opportunities for engagement and stewardship with current and potential philanthropic supporters. They will collaborate on print and electronic communications to build connections with this community.

Our communications platforms are growing, providing more information and resources for Engineering projects. In the coming months, staff will work together to create an archive of CBMM research and teaching resources.

Administrative and Engineering Staff

The leadership team's efforts to hire and promote administrative and engineering staff have been highly successful, and the Quest's current staff now comprise: Erik Vogan, Executive Director; Frances Hamilton, Senior Development Officer; Rachel Kemper, Senior Communications Officer; Jim Neidhoefer, Missions Project Manager; Brian Pierson; Senior Financial Administrator; Allison Provaire, Project and Events Administrator; and Carissa Leal, Senior Administrative Assistant. Katherine Fairchild leads the Software Engineering Team of Sam Winebrake, Deirdre Kelliher, and Ethan Pellegrini, all of whom are new hires at MIT. CBMM staff are Kathleen Sullivan, Managing Director, and Kris Brewer, Director of Technology. During this year, Rachel Kemper was promoted to Senior Communications Officer, and Martin Schrimpf accepted a faculty position at École Polytechnique Fédérale de Lausanne.

Future Plans

In the coming years, the Quest will build on the foundations of our core Missions and take steps to develop machine-executable models of core aspects of natural intelligence that work computationally, cognitively, and neurally. We will also develop and use new platforms for building and testing models of intelligence, and show that these models are more robust, efficient, and understandable than today's deep learning systems.

While work within a core set of Missions (Developing Intelligence, Embodied Intelligence, Language) has progressed significantly since they launched, these Missions are early in their life cycle and will need continued resources to develop further and to leverage the Cognitive AI platform. We expect that a core set of Missions will go through one or two full cycles of a feedback loop where natural intelligence benchmarks and associated hypotheses inform and guide systems engineering (i.e., machine-executable model building) resulting in hypothesis refinement or generation of new scientific hypotheses. We expect the five existing Missions will remain largely stable over this period, with some adjustments to focus resources on those areas that are most ready to scale.

The Quest plans to build an *Intelligence Observatory*, a human behavioral testing and benchmarking platform designed to work at scale. Human and non-human primate data captured from a variety of internal and external collaborators will fuel the natural intelligence benchmarking platform already under development, and that platform will in turn inform and guide the Missions. The Quest Engineering Team will support this work, alongside continuing to develop tools needed by the Missions and supporting infrastructure.

In addition, during this period we aim to facilitate the launch and growth of an external (to MIT) non-profit organization to advance the Cognitive AI platform. This non-profit will build on the work already accomplished by the Developing Intelligence mission and the Scaling Inference platform. Initial discussions on a gift to support this effort are underway and are expected to conclude soon.

To accomplish these goals, we will continue to pursue funding opportunities, build strategic partnerships, increase our efforts in fundraising and stewardship, hire and promote software engineers and project managers, and align with MIT's goals.

In the coming months, our offices will shift to the new SCC building. The building's physical location between Stata and Building 46 (BCS headquarters) will allow Quest staff to interact regularly with faculty and students and we hope to establish seminars and other programming in the new space.

James DiCarlo, MD, PhD
Director, MIT Quest for Intelligence
Peter de Florez Professor of Systems & Computational Neuroscience