Report to the President year ended June 30, 2024, The MIT Quest for Intelligence

The endeavor to understand intelligence in engineering terms drives The MIT Quest for Intelligence (The Quest). Our faculty, staff, and students are focused on research and applications at the interface of Natural Intelligence (NI) and Artificial Intelligence (AI). Over the past year, we have seen significant progress in the work; this progress is due, in part, to a stable team, advances in the research tools built by the Engineering Team, and increased support from inside and outside the Institute. We have issued another round of funding to our Missions — interdisciplinary teams of researchers, each spanning science and engineering, and each focused on a specific domain of intelligence.

Recent significant changes and milestones include launching the Perceptual Intelligence Mission, taking steps towards establishing the Intelligence Observatory, and seeing community adoption of the Brain-Score Platform as a research tool. Several ongoing hiring searches have been completed and workloads are better balanced. With the opening of the Schwarzman College of Computing (SCC), Building 45, our offices have reached their planned destination, allowing us easy and frequent access to colleagues and labs in Building 46 and the Stata Center.

Leadership and Affiliated Researchers

James DiCarlo, Peter de Florez Professor of Systems and Computational Neuroscience, is the Quest Director; Nicholas Roy, Professor of Aeronautics and Astronautics, is the Director of MIT Quest Systems Engineering; Professor Joshua Tenenbaum is the Director of Science; Leslie Pack Kaelbling, Panasonic Professor in the Department of Electrical Engineering and Computer Science is the Director of Research; Vikash Mansinghka, Principal Research Scientist is the Director of Modeling and Inference; and Erik M. Vogan is Executive Director. The Center for Brains, Minds, and Machines is co-led by Tomaso Poggio, the Eugene McDermott Professor.

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

- Schwarzman College of Computing (SCC), Department of Electrical Engineering and Computer Science (EECS): Associate Professor Jacob Andreas; William Freeman, Thomas and Gerd Perkins Professor of Electrical Engineering and Computer Science; Tomas Lozano-Perez, School of Engineering Professor of Teaching Excellence; Professor Martin Rinard; Russ Tedrake, Toyota Professor; Leslie Kaelbling.
- Computer Science and Artificial Intelligence Laboratory (CSAIL): Daniela Rus, Director, and Andrew (1956) and Erna Viterbi Professor of Electrical Engineering and Computer Science.
- Department of Brain and Cognitive Sciences (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; Associate Professor

Steven Flavell; Associate Professor Mehrdad Jazayeri; Professor Roger Levy; Jim DiCarlo; Josh Tenenbaum; Vikash Mansinghka; Tommy Poggio.

- 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; David Rand, Erwin H. Schell Professor.
- Department of Mechanical Engineering: Faez Ahmed, ABS Career Development Professor of Mechanical Engineering.
- Department of Aeronautics and Astronautics: Nick Roy.
- Department of Political Science: Adam Berinsky, Mitsui Professor of Political Science.

Research

Researchers in the Quest aim to understand intelligence in engineering terms — how brains produce it and how it can be replicated in artificial systems. This single grand challenge requires the organized, collaborative efforts of specialists representing science, engineering, the humanities, and social sciences.

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*," 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 Mission dedicated to Perceptual Intelligence and continued its support of the Developing Intelligence, Embodied Intelligence, Collective Intelligence, and Language Missions. Three key software platforms have been developed: Brain-Score has matured and is now an important research tool, and the Scaling Inference Platform and Intelligence Observatory are growing as well. The Quest provides institutional support, guidance, and engineering support for each Mission and these supporting software platforms.

Center for Brains, Minds, and Machines (CBMM)

CBMM is a multi-institutional, NSF-funded Science and Technology Center headquartered at MIT that was launched eleven years ago (currently in a no cost extension). CBMM's research, education, and outreach programs are 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, CBMM leadership is working with the Quest to continue to nurture and grow the community of intelligence researchers that CBMM nucleated and to preserve its legacy of educational and outreach programs.

The CBMM Diversity Program annually hosts two outreach programs: the Quantitative Methods Workshop, an intensive seven-day program on computational and cognitive neuroscience methods, held in January, and a ten-week summer Undergraduate Summer Research Internships in Neuroscience.

Engineering Team

The Quest Engineering Team develops and maintains software systems that accelerate and integrate the work of the Quest Missions. Earlier this past year, the Team developed the Quest Robotics/Embodied Intelligence procedural experiment creation and benchmarking system (QREI). QREI facilitates experimentation and benchmarking and easily scales to incorporate relevant natural intelligence benchmarks as the requisite data becomes available. Also, this past year, the Team focused effort on the Brain-Score Platform, assisting in the release of Brain-Score 2.0. They continue to develop tools for the Missions and supporting infrastructure.

Quest Platforms

Brain-Score

Brain-Score enables direct comparison between neural representations and behavior of humans/non-human primates and the artificial neural representation and behavior of computational models. Such 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. This past year saw the release of Brain-Score 2.0, which is significantly better in performance, scalability, stability, and overall user experience than version 1.0. The release of Brain-Score 2.0 enabled the launch of the 2024 Brain-Score Benchmark competition, which will reward benchmarks that highlight lack of alignment between the primate visual ventral stream and new or existing computational models of vision; winners will share $10,000 in prizes and will be recognized at the Cognitive Computational Neuroscience Conference to be held at MIT in August.

Intelligence Observatory

Steps have been taken toward establishing 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 and increasingly open data sources 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 includes three staff to support this work.

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 over the past year include:

- Developing and quantitatively testing the first theory of brain computation that simultaneously predicts quantitative datasets on synaptic physiology, long-range functional neuroanatomy, and spiking of cortical neurons from the Allen Brain Atlas. This theory of brain computation as real-time spiking neural Monte Carlo inference in probabilistic programs is widely considered to be the first computational theory with predictions that are sufficiently accurate to detect errors in the Allen Brain Atlas.
- Launching and collaborating with CHI-FRO.org, a non-profit Focused Research Organization supporting the Quest for Intelligence and the main goals of the Scaling Inference Platform via open-source research engineering.
- Developing GenJAX, the first probabilistic programming system with massively parallel programmable inference. There is broad demand for GenJAX across Quest Missions and other MIT faculty labs. For example, GenJAX libraries are being used by the Perceptual Intelligence, Developing Intelligence, and Embodied Intelligence Missions, and also by collaborative research projects with multiple faculty labs across multiple MIT departments and schools.

Research goals for the next year include:

- Releasing the OpenGen platform, along with libraries for core faculties of intelligence that grow out of collaborative projects with multiple Quest Missions, such as 3D scene perception, grounded language understanding, and language model probabilistic programming.
- Quantifying scaling of Gen models for 3D scene perception and grounded conversational AI expert systems, aiming to demonstrate ~100–1,000x scale advantages compared to neural network models. This work will draw on and contribute to the ChiSight and ChiExpert projects.
- Prototyping initial versions of two new collaborative projects, drawing on OpenGen and collaborations across the Quest. Specifically, developing prototypes of ChiMind, the first boundedly-rational agent models of human minds and actions in digital environments, and ChiBrain, the first models that simultaneously predict human inferences, brain

activity, and neuroanatomy, tested in domains drawn from the Perceptual Intelligence, Developing Intelligence, and Embodied Intelligence Missions.

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, currently focused on the table-top scale.

At the "table-top scale", the PIs have formulated "Find the Grape", a set of manipulation problems under partial observability 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.
- Improving and implementing robot infrastructure, and developing novel planning algorithms for efficiently handling uncertainty and reasoning about unobserved space.

Human studies suggest that simple perceptual properties do not account for their action choices in 3D situations. Project researchers have devised computational models, based on volume estimation and Bayesian reasoning, that more accurately account for human reasoning, although humans also do not seem to be exactly rational. Work is underway on implementing several other computational strategies that more closely match the humans' perceptual input, with a paper likely to arise soon from this work.

In the robotics area, the group has designed and implemented an artificial "cognitive architecture" that integrates perception, world-state estimation, planning, and control. This serves as a locus for combining the work of several PIs and increasing focus on building whole functional systems. With a new hardware platform expected to arrive in October, this architecture is poised to take on the same tasks as human subjects and allow more fine-grained comparison between artificial and natural intelligence in problems of spatial reasoning and decision-making.

In the area of human 3D object perception, we have shown that standard artificial neural networks fail to match human abilities to infer the 3D shapes of novel objects whereas 3D Light Field Network models most closely resemble human performance, suggesting that building in 3D inductive biases increases human-model alignment.

Developing Intelligence (DI): Scaling AI the Human Way

Lead PI: *Tenenbaum*
Collaborators: *Schulz, Saxe, Mansinghka*

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”
- A first large-scale study of the evolutionary and developmental origins of the human capacity to learn complex, structured “program-like” representations, comparing the performance of adults, children across ages, and macaque monkeys in open-ended spatiotemporal sequence prediction tasks with a range of computational models of learning including Bayesian symbolic program learning systems as well as artificial neural networks models.

Collective Intelligence (CI)

Lead PIs: *Malone, Almaatouq, Rand, Rus*
Collaborators: *Ahmed, Berinsky, Horton*
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. An important goal for the Mission is to identify similarities in how intelligence emerges in these different kinds of groups. To do this, the team is developing an “ontology of collective intelligence” — a body of theory that identifies a wide range of possible processes and performance prediction models for tasks such as group decision-making.

The team also develops and tests innovative examples of how humans and AI systems can work together: 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 design visual objects, and write short stories. This work is expected to help better understand both the science of collective intelligence and its application to important practical problems.

Language

Lead PIs: *Fedorenko, Andreas, Levy*

Large language models (LLMs) 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, using our understanding of human cognition to make models better, and using models as a tool to understand human language and cognitive processing. The team

- has developed a new evaluation benchmark (Elements of World Knowledge) that evaluates social and physical reasoning in LMs. This benchmark has been released publicly and is being incorporated into major LM evaluation projects, including Stanford's HELM.
- has developed new techniques for more developmentally plausible training of language models with perceptual grounding, enabling learning of word meanings from small amounts of language data. (This work received the best paper award at ACL 2024.)
- is developing a cognitive modeling framework, Rational Meaning Construction, in which utterances are interpreted as updates to, and queries against, a probabilistic model that is constructed on-the-fly to support reasoning about a situation described in language.
- is developing new algorithms for learning language interactively via targeted queries to a linguistic informant.
- has found that LLM representations capture human neural responses accurately enough to allow for non-invasive 'control' of the language circuits.
- has reviewed the state of the art and challenges in relating LLM representations to human neural responses to language.
- is designing a new fMRI benchmark (unprecedented in size) of responses to 10,000 sentences to be collected in 2024–25, with the goal of developing the most accurate encoding model of human responses to language to date and evaluating the importance of syntactic structure in aligning model representations to human neural responses.
- is designing a combined fMRI / model evaluation approach to investigating arithmetic reasoning in humans and machines.
- continues deep analysis of how LLMs represent fundamental grammatical generalizations in language, as well as the limitations of LLMs in this regard.

Perceptual Intelligence

Lead PIs: *DiCarlo, Tenenbaum, Mansinghka*

The Perceptual Intelligence Mission brings together computer scientists, cognitive scientists, and neuroscientists with a shared goal of producing the first machine executable models of human visual intelligence that work computationally, cognitively, and neurally.

Concretely, the interdisciplinary team aims to produce new machine executable models that take sensory data as input — any possible movie or sequence of images and achieve all three of the following: First, match or exceed the data efficiency, generalizability, and robustness of human vision, unlike current industry "AI " approaches that require large training datasets and struggle with difficult viewing conditions and often fail to transfer from the factory to the field; Second, these models are theoretically well-grounded in that they are provably constructed by design to implement sound approximate probabilistic inference in generative models that accord with core principles of optics and mechanics, with controllable tradeoffs between perceptual latency and error rate, for fixed-size circuit implementations, unlike industry neural network architectures for vision; And third, these models have internal components that simultaneously align with internal brain responses measured across the multiple brain areas in primates (including humans) that are thought to support scene understanding and thus are the leading scientific models of the brain mechanisms of visual perception.

If this Mission is successful, it will:

- Produce theoretically grounded, open-source software and tools to catalyze new AI approaches that scale far better in data and compute than current industry approaches. This will allow transformative AI applications far beyond human visual capabilities that this Mission aims to explain.
- Unlock the ability to use the brain-aligned models of human visual processing and perception for biomedical applications. For example, the development of non-invasive methods for targeted, precision deep brain modulation.
- Enable other scientific teams to use the same approach to explain human auditory and tactile perception.
- Form the foundation of other machine executable models of other domains of intelligence that draw on the contents of visual perception, such as planning and navigation.

Administration

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 $4.1 million, and the MIT-Liberty Mutual Insurance Collaboration had $2.3 million in total secondary research volume.

Development

Over the past year, efforts have been made to strengthen relationships with the Quest donors and to work closely with central RD on strategy and engagement. As part of these efforts, the Quest held a donor dinner in Cambridge and a reception at the Greenwich Country Club for MIT alums and friends who are interested in supporting the Quest. Several donors have committed initial gifts to establish the Fund for Future Leaders in the Field.

Education

In August 2023, CBMM hosted the tenth 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 enrolls 35 students who are advanced PhD students and postdocs from the top universities around the world.

The Quest trains and mentors undergraduates interested in neuroscience, psychology, and software engineering through UROP. This year, we supported eight fall UROPs, four IAP UROPs, five spring UROPs, and we are currently supporting three summer UROPs (one of whom is a volunteer). A past UROP, Khaled Shehada, worked briefly with the engineering team after completing his MEng last summer; he is now working for Meta.

The Quest will continue to support 50% of a postdoc fellow working in SERC in SCC, Kevin Mills, through the end of his appointment in June 2025.

Events

The Quest provided financial and organizational support to events and student activities to benefit the AI community at MIT and beyond. In October, CBMM celebrated their tenth anniversary with a two-day workshop, attended by friends and colleagues from around the world. The Quest and CBMM co-sponsored seminar series hosted nine speakers, representing industry and academia in neuroscience, cognitive science, and computer science. We continued to organize “Research Meetings” and “Mission Updates” at which members of our community gather in an informal setting to present and discuss recent areas of effort.

Communications

Over the past year, communications efforts have focused on connecting with our community, at MIT and in the field. Improvements have been made to our website to allow us to provide better content about researchers and their work. The Senior Communications Officer continues to work closely with the Senior Development Officer 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.

Personnel

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, Financial Officer; and Valerie Hoke, Senior Administrative Assistant. The Quest Software Engineering Team is made up of Sam Winebrake, Deirdre Kelliher, and Ethan Pellegrini, who work closely with new hires Tiffany Luong, Research Scientist; Esther Alter, Intelligence Observatory Software Engineer; Michael Lee, Intelligence Observatory Research Scientist; and Mike Ferguson, Research Software Engineer for the Brain-Score Platform. Kartik Pradeepan has accepted a position as Brain-Score Research Scientist, starting July 1. CBMM staff are Kathleen Sullivan, Managing Director, and Kris Brewer, Director of Technology. During this year, Brian Pierson was promoted to Financial Officer. Two positions are currently vacant: Lead Software Engineer and Events and Projects Coordinator.

Future Plans

The Quest's goals are to continue to build on the foundations of our Missions and to develop machine-executable models of core aspects of natural intelligence that work computationally, cognitively, and neurally. We will also continue to 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.

In the near future, we will acquire equipment to outfit the Intelligence Observatory. Plans include a VR laboratory, a humanoid robot, and a Spot robot dog from Boston Dynamics. These will be valuable resources for the Embodied Intelligence Mission and other Missions.

In November, we will host Matter of Minds: Building the Science of Natural and Artificial Intelligence, a one-day event that will share the Quest's work with alumni/ae and donors.

In the coming year, we will continue to pursue funding opportunities, build strategic partnerships, increase our efforts in fundraising and stewardship, hire and promote software engineers, and align with MIT's goals. Steps are being taken to reorganize the corporate program, and we hope soon to hire a staff person to manage that endeavor.

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