The MIT Quest for Intelligence
Report to the President 2021–2022

Over the past year, the MIT Quest for Intelligence (the Quest) has executed a significant metamorphosis of the entire organization, in its focus, its administration, and its staffing. At this point, the Quest community is stronger than it has been since its launch in 2018 and is well poised to make an impact on research breakthroughs and applications at the interface of Natural Intelligence and Artificial Intelligence (AI) for years to come. In summer 2021, we began funding Quest Missions that bring together interdisciplinary teams of researchers focused on specific areas of inquiry; in fall 2021, the Quest leadership team was reorganized; and in winter 2021, the Quest incorporated MIT's Center for Brains, Minds, and Machines (CBMM), a closely aligned NSF-funded science and technology center focused on the interdisciplinary study of intelligence and how it can be replicated in machines.

Leadership and Affiliated Researchers

James DiCarlo, Peter de Florez Professor of Brain and Cognitive Sciences (BCS), is the Quest Director, and Nicholas Roy, Professor of Aeronautics and Astronautics, is the Director of MIT Quest Systems Engineering. CBMM is led by Professor Tomaso Poggio, the Eugene McDermott Professor (BCS).

Quest researchers represent units across MIT:

- Schwarzman College of Computing (SCC), Electrical Engineering and Computer Science Department (EECS): Asuman Ozdaglar, Deputy Dean of Academics, SCC, head of EECS; Leslie Kaelbling, Panasonic Professor; Associate Professor Vivienne Sze; Pulkit Agrawal, Steven G. (1968) and Renee Finn Career Development Professor; Jacob Andreas, X-Window Consortium Career Development Professor; Jesus del Alamo, Donner Professor; Joel Emer, Professor of the Practice; Associate Professor Phillip Isola; Tomas Lozano-Perez, School of Engineering Professor in Teaching Excellence; Professor Aleksander Madry; Professor Russ Tedrake.
- 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, Department Head; John Gabrieli, Grover M. Hermann Professor; Associate Professor Mehrdad Jazayeri; Professor Roger Levy; Vikash Mansinghka, Principal Research Scientist; Professor Alexander Rakhlin; Mriganka Sur, Newton Professor; Professor Joshua Tenenbaum; Assistant Professor Guangyu Robert Yang.
- Sloan School of Management: Assistant Professor Abdullah Almaatouq; Thomas Malone, Patrick J. McGovern Professor; Professor David Rand.

- Departments of Nuclear Science and Engineering (NSE) and Materials Science and Engineering (DMSE): Professor Bilge Yildiz; Professor Ju Li.
- Economics: Daron Acemoglu, Elizabeth and James Killian Professor.
- Media Arts and Sciences: Associate Professor Cynthia Breazeal.
- Harry Lee, Advanced Television and Signal Processing Professor of Electrical Engineering.
- Political Science: Professor Adam Berinsky.

Research

MIT Quest Research aims to understand intelligence — how brains produce it and how it can be replicated in artificial systems. The Quest approaches this as a single grand challenge requiring the organized, collaborative efforts of science, engineering, the humanities and beyond. Researchers from labs, centers, and academic departments across the Institute are involved in research in these areas, most notably CBMM.

To execute on its vision, Quest research is organized around "*Missions,*" long-term collaborative projects that are rooted in foundational questions in intelligence—animal, human, or collective. Each Quest Mission supports research organized around a domain of natural intelligence. New models of intelligence proposed along the way will be integrated into theoretical or computational platforms designed both to test the models within a coherent system of intelligence by assessing the system's ability to explain and account for phenomena of natural intelligence in the domain (and thus gauge overall progress) and demonstrate the utility of the system in solving real problems of perception, reasoning, or decision-making on campus or in the world at large. In 2021, the Quest made a request for proposals to specific research teams that align with the vision of solving foundational questions of intelligence in engineering terms; this RFP resulted in 31 proposals from across the Institute. In September 2021, $2.6M in funding was committed for nine of these proposals for incubation at a range of funding levels (four at the Mission level, and five at the seedling level, with renewals planned roughly annually). The Quest provides institutional support, guidance, and protocols to support the development of each Mission and seedling project.

Center for Brains, Minds, and Machines (CBMM)

Now entering the final year of its NSF funding, CBMM's research, educational, and outreach programs have become an important part of the MIT environment. CBMM researchers and collaborators come from MIT and several affiliated universities (https://cbmm.mit.edu/about/people). The center has meaningful relationships with industrial partners — DeepMind, Google, IBM Research, and Siemens — and researchers in federal laboratories. Research collaborations have continued with Google DeepMind, Fujitsu Laboratories, and IBM. Representatives from these partners and from other industries have been invited to scientific workshops and symposia held throughout the year.

CBMM's speaker series hosted fifteen speakers, including neuroscientists, cognitive scientists, and computer scientists representing visitors and units across the Institute.

In August 2021, CBMM's eighth annual summer course was attended by 20 students representing universities around the world, attending a three-week intensive course on the intersection between biological and computational aspects of learning and vision and on high-level social cognition and ÅI, as well as audition, speech, and language processing, taught by CBMM faculty and colleagues from both academia and industry. The ninth summer course will be this August, on the campus of the Marine Biological Laboratory in Woods Hole, MA.

Quest Missions:

<u>Embodied Intelligence Mission (EI)</u>
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 "tabletop" 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. They have made progress on infrastructure in three settings:

- Designing and beginning to fabricate apparatus for monkey experiments
- Creating a suite of simulated domains, which includes those that will be used for monkey experiments, that can be interacted with by humans in VR and by robots in simulation.
- Improving and implementing robot planning algorithms for handling uncertainty and reasoning about unobserved space.

Scenarios in all three modes (monkeys, humans, robots) will run in early Fall 2022. Simple pilot experiments will reveal what is feasible in each mode and enable design of a more systematic set of experiments and evaluation / comparison metrics, to have real comparative data and new insights next year.

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.

<u>Developing Intelligence: Scaling AI the Human Way</u>
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. Can we 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. The team has made multiple advances, including: creating online platforms for automated, scalable developmental experiments with infants and young children that are now being used by dozens of labs around the country; the development of probabilistic programming systems in robotics 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; and the first computational models of curiosity in infants and adults that quantitatively predict how visual engagement depends on novelty and complexity of stimuli, recognized with a Cognitive Science Society Outstanding Computational Modeling Paper award.

<u>Collective Intelligence (CI)</u>
Lead PIs: *Malone, Andreas, Madry*
Collaborators: *Rus, Almaatouq, Acemoglu, Berinsky, Ozdaglar, Rand, Athulya, Fedorenko, Levy*
This Mission focuses on one of the most important types of intelligence: the collective intelligence that arises in groups of individuals, whether those individuals are people, computers, animals, and/or other kinds of entities. The research addresses optimizing human-AI group decision making, creating super-intelligent groups that operate effectively in dynamic environments, and how highly polarized groups make decisions. The PIs are defining benchmarks to define superintelligence, determine character types and evaluation performance of different tasks and groups. Their goals include developing modeling tools for group design, creating a group configuration library, and establishing predictive models to evaluate performance. This Mission incorporates several projects: how to create "superintelligent" groups of humans and computers; how groups of humans and/or computers communicate meaning using language; and how to reduce the ways AI-powered social media can make societies less intelligent.

Creating Superintelligent Groups: this project aims to understand CI well enough to create innovative—and dramatically better—forms with an initial focus on creating superintelligent human-computer groups. The team has developed a test—analogous to the Turing Test—for quantitatively evaluating the ratio of performance improvement in human–computers groups relative to humans alone, computers alone, or other baselines. They developed an innovative method for writing software using a combination of humans and GPT-3, a massive, state-of-the-art AI system. Using this method, programmers generate software about 30% faster; more interestingly, non-programmers, incapable of writing software on their own, could do so about as quickly as professional programmers. They have also applied a similar method to generate short romance stories and are experimenting with prototype versions of human–computer systems for visual design. They will next develop and evaluate more human–computer systems for performing tasks, forming the beginnings of a coherent theory about how different kinds of human–computer groups perform different types of tasks, and extending this work into a computer-aided design (CAD) tool that helps generate and evaluate ideas for configuring collectively intelligent groups. This activity has potential to unify and accelerate the whole project.

Understanding Meaning in Language Models and Minds: 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. They are important in silico models of human language processing. 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. This research aims to provide empirical and theoretical foundations by deeply understanding meaning representation, world knowledge, and reasoning in computational models, and using our understanding of human cognition to make them better. The team has designed a suite of language modeling benchmarks targeting key aspects of world knowledge and reasoning including physical situations, events and causality, and theory of mind. Motivated by these results, they have developed a new modeling framework, in which language models perform next-word prediction by generating short probabilistic programs encoding the meaning of user requests, then executing these to generate results.

Taming AI—Analysis and Remedies of the Unintended Consequences of AI-Powered Social Media: Work includes modeling user behavior (belief shaping) and then validating this model; and building an experimental infrastructure for testing behavioral aspects of human decision making (information sharing). Social media platforms generate personalized content for each user, and data show that users' preferences evolve and are shaped by their past social media content. The content a user sees today may affect what they engage with tomorrow. Using both analytical results and simulations, the team examines how platforms inject content from external outlets to maximize engagement, how this content affects short-term and long-term beliefs of users, and how the injected content depends on the available content sources and structure of social networks. They have learned that patient platforms are more incentivized to bring users to consensus; that, when platforms have access to extreme sources, they are incentivized to use them; and that network structure matters, with more connected networks yielding consensus and less connected networks yielding homophily. They are also investigating the sharing behavior of social media users after reading about political ideology sharing misinformation.

<u>Brain-Guided Intelligence Hardware</u>
Lead PIs: *Yildiz, del Alamo*
Collaborators: *Fee, Li, Sze, Lee, Emer*
Electrochemical synapses have potential to enable a new domain of hardware for machine intelligence. These devices work with a charge-controlled cation-intercalation mechanism and present very high energy efficiency. In addition to powering analog neural networks, they can serve as building blocks in novel "brain-guided" machine intelligence architectures, representing the learning rules being developed in neuroscience. The architectures in return can advance the understanding of learning rules in the brain. This Mission seeks to answer the questions: what are

the intrinsic dynamics expected from these electrochemical synapses, how can these dynamics be utilized to reflect the learning rules in biological brains, and how to design materials, devices, and circuits to realize brain-guided machine intelligence platforms? This team is first examining the rule underlying how songbirds learn to sing, developed by Fee.

The Yildiz and Li group has developed an equivalent circuit model for the electrochemical synapses to understand their intrinsic dynamics, and to implement an ionic control circuit around the synapses in line with the songbird learning rule. The model captures the ion transport, charge transfer, and dopant redistribution processes in the devices, with the goal of predicting behaviors of the devices from various excitations. It also provides guidelines for the choice of ions, material systems, and device structures to achieve various dynamics on different time scales. Next the team will apply the model to design electrochemical synapses that implement learning rules to emulate the way songbirds learn to sing. The del Alamo group experimentally studied the dynamics of these electrochemical synapses based on proton intercalation using PSG as electrolyte. These CMOS- and BEOL-compatible nanoscale devices exhibit excellent characteristics in terms of high operation speed, high energy efficiency, many nonvolatile conductance states across a large dynamic range, linear and symmetric modulation, and high endurance. The Yildiz, Li, del Alamo, and Fee group studied potential device structures and circuit designs that use electrochemical synapses to perform delayed coincidence detections that are critical in emulating a vocal behavior in the songbird.

Among specialized computing hardware, Analog Neural Networks (ANN) most closely approach the operation of the brain because they compute "in memory", where the data that describe the model reside. Also, ANNs process signals in the analog mode that span a continuum of values. Further, "neurons" in ANNs communicate with many other neurons through programmable analog weights. This parallelism has raised the hope that ANNs can demonstrate similar accuracy as Digital-DNNs on many AI tasks with much lower energy consumption and latency.

Seed funding
In addition to the Missions, the Quest also provided funding to five early-stage projects, expecting that they may dovetail with existing or evolve into new Missions.

<u>Building and evaluating multi-system functional brain models</u>
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, including visual, auditory, cognitive, episodic memory, motor control, reinforcement learning systems, etc. Substantial progress has been made on two critical components of this project: high-throughput evaluation of many models on many experimental datasets and building multi-system network models for elemental cognition.

A high-throughput pipeline rapidly and scalably compares 100–1,000+ network models with ~10 datasets of neural activity data from prefrontal and other areas when monkeys or mice perform cognitive tasks. Using this pipeline to test what kinds of neural network models can better explain experimental data, PIs have identified important aspects for building brain-like neural network models, including the use of realistic output structures.

They will build multi-system neural network models of elemental cognitive functions, each containing at least one sensory processing system and an executive cognitive system. Models have been built for a wide variety of tasks probing working memory, a fundamental elemental cognition. The networks reproduce a wide range of behavioral and neural findings from both classical and notable recent experiments.

Yang was awarded the Searle Scholarship (~15 early career faculty/year throughout biomedical sciences) based on this research program and preliminary results, which have been presented at various conferences and invited talks.

The project's goal is to build neural network models consisting of multiple (3+) systems with modules for thalamus, hippocampus, basal ganglia, etc., to be compared to key experimental datasets across a much wider range of brain functions, and in close collaboration with domain experts from many subfields.

<u>Bootstrapping Emergent Intelligence</u>
PI: *Isola*
This effort to understand the genesis of intelligence, bootstrapped from simple rules interacting in complex environments made progress on multiagent reinforcement learning, emergence in language models, and learning from procedural and self-generated data.

Unified Automatic Control of Vehicular Systems with Reinforcement Learning: Deep reinforcement learning (DRL) is "model-free" and flexible. Nonetheless, applying DRL commonly relies on painstaking specialization hindering the discovery of emergent behaviors in diverse multi-agent traffic scenarios. This work contributes a streamlined DRL methodology for vehicular microsimulation and may discover high-performance control strategies with minimal manual design. The methodology shows empirical improvement, typically 15–60% over no intervention, across diverse scenarios with connected and automated vehicles in the presence of human drivers.

Characterizing emergent behavior in pre-trained language models: A new dataset contains a large number of procedurally generated, compositional tasks, characterizing the relationship between generalization, memorization, and calibration in few-shot LM learning. Their results offer first steps towards explaining that in-context learning uses a batch algorithm, with a regularizer, that performs improper learning and has hard lower bounds on both the depth and width of the network required to support it.

Investigating procedurally generated data in robotics: Pre-training using noise and knowledge of self-motion leads to moderately successful navigation policies. Strategies for procedural generation of environments and goals help robotics systems achieve superior performance when dealing with competing constraints. The team's real-world robotic system surpassed the performance limits of prior work while maintaining robustness to wide environmental variations.

A platform for multiagent intelligence and artificial life: The team made progress on a Neural MMO platform, introducing more resource types, tools, equipment, and a trade system.

Reimagining Reinforcement Learning via Memory-Based Biological Learning
Lead PI: *Agrawal*
Co-PIs: *Rakhlin, Sur, Fiete*
This team aspires to understand the role of memory in animal cognition and how it enables rapid multi-task learning without forgetting previous tasks and then to build decision-making algorithms that learn online and can quickly solve new tasks. Their computational models may explain animal behavior in rich naturalistic settings, offering tools to study and mitigate memory loss and its consequences. There has been significant progress since funding began.

- Advances in Exploration–Exploitation: Rakhlin's lab is developing a theoretical foundation for sequential decision-making, a learning paradigm that includes reinforcement learning. Agrawal lab has developed an algorithm that trades off exploration bonus against task reward to achieve state-of-the-art results across many standard benchmarks.
- Exploration via Inductive Extrapolation: Fiete lab developed approaches for fast exploration of new spaces by identifying regularities in a given environment. In collaboration with Agrawal lab, they developed a new computational model for exploration that leverages long- and short-term memory for quick learning.
- To complement the theoretical framework from the biological perspective, Agrawal, Sur, and Rakhlin labs are exploring a novel mechanism of exploration–exploitation in the brain.

Based on these and other results, some of the PIs are looking for additional funding to support future efforts.

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. While persistent neural spiking activity is believed to be a well-established neural correlate of WM, recent observations suggest that neurons involved in WM maintenance fire in sparse pattern instead of firing persistently, and perhaps WM maintenance may involve populations of neurons. This project is investigating the formation of working memories in a recently introduced model of brain computation based on assemblies of neurons, a computational model at a level of abstraction between models of neurons and whole brain models. While modern deep neural networks are inspired by models of the sensory parts of the brain, especially the visual cortex, assemblies of neurons model intermediate levels of computation beyond sensory information processing and can represent memories and cognitive concepts.

The researchers performed a delayed-match-to-sample task on a neural assembly model with Hebbian plasticity by presenting it with images and computing the distance between neural trajectories and found similar results to the neural trajectories in the experiment in which a monkey is shown a sample object and must choose the matching object after a variable length delay.

While preliminary results are promising, further efforts will investigate the robustness of the neural assembly model to the ablation of synaptic connections. \

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

PIs: *Breazeal and Gabrieli*

This project studies how to responsibly design AI technologies that enhance human–human interactions and relationships at home, particularly to support young children's learning with parents in a home setting. High-quality, reciprocal parent–child interactions play an essential role in children's development. Many parents do not know to engage children in this way, particularly in lower socio-economic communities. Robot-assisted parent–child interaction could be a viable future strategy for promoting children's in-home education with parents. This new robot–child–parent design direction provides inspirations to develop social robots capable of fostering long-term parent–child social and affective interaction at home.

The researchers investigated this question through "lived technology experiences" and interviews. They deployed and remotely teleoperated a social robot in the homes of 12 families with 3–7-year-old children to engage in six 25-minute sessions of triadic story reading with both parent and child. Before and after deployment, semi-structured interviews were conducted with participants regarding their triadic interaction experience and desired robot design features. Qualitative analysis showed that social robots in parent–child co-reading at home benefits children's comprehension and interest in reading, improves parents' participation behaviors, and diversifies parent–child interaction styles.

The team proposed design guidelines to launch a long-term, multimodal parent–child–robot interaction in family homes with a goal of understanding what and how various robot interaction strategies facilitate parent–child social and affective interaction during co-reading.

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 and collaborate with the MIT-IBM Watson AI Lab, the MIT Energy Initiative, MIT Mobility Initiative, and the MIT Climate and Sustainability Consortium.

In 2021–2022 our total spent fund volume was $2.8 million, and the MIT-Liberty Mutual Insurance Collaboration had $1.6 million in total spent secondary research volume.

Community Outreach and Activities

Education

The Quest trains and mentors undergraduates interested in neuroscience, psychology, and software engineering through UROP. We sponsored 8 fall UROPs, 11 spring UROPs, and 11 IAP UROPs.

The Quest is supporting 70% of a post-doc fellow working in SERC in SCC, Miriam Boulicault.

Events
We provided financial and organizational support to events and student activities to benefit the AI community at MIT and beyond. In April 2022, CBMM and Quest co-sponsored "Using AI to Accelerate Scientific Discovery," a special lecture from Demis Hassabis, the Founder and CEO of DeepMind, which was open to the public, streamed, and recorded. Over the course of the year, we collaborated with the AI at MIT Club to sponsor several sessions for undergraduates to learn about AI research at MIT and about Quest UROPs. Also, PIs share ideas and build collaborative relationships at monthly Quest social events.

Communications
Redesign of the Quest website began in January 2020, and the new site launched in summer 2021. We continue to refine and add new features, which will include a repository of past and current engineering projects, with links to their code or other resources.

Administrative Staff
For several years, the Quest shared some administrative and support staff with the MIT-IBM Watson AI Lab. As of the end of this fiscal year, that practice has ended, and all current staff are fully committed to the Quest and its work. Over the past year, the leadership team has carefully considered staffing responsibilities and how best to serve the Quest's needs, and they developed a hiring plan, promoted some individuals, and reassigned some duties. The current staff are Erik Vogan, Executive Director; Rachel Kemper, Communications Officer; Brian Pierson; Financial Officer; and Allison Provaire, Project and Events Administrator. Katherine Fairchild leads the Software Engineering Team of Dhaval Adjodah, Kyle Keane, and Martin Schrimpf. CBMM staff who are now working with the Quest are Kathleen Sullivan, Managing Director, and Kris Brewer, Director of Technology. Transitions during this year include promoting Erik Vogan from Managing Director to Executive Director, and promoting Katherine Fairchild to Engineering Team Lead; Gilbert Cordova, Administrative Officer, accepted a new role in another department at MIT and will leave the Quest on July 1; Rachel Kemper and Brian Pearson joined the staff, coming from other positions at MIT, and Martin Schrimpf joined the Quest after completing his doctorate in Brain and Cognitive Sciences. There are three open positions: Senior Development Officer, Missions Project Manager, and Senior Administrative Assistant.

Future Plans
The past year has seen enormous progress as the Missions have coalesced around common, shared understanding of the research problems to be addressed and a common language for describing the experiments to be undertaken. While we understood it would be challenging to bridge disciplinary divides, for neuroscientists to learn to speak with computer scientists, and for cognitive scientists to dialogue easily with roboticists, we underestimated the challenges of developing shared nomenclature and mental models in the aggregated research teams. These efforts and our patience have begun to pay off, with several Missions embarking on interdisciplinary research that would not have been possible without this period of development. This highlights the need to fund these initial efforts with patient capital, most commonly in the form of unconstrained philanthropic funding, as the impact of such interdisciplinary research is likely to be significant, far beyond the traditional single-lab approach, but may take longer to show progress as new research approaches are sought. Growth and development are also apparent in the seed funding Quest provided, with at least one seed-funded project likely being

promoted to a Mission during the next year, and additional seedling funding is being contemplated.

In May, the Quest submitted a proposal to the NSF to establish a National AI Research Institute, led from MIT, but partnered with 10 other institutions and universities. We will learn if the proposal is successful in late 2022. If funded, the Institute will create a program of nationally distributed teams bringing together more than 40 internationally recognized research leaders from computer science, cognitive science, and neuroscience and a wide range of trainees to address a single grand mission: a scientific understanding of human core embodied intelligence, in engineering terms. The Institute's core premise is that while many natural sciences use AI, scientific advances in neural-cognitive systems are the foundations of future AI. We will pursue integration of neural and cognitive levels through a commitment to engineering AI models that simultaneously function as joint neural-cognitive models: they implement—and thus explicate—key cognitive capabilities in those scientific models with neurally-mapped components. This unified scientific approach to understanding the foundations of intelligence seeks not only payoffs in AI, but also payoffs in health care and education.

The first few years of the MIT Quest for Intelligence have not been without challenges, many of which are faced by any new organization, particularly one that is founded with ambitious goals and that includes a remarkably diverse group of investigators. We have learned from our efforts and have made adjustments. With the refinement of the Quest vision, it became clear that the Quest corporate program would need to adapt to support that vision. The past year has seen productive discussions, both internal and external, that will lead to a relaunch of that program in the coming months. In 2023, we expect to relocate to the new SCC building. Returning our offices to the main MIT campus will bring about improved interpersonal communication, revitalize our relationships with UROPs and graduate students, and spark new opportunities for collaboration.

Given recent enormous challenges, not just for the Quest but for MIT, the nation, and the world, we are incredibly encouraged by our progress this year, and look forward to a future where a fully developed understanding of natural intelligence enables its replication in engineered systems to the benefit of both MIT and the world.

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