A National Science Foundation Science and Technology Center

Center for Brains, Minds and Machines (CBMM)

First Annual Report
Award: CCF-1231216
September 1, 2013 – May 30, 2014
Submitted to NSF: May 30, 2014

Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
Phone: 617-253-0551; Fax: 617-253-2964
http://cbmm.mit.edu/

Table of Contents

I. General Information

II. Research

III. Education

IV. Knowledge Transfer

V. External Partnerships

VI. Diversity and Broadening Participation

I. GENERAL INFORMATION

1a. Primary Partner Institutions and Contact Information

Date submitted	May 30, 2014
Reporting period	September 1, 2013 – May 30, 2014
Name of Center	Center for Brains, Minds and Machines (CBMM)
Name of Center Director	Tomaso A. Poggio
Lead Institution	Massachusetts Institute of Technology (MIT)
Address	McGovern Institute for Brain Research at MIT 77 Massachusetts Avenue, Cambridge MA 02139
Phone Number	617-253-0551
Fax Number	617-253-2964
Email Address of Center Director	tp@ai.mit.edu
Center URL	http://cbmm.mit.edu/
Institution 1	Harvard University
Address	Dr. L. Mahadevan Harvard University Pierce Hall 29 Oxford Street Cambridge, MA 02138
Phone Number	617-496-9599
Fax Number	
Email	lm@seas.harvard.edu
Role of Institution at Center	Harvard provides leadership on the Technology and Knowledge Transfer Program

1b. Biographical Information for New Faculty

Biographical information for the following new faculty is attached as Appendix A.

Ellen Hildreth, Wellesley College

Lizanne DeStefano, UIUC

1c. Annual Report Primary Contact Person

Name of Individual	Dr. Kenneth Blum
Center Role	Managing Director
Address	Center for Brain Science Harvard University, Northwest Building 52 Oxford Street Cambridge, MA 02138
Phone Number	617-496-1478
Fax Number	617-495-0524
Email Address	kenneth_blum@harvard.edu

2. Context Statement

From its inception, the main goals of the proposed Center for Brains, Minds, and Machines (CBMM) were to better understand human intelligence by determining how the brain and mind perform intelligent computations, and how to make smarter machines. Our vision of the Center's research at the time of our proposal integrated cognitive science, neuroscience, computer science, and artificial intelligence. We believed that there is no greater scientific and engineering mission than understanding intelligence, because any progress toward improving intelligence in brains and machines can contribute to other great challenges in science and society, both theoretical and practical.

These beliefs have remained intact and actually grew stronger because of specific developments in our own research and in the commercial world of high-tech companies (see later). In our Strategic Plan drafted last February we distilled these beliefs into the following Vision and Mission statements:

Vision Accumulated knowledge and technology, now in place, enables a rapid leap in our scientific understanding of intelligence and our ability to replicate intelligence in engineered systems.

Mission We aim to create a new field by bringing together computer scientists, cognitive scientists and neuroscientists to work in close collaboration. The new field – the Science and Engineering of Intelligence – is dedicated to developing a computationally centered understanding of human intelligence and to establishing an engineering practice based on that understanding.

In support of our mission, we have developed specific goals and plans in four main areas: Collaborative Research, Education, Human Resources/Diversity, and Knowledge Transfer.

- Our scientific research goal is to discover how intelligence is grounded in computation, how those computations develop in childhood, how those computations are implemented in neural systems, and how social interaction amplifies the power of those computations. As we progress, we will aggressively pursue opportunities to discover and develop unifying mathematical theories. To encourage collaboration, we will jointly develop top-to-bottom computational models powerful enough to explain visually-perceived situations the way humans do, by answering questions about who, what, why, where, how, and why, as well as with what motives, with what purpose, and with what expectations. In support of visual understanding, we will develop computational models of what children know and learn about physical objects and intentional agents and how they learn so much so rapidly; we will develop computational models of learning, reasoning, memory, and concept formation that are consistent with behavior, neural systems, and neural circuits; we will develop computational models that enable computers to think new thoughts, imagine scenes, form hypotheses, propose interventions, and compose narratives; and we will develop new methodologies and new technologies.

- Our education goal is to ensure that our new knowledge is packaged in accessible ways, including model courses at the graduate and undergraduate levels.
- Our diversity goal is to ensure that our new field is broadly inclusive.
- Our knowledge transfer goal is to ensure that our new knowledge is quickly and broadly disseminated and brought to bear on the great challenges of the 21st century, so as to serve the people of the nation and the world.

We now provide some more detailed context for the first seven months of our Center, focusing first on changes and improvements we have made to the organization relative to the proposal.

Focus on Centerness and Collaborations

During and after the pre-award site visit in November 2012, we decided to emphasize our focus on collaboration and "centerness": we decided that the CBMM will exclusively fund collaborative projects that cannot be done in a single lab with typical single investigator grants but only in a Center like ours. In particular, we decided that no single PI will have students or postdocs funded by CBMM. Instead, projects will have to be collaborations between two or more PIs. The specific policy is that thrust leaders are in charge of hiring postdocs/students for the cooperative projects in their thrust—after consultation with the PIs in the thrust, the Research Coordinator (PHW) and the director (TP). All PIs understand that their funding level will likely go up or down as new opportunities emerge, interests change, or expectations come up short: we want our Center to remain a vibrant enterprise open to new projects and new researchers. We also pledged to share stimuli and data within labs in the Center, consistently with the idea of centerness. The external evaluation will assess the extent to which collaboration is thriving across the Center. Key metrics used to determine annual resource allocations and continued involvement with CBMM are: faculty engagement with CBMM activities (research, education, diversity, and knowledge transfer), and extent of CBMM collaboration (exchanges, presentations, publications, external funding).

New initiatives in Education, Diversity, Knowledge Transfer

We are focusing on a small number of high priority projects that cut across the three areas of Education, Diversity and Technology Transfer:

1. Summer School. A two-week summer school at Woods Hole on the Science and Technology of Intelligence, May 28–June 12, 2014; see http://cbmm.mit.edu/event/summer-course-mlb-woods-hole-brains-minds-machines/. We received more than 130 applications for 25 slots; interestingly all of the accepted applicants accepted our offer. Advanced graduate students from across the country will participate in lectures, hands-on sessions, and networking with the goal of building a professional community. The evaluation will track participants over time

to assess the impact of the Summer School on scholarship, professional networking, careers, and the development of the new discipline.

2. Technical Workshops. A small number (~ 2/year) of two-day workshops focused on a technical topic of central interest to the Center. One of the PIs proposes and organizes a session. The workshops include speakers from outside CBMM and are open to CBMM PIs and their students and postdocs. Our industrial partners are invited to participate as well as our international collaborators who may help organize occasional workshops in their home Institution. The first workshop on "*Learning Data Representation: Hierarchies and Invariance*" took place at MIT Nov. 22–24, 2013, organized by Lorenzo Rosasco and Tomaso Poggio with joint support from CBMM and one of our international partners (IIT in Genova, Italy); see http://lcsl.mit.edu/ldr-workshop/Home.html and http://lcsl.mit.edu/ldr-workshop/Schedule.html. Participants rated the quality and format of the workshop as "outstanding", indicated that the content was not readily available from other sources, and stated that they would readily attend similar workshops. The evaluation will continue to assess the extent to which the workshops promote networking and collaboration within CBMM and with industrial and international partners.

3. Workshops for our Diversity Partners. The first such workshop was organized by Mandana Sassanfar and Ellen Hildreth at MIT in January, 2014; see http://cbmm.mit.edu/wp-content/uploads/2014/01/CBMM_Report_Jan14_Workshop.pdf Faculty attended from all of the CBMM partner institutions for the broader participation of women and minorities, including Howard University, Hunter and Queens Colleges, the University of Puerto Rico at Rio Piedras, the Central University of the Caribbean, and Wellesley College. Together with CBMM faculty from MIT, Harvard, and Cornell, participants explored ways to create new research and educational opportunities for faculty and students in the emerging field of the science of intelligence. Participants are committed to develop and implement new undergraduate curricula, based on Science and Engineering of Intelligence, at their home institutions. The evaluation will monitor the development of coursework and implementation at partner institutions.

4. Summer Research Program (SRP). The first summer program will be offered in 2014 and will host 12–14 students from under-represented groups (women, under-represented minorities, non-traditional students). They will be placed in labs which are affiliated with CBMM or which have research relevant to the CBMM goals. Participants will be tracked longitudinally to assess graduate school participation and impact of the SRP experience. Faculty and graduate student engagement and satisfaction will be assessed.

5. New Industrial partners. We have added three new partners who are very interested in expanding their research in the new field of intelligence research represented by CBMM: Siemens, General Electric, and Schlumberger. We had meetings with exponents of them (for GE with the CEO, for Siemens with the director of the Princeton lab). We lost two of the small partners—both bought by Google since we started the Center last September: DeepMind and Boston Dynamics. This is another indirect sign that corporations feel as enthusiastic as we do about the technology of intelligence represented by CBMM and its members. We will track the growth and quality of industrial partnerships over time.

6. Website development. The CBMM website is our main platform for communicating within CBMM and with the outside at the level of Education (we currently have links to the material for 11 CBMM courses, including syllabus, slides and videos of classes), Technology Transfer, Outreach and Research. Web analytics will be tracked quarterly to assess usage and user feedback will guide development.

7. CBMM Memos. We have started a new series of CBMM memos, encouraging CBMM members to publish in original contributions, including white papers and preprints of publications, and thus begin establishing a growing body of research on the field of intelligence. For the motivation and the roots of this initiative, see http://cbmm.mit.edu/updates/announcing-new-cbmm-memo-series/

Measuring progress in research: The *CBMM Challenge*

Probably the most significant change we made to the original plan was first formulated in our response to the site visit comments. We sketched a new way for measuring our scientific progress during the course of the research, motivated by the IBM experience in developing Watson.

At the same time, a way to measure progress can be a way to motivate collaborations and team spirit. One way to encourage collaboration is to focus on a challenge problem that inspires good science, lays a foundation for future applications, enables measurable progress, and is broad enough to engage all participants.

Thus we devised a ***CBMM Challenge*** that spans all our thrusts: develop an account of the intelligence needed to understand images and video, and to build systems based on that account. Successfully addressing this challenge will require us not only to build systems that demonstrate understanding, but also to ensure that such systems are consistent with our evolving understanding of human behavior, brains, neural systems, and development.

To undergo the CBMM Challenge, a system has to answer a range of questions when presented with images and videos. Consider the image above. Representative questions range from those focused on basic recognition to explanation and anticipation of the sort that require models of the social and physical world.

- *What is there?*
- *Who is there?*
- *What are they doing?*
- *Are they friends or enemies?*
- *Do they trust each other?*
- *How, in detail, are they performing actions?*

- *Have you seen anything like this before?*
- *Why are they there? What will they do next?*

Of course human-like behavior is not enough: our computer models should also be human-like at the level of the implied physiology and development. Thus, the CBMM Challenge will be applied at several different levels: absolute (computational) performance, consistency with human performance/behavior, with human and primate physiology, and with human development. Tests at different levels will measure our progress. This will be one way—together with more standard metrics such as publications and evaluations—through which we will measure the progress of our overall project over time and in particular during the five years of the Center.

The series of questions is open-ended; we will rely on a set of stimulus databases and add to them during the next five years, keeping a separate test set. The initial problems—e.g., face identification/object recognition—will involve tasks that computers begin to do and for which we already have encouraging models and theories of how they are performed by the brain; the subsequent problems—e.g., models of how we can watch a video and say who is doing what to whom and why—are (ambitious) goals for the next five years of the Center or even beyond.

Each research thrust is expected to make progress on the challenge problem. For example, in the development thrust, the challenge is to create systems that represent objects in the same manner as a 3-month-old infant. In the neural thrust, there will be an analog challenge to develop models and theories which fit the physiological data on a specific region of the ventral stream.

The purpose of the CBMM Challenge is to measure progress over time and to stimulate collaboration aimed at understanding human intelligence, not to compete with commercial systems. In the engineering dimensions, our contributions are aimed at tomorrow, not today.

In summary, we will measure our progress over the life of the CBMM by evaluating:

- How well our systems perform.
- How consistent our systems are with human performance.
- How well our systems reflect models of human and primate behavior, neural systems, and neural circuits.

Databases for measuring progress in research

Given the CBMM Challenge (above) we decided that a top priority for the Center is to develop a set of databases (mainly images and videos) that will be used across different labs and techniques to measure performance of the mind and of the brain for recognition/perception of objects, of people, of interactions between people and objects, of people’s actions, and of people’s social interactions. The same data will be used to measure how well our models and our computer systems perform in absolute and relative terms. Gabriel Kreiman (leader of one of the thrusts) will be the responsible for coordinating this key effort. We will use these tests to evaluate progress of our work and our theories over the years.

Critical to the efforts in this scenario will be the rigorous evaluation of performance (by subjects, by computers, by neurons) in understanding non-visual social perception. To this end, the team will be creating, curating, annotating, and evaluating databases that will be shared with the research community at large as a “CBMM Challenge” for NVSP (Non-Verbal Social Perception). These efforts follow on related efforts with the recognition of objects and faces in static images in the computer vision community, but here we propose to extend these databases to more challenging and dynamic stimuli. There will be six main databases:

- In the work by Poggio and Sheinberg, investigators used three computer vision databases to evaluate the performance of the model for recognizing animal and human actions. Each of these datasets contains 5–10 actions and 50–100 videos per action. The datasets have been expanded by including a large number of video sequences from clips available on the Internet, which we have manually labeled for their content. This fourth database is the largest action video database to date, with 51 action categories and approximately 7,000 manually annotated clips extracted from a variety of sources ranging from digitized movies to YouTube (Jhuang, Serre, Poggio, 2011, ICCV).
- The database developed by Saxe and Kanwisher tests the computational models developed in Thrust 3 and Thrust 1 to account for specific key "social" tasks such as recognizing the expression of a face using gaze direction or hand pointing to associate an object image to its name.
- The database being developed by Kreiman consists of six episodes of a TV series ("24"), which will be characterized by using a combination of computational tools and manual annotation. The database consists of approximately 9,000 shots (with an average of ~ 45 frames per shot and ~ half a million frames total). These video shots are described in terms of the presence of specific characters and their viewpoints (33 different main characters and 4 viewpoints), objects (30 different objects), emotions (12 different emotional labels distinguishing between emotions depicted by the actors and emotions elicited in the viewer), locations, actions, and other features. Future efforts will expand this database to explore other aspects of social interactions. The database will also include eye-tracking information from viewers.

The same databases will be used to examine the underlying neural circuitry at different spatial scales from coarse (fMRI, MEG, subdural field potentials) to fine (single neurons), and at different temporal scales from seconds (fMRI) to milliseconds (MEG, field potentials, single neurons) in different species (monkeys, humans). The development of novel technologies for scalable, multisite, high-density recording (Boyden) will enable powerful studies of neural dynamics across these scales.

These databases will also provide rigorous quantitative benchmarks to compare the performance of computers and humans. fMRI and MEG work in humans and monkeys are ongoing together with field potential and single unit recordings in humans and single neuron recordings in the macaque monkey. Testing some of the theoretical predictions by manipulating circuits is problematic in humans. To this end, we are beginning to use optogenetics in monkeys, which could extend our understanding of the neural circuits and provide direct testing of some theoretical predictions. Without CBMM, sharing stimuli, databases, and algorithms, and comparing notes across species and techniques would be impossible.

The development of these databases is intrinsically collaborative and will lead to research efforts cutting across thrusts. Many of the databases discussed here are beginning to be used to compare notes across neural circuits, functional imaging, behavioral measurements and computational models.

Performance and management indicators

Management of a highly distributed Science and Technology Center is by nature difficult, so it is important from the beginning to be clear about what can be expected of Center management, especially with respect to evaluation principles. In the expectation dimension, the following particularly deserve mention:

- Management will be without a crystal ball. As the Center is attempting to do what has not been done before, we cannot say with absolute certainty that what we believe and expect today will be what we believe and expect tomorrow. Center leadership will have to manage through a changing landscape. No simple formulas constructed now would serve to guide us all the way through the next half decade, and, we hope, full decade. Accordingly, at the Center level, the director's responsibilities will include making strategic shifts of emphasis and funding among the thrusts, in consultation, of course, with the rest of the Center's leadership, the Center's external advisors, the thrust leaders and the Research Coordinator.
- Adherence to proposal promises is expected, but with the understanding that the promises will evolve so as to better align the Center with new opportunities as new opportunities emerge.
- Distributed decision-making is expected. The Center director will expect the Research Thrust Leaders to recognize new opportunities and, in consultation with the Research Coordinator, to make appropriate resource adjustments.
- Transparency of decision-making is expected. When opportunities and disappointments require funds to be moved by a Thrust Leader or the Director, then the Thrust Leader or the Director will coordinate with the Research Coordinator and carefully cite the reasons and principles guiding the decision.
- Stability of project funding is expected. The Thrust Leaders will be mindful of the need to commit to well-performing graduate students and postdocs for reasonable time periods.

In the evaluation dimension, we have collectively discussed, developed and refined evaluation principles over the course of our proposal-writing effort and up to the present time. Among these, five evaluation principles lie at the core of how participation will be evaluated:

- Contribution to the Center's objectives. We aim to better understand human intelligence, to make smarter machines, and to establish a new Science and Engineering of Intelligence. Thus, participants are expected to advance our understanding of how intelligence develops in early life, how it grounds out in neural hardware, how it works at a computational level, how it rests on social interaction, and how our understanding can be magnified via unifying mathematical theories.
- Collaboration within and among the thrusts. We believe that seminal contributions are most likely to emerge from collaborative efforts. To increase the likelihood of success, we will operate exclusively in terms of collaborative projects between and among participants, rather than in terms of efforts limited to the research group of an individual participant.
- Centerness. Our Center funds only collaborative projects that cannot be done in a single lab with typical single investigator grants. All projects—as a general rule—should be a key component of a thrust and pass the litmus test provided by the CBMM challenge. No single PI will have students or postdocs funded by CBMM. Instead, projects will have to be collaborations between two or more PIs. Thrust leaders have the responsibility of hiring postdocs/students for the cooperative projects in their thrust—with the help of thrust members, the Research Coordinator and the director.
- Community growing. We believe that our common objectives are best reached by establishing a new field of study. To further this end, we will work to encourage Center

participants, and especially our students, to have broad interests and participate energetically in all Center activities including Research, Education, Diversity, Outreach, and Knowledge Transfer.

- Commitment to diversity. We believe that diversity is intrinsically valuable. To seize opportunities for bringing diversity into our new field, all Center faculty have committed to contribute to at least two CBMM outreach activities per year.

Respect for contribution, collaboration, centerness, commitment, and community will be the focus of the Director's thinking, in consultation with the Research Coordinator, as he supervises the decisions made by the management team through our changing landscape. To promote transparency and evidence-based decision-making, CBMM is using an online annual reporting system (ARS) in which PIs, trainees, and staff report activities, products, and contributions using the five evaluation areas described above. Reports are generated by PI, thrust, and area (research, education, diversity, outreach, and knowledge transfer). The ARS provides the management team with extensive information on engagement, collaboration, and productivity to guide resource allocation and decision-making and promote fairness and transparency in management.

II. RESEARCH

1a. Research Goals and Objectives

The general goal of our research did not change from the proposal stage and is described best as the development of theories based on experimental data that allow a system to pass the CBMM Challenge. Three different versions of our goal were formulated during our Strategic plan session. They correspond to combinations of our five thrusts and reflect different levels of understanding the problem of intelligence: the computational level, the developmental level, and the level of the neural circuitry.

- A computational system, grounded in models of behavior, neural systems, and neural circuits, that interprets and describes visual scenes, and answers questions about them, the way humans do, thus passing a kind of visual "Turing Test."
- A computational system, grounded in models of child development, that constructs intuitive theories of physical objects and intentional agents as effectively as a child, using the same kind of information that is available to a child.
- Validated quantitative models of how knowledge, concepts, memory, learning, and reasoning are represented and processed in the brain.

1b. Research Performance Indicators

In our Strategic Plan the Center defined outcomes for its Collaborative Research Program and for each of the five Thrusts. At that time we also established performance indicators to assess progress in meeting our Collaborative Research goals and research goals we want to achieve in each thrust. The performance indicators include standard measures of publications, participation in conferences, co-authoring and collaboration, and especially measures associated with the CBMM Challenge. We also established a set of milestones for each thrust in the near and long term that we provide in Section 2a.

1c. Research Problems

We have not encountered any problems with respect to implementing our research plan during the current reporting period.

2a. Research Thrust Areas

Human intelligence can be defined in a variety of ways; central to it is the ability to acquire and apply knowledge in order to perform better in a specific environment and generalize to new situations. The definition of visual intelligence provided by the CBMM Challenge best captures the ultimate goal of our research: to understand the brain and to replicate the human mind in machines. We rely on a collaborative approach combining experimental techniques in neuroscience and cognitive science with computational modeling. This approach will also combine the study of the different aspects of perception, action, and cognition. We believe that real progress will come only from a thorough investigation of four aspects of intelligence: the integrative aspect of intelligence, its development, the wetware, and social intelligence, together with the development of a unifying mathematical framework. We have therefore constructed a research agenda organized into five major research thrusts:

- Development of intelligence
- Neuronal circuits underlying intelligence
- Integrating intelligence: vision, language, and social interactions
- Social intelligence
- Theory for intelligence

Our Research plans were detailed in the Center Strategic Plan. We briefly describe our plans and goals for each Research Thrust below.

Thrust 1: Development of Intelligence

The goal of Thrust 1 is to understand the roots of human intelligence by studying how it begins and how it develops in young children. The last two decades of developmental research, including seminal work by CBMM members, has suggested two fundamental insights.

Insight 1: How we start, or "The common-sense core". From the earliest ages, human beings organize their experience of the world around a basic understanding of physical objects, intentional agents, and their interactions. This knowledge is perceptually grounded—it guides visual as well as haptic, auditory and other perceptual modalities—but it is deeply conceptual, structured in the form of intuitive theories of common-sense physics and psychology. Much like scientific theories, but less formal and explicit, these intuitive theories comprise systems of abstract concepts such as forces and masses that explain objects' motions via principles of force transfer, or beliefs and desires that explain agents' actions via principles of efficient planning. These theories in turn build on core representations of basic aspects of the external world, such as space, time, and quantity.

Insight 2: How we grow, or "The child as scientist". Just as children's knowledge is deeply analogous in form and content to early scientific theories, so does it appear to grow in ways similar to how scientists revise and develop their theories of the world. In contrast to the mechanisms that underlie conventional machine-learning systems, children learn new knowledge not simply through finding patterns in big data sets, but through active hypothesis-testing and explanation-seeking, often starting from very sparse data. They actively explore, devise new experiments and ways to explore. They learn causes, not only correlations. They learn abstractions on multiple levels. They learn to learn, and use their new concepts to access new ways of thinking or ask new questions that were not even conceivable to them before.

Our goal is to capture these insights in engineering terms—in quantitatively testable and usable computational models—in order to better explain the phenomena of cognitive development, and to build more human-like machine intelligence and learning. This goal integrates tightly with all the other thrusts in order to advance our work on the CBMM challenge. The representations and algorithms used in our system for visual understanding in Thrust 3 will build in part on Thrust 1's models of common-sense physics and psychology, grounding our system in the visual scene understanding capacities humans have from early childhood. These models also provide a computational grounding for Thrust 4, in the form of ideal observer models for nonverbal social perception. Our models in turn build on formal approaches developed under Thrust 5, for building rich generative models using probabilistic programs, and for learning these models using methods of program induction and program synthesis. Finally, the aims of Thrust 2 to discover neural circuits underlying how we perceive physical objects and social agents, and how we learn about our environment via exploration, will illuminate the biological basis of the core cognitive capacities Thrust 1 studies, and will be guided in part by looking for the computational targets that Thrust 1 identifies. By the end of Year 1, we expect to have postdocs or graduate students working jointly between Thrust 1 and all other Thrusts of CBMM.

A priority for Thrust 1 in CBMM's first year has been to begin constructing a roadmap of cognitive development over the first several years of life: a coarse timeline of what knowledge and learning mechanisms are available to children at which ages, with a focus on knowledge about physical objects and intentional agents, and science-like exploratory learning mechanisms, as described above. To facilitate this goal, the PIs affiliated with Thrust 1 (Tenenbaum and Spelke, with the assistance of Schulz, as well as Susan Carey at Harvard) taught a special graduate seminar on computational models and cognitive development in Spring 2014, cross-listed between MIT and Harvard, with remote participants joining by Skype. This seminar met 3.5 hours/week and attracted approximately 25 participants (almost exclusively PhD students and postdocs, with a small number of undergraduates and others). Together, the students and faculty have begun preparation of three enduring sets of materials that will constitute the core of our roadmap:

- Detailed powerpoint presentations outlining the basic phenomena of cognitive development and multiple computational approaches, in the areas of (i) objects and intuitive physics, (ii) number, (iii) space, (iv) agents and intuitive psychology, (v) social relations, (vi) logic, (vii) language, and (viii) learning mechanisms.
- A database of approximately 150–200 papers from the developmental literature that capture the key empirical findings in the above domains, indexed by both domain and age-range.
- Roughly 15–20 proposals for specific experimental, computational, and joint experimental-computational studies that could be performed by CBMM participants over the next few years to test key aspects of the roadmap, with a focus on filling in holes in what we know about the cognitive capacities most relevant to the visual Turing test, when they develop and how they are learned.

In addition to these foundational materials, a wide range of specific research results have already been produced by members of Thrust 1, on their own and in collaboration with other thrusts. Some highlights include:

- We have built the first computational models of how aspects of intuitive physics may be learned from observations of dynamic events. These are formalized as hierarchical Bayesian models defined over probabilistic programs, in collaboration with Thrust 5, and they have been tested quantitatively in experiments with adults. A preliminary paper presenting these results (Ullman et al., 2014) is being presented this summer at the Cognitive Science conference, and a journal paper is in preparation. In Year 2 of CBMM, a new postdoc will be joining Thrust 1 (primarily in the Spelke lab at Harvard, but collaborating with the Tenenbaum and Schulz labs) who will help to develop touchscreen (iPad) methodologies to test these models in young children. We will also explore methods for testing them with infants using the Lookit system (described below) for online experimentation.

- We have developed preliminary models of several milestone stages in two domains of intuitive physics in 0–12 month old infants, inertial collisions and stability/support relationships, as well as the possible transitions between these stages driven by children's growing experience. These models use the hierarchical Bayesian program learning framework described above. In Years 2 and 3, we plan to develop more mature versions of these models, evaluate their ability to explain previously published results in the developmental literature, and test several novel experimental predictions of them.

- We have begun to develop a general account of children's reasoning about other agents that we call the naïve utility calculus. People explain and predict other agents' behavior using mental state concepts, such as beliefs and desires. Computational and developmental evidence suggests that such inferences are enabled by a principle of rational action: the expectation that agents act efficiently, within situational constraints, to achieve their goals. Here we propose that the expectation of rational action is instantiated by a kind of utility calculus sensitive to both agent-general (or objective) and agent-specific (or subjective) aspects of costs and rewards associated with actions. We show that children can infer unobservable aspects of costs (differences in agents' competence) from information about subjective differences in rewards (i.e., agents' preferences) and vice versa. Moreover, children can design informative interventions on both objects and agents to infer unobservable constraints on agents' actions. A conference paper on this work will appear at Cog Sci 2014, and a journal paper has been submitted. In ongoing work, we are developing paradigms to test these models in younger children, and to test them more quantitatively with adults.

- Theory of mind research has looked at how learners infer an agent's unobservable mental states from observable actions. However, such research has tended to neglect another observable source of data: the agent's emotional reactions to events. In particular, the agent's facial reactions might provide important information about her mental states that are otherwise ambiguous given her actions. We have begun building a Bayesian framework to explain these inferences, and tested this model in a behavioral study with adults asking them to use an agent's facial reactions to reason backward about her beliefs and desires. We found that participants' joint inferences of belief and desire from facial expressions were well-predicted by a Bayesian model analysis, based on integrating the likelihoods of the observed facial reactions and the observed action with their prior over mental states. In a paper to be presented at Cog Sci 2014 (Wu et al.), we argue that people's naïve theory of emotional reactions is structurally and causally intertwined with theory of mind in a way that allows forward prediction and backward inference.

- We have piloted a new experimental paradigm for studying children's exploratory learning. Studies of children's causal learning typically provide learners with clear evidence for direct causal relations, e.g., a machine that activates when a toy is placed upon it. But causal systems in the real world often present indirect perceptual evidence generated by interactions between hidden variables: Consider a child trying to figure out what's inside a box by shaking it. We propose that effective learning and exploration depend on being able to interpret evidence through the lens of intuitive theories—theories of both the physical world and one's own perceptual apparatus—to imagine how one's actions might change the state of the world and what kinds of changes would be most perceptually discriminable. In a paper to be presented at Cog Sci 2014 (Siegel et al.), we describe three studies exploring these capacities in young children, and suggest how they could support powerful and sophisticated inferences about hidden causes. Extensive development of this work is currently in progress.

- We have hypothesized that children's ability to generate good hypotheses for causal learning is driven in part by their ability to imagine possible causes whose abstract type is consistent with the abstract type of some observed effects, and abstract types of functions relating causes to effects (Magid et al., in press, Cognitive Development). For instance, a binary outcome is likely to be the effect of a discretely varying binary cause, and a continuous outcome is likely to be the effect of a continuously variable cause, under the assumption that the cause-effect relation is one-to-one and invertible. We have shown experimentally that young children can use this mode of reasoning to generate effective interventions for causal learning.

We intend to continue developing these projects over the course of Year 2, in accord with the milestones outlined below. In addition, our Year 2 plans include several core efforts that were not originally part of our thrust's research plan, but which have emerged as critically important opportunities during planning discussions in Year 1.

- Extending our models of intuitive physics from solid objects to liquids and containers. Our models are based on approximate probabilistic simulations, and originally we expected this approach would work only for solid objects. However, during the course of Year 1 meetings, it became clear that infants' early intuitions about liquids, and the interactions of solid objects with liquids, are important landmarks in the development of physical knowledge—and that our simulation approach could perhaps be extendable to these cases. We have thus hired a part-time research assistant to construct models and stimulus materials for testing these models in Year 2 and going forward.

- Enhancing our models of physical object perception and reasoning beyond the visual sense modality to cross-modal perception—especially to integrating vision and optics, and vision and audition. From the beginning of our work, we have been keenly aware that infants' early physical object concepts are not confined to the visual sense modality, but rather are amodal—they are abstract*, physical* representations, in terms of masses of stuff and the shape that takes, its properties, necessary to predict how it will move, how it interacts physically with other objects and stuff, and how it can be acted upon. However, until recently we did not have the ability to simulate with reasonable richness the haptic and auditory dimensions of physical object experience. Very recently we have figured out how to do this, building on new software tools developed by others in the computer game and computer graphics fields. Going forward we will use these tools to build and test cross-modal models of object perception, in adults and especially children. This is an especially exciting development because it is also one of our most promising routes to collaborate with Thrusts 2 and 4.

- Integrating infants' intuitive physics and intuitive psychology, via models of counterfactual reasoning and causal attribution. During our Year 1 seminar meetings, it became apparent that reverse-engineering models of infant cognition should focus not only on core domains of physics and psychology, but also on how these domains integrate, and the cognitive capacities that allow infants to reason flexibly across and between them. This motivated a new line of studies on causal and counterfactual reasoning, including models of these abilities that can be tested quantitatively in adults and older children, and experiments testing their predictions qualitatively in infants. We have published initial modeling efforts in Year 1 (Gerstenberg et al.), and have piloted infant studies which we plan to develop more fully in Year 2.

- Developing the Lookit project: a web-based platform for infant looking experiments at home. This project was begun by Kim Scott and Laura Schulz, two members of Thrust 1, in the months before CBMM officially launched. Its goal is to develop a platform for running experiments that test infants' knowledge and utilities through showing them visual displays and recording various looking measures, such as looking time to a display or part of a display, looking-away time, and habituation and dishabituation measures, which can indicate surprise or familiarity, curiosity or boredom, and preference or subjective valuation. Infants are tested in their homes, or wherever and whenever is convenient for parents, and looking data are recorded over webcams and coded remotely in the lab. Thinking this system could be a valuable resource for other center participants, Scott was invited to present the concept and initial design at one of the first weekly CBMM meetings in October 2013. This presentation generated great interest, and after further discussions with multiple PIs, it was decided to make Lookit a core part of Thrust 1's efforts going forward. This will enable us to scale this project up in ways that would otherwise be impossible and will be of use to other thrusts as well. Specifically, we plan to invest resources in Year 2 and beyond in Center personnel who will work to develop additional experiments and methods for Lookit, manage its use, and in collaboration with Thrusts 3 and 5, develop automated (computer vision-based) systems for coding infants' looking behavior. We expect Lookit to be transformative for CBMM and more broadly for all of infant cognitive research: it will support the first methods with sufficient scale to quantitatively test computational models of infant behavior, and it will provide the first low-cost approach to running multiple studies across different domains on the same set of babies, and tracking their performance longitudinally across development, necessary to build a truly integrative picture of cognitive and brain development.

Milestones, near term, 1–3 years:

- Construct a "roadmap" of how intuitive physics and intuitive psychology develops, and which learning mechanisms become available in young children, with a focus on ages 0–5 years and on behaviors that can be tested in developmentally appropriate versions of the CBMM challenge. Perform experiments with children to test key claims about core knowledge and learning mechanisms.

- Create developmental computational models for age-appropriate visual Turing tests at 0–12 months, in roughly 3 month intervals.
- Develop stimulus materials and experimental methods to validate models qualitatively for children 0–12 months, and quantitatively in experiments with older children and adults.

Milestones, mid to long term, 4–10 years:

- Create developmental computational models for age-appropriate visual Turing tests at 24–36 months, in roughly 6 month intervals.
- Develop stimulus materials and experimental methods to validate models qualitatively for children 24–36 months, and quantitatively in experiments with adults.
- Develop robust experimental methods for testing computational models quantitatively with children 0–36 months.
- Towards the "what happens next" question: test intuitive physics and intuitive psychology models increasingly integrated with language, sensitive to increasingly complex and abstract physical properties and mental states.

Thrust 2: Circuits for Intelligence

Abstract thinking and complex problem solving constitute paradigmatic examples of computation emerging from interconnected neuronal circuits. Progress towards a quantitative understanding of emergent intelligent computations in cortical circuits faces several empirical challenges (e.g., simultaneous recording and analysis of large ensembles of neurons and their interactions) and theoretical challenges (e.g., mathematical synthesis and modeling of the neuronal ensemble activity). Understanding neuronal circuits that implement solutions to the central CBMM challenges including "What is there?", "Where is that object?", Where am I?", "What will happen next?", is an essential part of scientific reductionism, leading to insights useful for developing intelligent machines.

In Thrust 2, we combine neurophysiological recordings in multiple species (rodents, monkeys, humans) across different scales (single neurons, field potentials) with computational models. We are working on paradigmatic examples of complex problems representing different aspects of intelligence including:

1) Invariance in computations underlying recognition of objects and people;
2) Understanding interactions among people, objects and scenes and
3) Interactions among brain areas.

Following this conceptual framework we are working on:

Tools to probe neural circuits at unprecedented resolution

We are developing the neuro-technology to build functional microscopes and tools to interrogate dynamical interactions in the brain including optogenetics, electrophysiology, and electrical stimulation.

- High-density electrode arrays consisting of thousands of microwires to simultaneously monitor the activity of large neuronal ensembles in multiple brain areas and the algorithms to interrogate the ensuing big data sets [Boyden, Wilson, Desimone, Kreiman]
- Optogenetic tools to manipulate neural circuits [Boyden, Wilson, Desimone]

Neural circuits implementing invariant computations (en route towards examining circuits that can solve “What is there? / Who is that person?”)

- Initial steps in mice to examine invariance to affine transformations [Koch, Buice, Kreiman, Poggio]
- Invariance to face transformations in macaque monkeys [Freiwald] and humans [Kreiman]
- Invariance to object occlusion [Freiwald, Kreiman]

These efforts will be linked to the theory developed in Thrust 5.

Neural circuits involved in spatial navigation

- Ensemble recordings in the rodent hippocampus and neocortex to investigate how neurons encode the answer to questions including “Where am I? Where have I been? Where am I going?” [Wilson]

Several ongoing efforts are aimed at understanding how different brain areas work jointly to implement intelligent computations in the context of attention and target selection [monkeys, Desimone; humans, Kreiman], navigation [rodents, Wilson], face recognition [Freiwald, Kreiman]

The main results so far are described in publications of the Center, especially in the series of technical reports called *CBMM memos*, see http://cbmm.mit.edu/publications-code-data/. Here we list a few notable ones:

- Prof. Ed Boyden has developed an imaging system that is capable of recording neural activity at millisecond timescale and neuronal resolution for an entire organism (Nature Methods 2014). This transformative technique pushes the frontiers in our capability to interrogate the circuits responsible for intelligent computations.
- In collaboration with Thrust 5 (Enabling Theory) we have connected the invariance theory (Poggio) to specific cortical computation as measured electrophysiologically by Freiwald. (~ “What is there? Who is that person?”)
- How to decode animal position from ensemble recordings in the rodent hippocampus [Wilson]. (~ “Where am I? Where am I going?”)
- Recognition of objects from parts presents a significant challenge for theories of vision because it requires spatial integration and extrapolation from prior knowledge. By recording intracranial field potentials from the human brain we have shown that higher visual areas remain selective (invariant) to strong degrees of occlusion. However, these visually selective signals emerged ~100 ms later for partial versus whole objects. These

results provide spatiotemporal constraints on theories of object recognition that involve recurrent processing to solve the pattern completion problem. [Kreiman] (~"What is there? Who is that person?")

- When searching for an object in a scene, how does the brain decide where to look next? We describe a simple mechanistic model of visual search that is consistent with neurophysiological and neuroanatomical constraints, can localize target objects in complex scenes, and predicts single-trial human behavior in a search task among complex objects. This model posits that target-specific modulation is applied at every point of a retinotopic area selective for complex visual features and implements local normalization through divisive inhibition. The combination of multiplicative modulation and divisive normalization creates an attentional map in which aggregate activity at any location tracks the correlation between input and target features, with relative and controllable independence from bottom-up saliency. We first show that this model can localize objects in both composite images and natural scenes and demonstrate the importance of normalization for successful search. We next show that this model can predict human fixations on single trials, including error and target-absent trials. We argue that this simple model captures non-trivial properties of the attentional system that guides visual search in humans. [Kreiman] (~ "What will happen next?")
- Development of high-density 3D probes for neural recordings [Boyden] (described in Caroline Moore-Kochlacs, Jorg Scholvin, Justin P. Kinney, Jacob G.Bernstein, Young Gyu Yoon, Scott K. Arfin, Nancy Kopell, Edward S. Boyden (2014) Principles of high-fidelity, high-density 3d neural recording, CNS.)

Milestones, near term, 1–3 years:

- Develop neurotechnology required for the longer-term goals including: (i) novel high-density multi-electrode arrays to interrogate neural circuits in rodents, monkeys, and humans and (ii) optogenetic tools to activate/inactivate sub-circuits (e.g., cortico-cortical feedback) to evaluate and constrain computational models.
- Develop stimulus sets (still images and video sequences) and experimental designs that can be used across labs and thrusts (joint effort with Thrusts 3, 5). These stimulus sets will initially focus on recognition of actions, faces, objects and interactions among them. These datasets will include annotations to be used in the experiments and computational models across the center in multiple different efforts.
- Evaluate the hypothesis that rapid recognition (people, objects, actions) can be described, to a first approximation, by a bottom-up architecture (Thrust 5).

Milestones, mid to long term, 4–10 years:

- Compare neural circuit data (neurophysiological recordings in rodents, monkeys, humans) with behavioral data and computational models (Thrust 5) in invariant recognition of actions, objects, people, and interactions among them, in making intelligent predictions about future behavior including where monkeys/humans will saccade next in the context of cluttered scenes and/or natural videos; in rodent/human navigation; and in evaluation of social interactions (monkeys/humans) (Thrust 4).
- Compare neural circuit data (neurophysiological recordings in rodents, monkeys, humans) with psychophysical data and computational models to constrain and inspire computational models (Thrust 5) in tasks that involve answering the Central Challenge questions including:

- “What is there?” – Neural circuits for invariant representation of objects and people
- “What is the person doing?” – Neural circuits for invariant representation of actions
- “What will happen next?” – Neural circuits involved in predictive coding
- “What happened before?” – Investigate how neural circuits can support inference of causal relationships including elements of intuitive physics (Thrust 1) and social interactions (Thrust 4)
- “Who is doing what to whom and when and why?” – Combine elements of the above questions into a mechanistic understanding of how such intelligent inferences can be instantiated in neural hardware

Planned connections between thrusts

T2–T1. Tenenbaum and Freiwald are hiring a joint postdoc to investigate statistical aspects of the neural code for visual stimuli.

T2–T3. Kreiman and Katz are beginning to characterize brain areas and signals involved in language understanding in human cortex.

T2–T4. Kanwisher and Freiwald are beginning to examine the neural correlates of social signals in macaques and humans

Kanwisher and Kreiman have been interacting via postdoc (Fedorenko).

T2–T5. Desimone and Poggio are working together to decode neural signals from the macaque brain.

Freiwald and Poggio are working together to decode neural signals from the macaque brain.

Kreiman and Poggio are working together to study invariance in the representation of visual information in computational models and physiological recordings.

Thrust 3: Visual Understanding

The goal of Thrust 3 is to combine vision with aspects of language and social cognition to obtain and communicate complex knowledge about the surrounding environment. To obtain full understanding of visual scenes, our computational models should be able to extract from the scene any meaningful information that a human observer can extract, about actions, agents, goals, object configurations, social interactions, and more. We refer to this capability as the ‘Turing test for vision’—being able to use vision to answer a large and flexible set of queries about objects and agents in an image or a video in a human-like manner. Queries can be, for example, about objects, their parts, spatial relations between objects, actions, goals, and interactions between agents. Understanding queries and formulating answers requires interactions between vision and natural language. Interpreting goals and interactions requires connections between vision and social cognition. Answering queries requires task-dependent processing, i.e., different visual processes to achieve different goals, combining bottom-up with top-down processing.

We aim to go beyond answering questions about visual scenes in order to model other aspects of cognition. A key aspect of social cognition is engaging in a discourse to come to a mutual understanding of a complex scene. To this end, we will model exchanges between agents consisting of a number of queries or statements about one or more visual stimuli. Such exchanges can be in the form of discussing a longer sequence of activities with a coherent narrative or exchanging, in language, detailed information about a specific activity. This ability to engage in a longer dialog with multiple stimuli facilitates learning visual and linguistic representations in a cognitively plausible fashion: learning by observing visual and linguistic stimuli and asking questions about them. We will interact with Thrust 1 in order to incorporate structures and biases derived from human developmental cognition into this learning process. To achieve our goals, we are developing methods for extracting meaningful information from images and videos based on extended interpretation and goal-directed processing. The first stage in this process will construct in a bottom-up manner an initial interpretation of the scene, and the second will generate and apply an interpretation in a task-dependent manner. Even with perfect bottom-up processing, generating task-dependent representations is crucial for a number of reasons, including the fact that the number of relationships between all objects that can be named in a scene is often astronomical. Developing such approaches will require novel algorithms and representations to bridge the gap between low-level processing and high-level cognition.

The integrative nature of this endeavor relies on interactions with the other thrusts. To gain insight into how the brain deals with scene recognition problems, we will collaborate with Thrust 2 to investigate the neural representations of both vision and language. To understand visual representations, we will investigate modeling cortical mechanisms of hierarchical representations for object and activity recognition and connect the computational constraints with neuronal circuits. To understand linguistic representations, we will investigate the neural basis of language processing in order to find representations which are amenable to learning and to serving as top-down constraints on visual processing. We will also collaborate with Thrust 4 in understanding actions, goals, and agents' interactions, and with Thrust 1 in incorporating useful structures and biases derived from human developmental cognition.

Multi-sentence event recognition

We have started working on understanding video events in context. Existing approaches handle individual events without understanding the relationship between multiple events and the objects that participate in them. We have developed a prototype approach that can recognize and describe sequences of events. Given a video and one or more English sentences describing one or more events, this approach can determine if the video depicts those events. It addresses several problems simultaneously in a single globally-optimized cost function by jointly detecting the objects, tracking them, and recognizing the events. This allows high-level knowledge from event recognizers to affect the low-level object detection and tracking. The current approach was tested on a small dataset, and is currently limited to a small set of actions and participating objects, performed by upright agents with large objects in a controlled environment. [Barbu, Berzak, Harari, Katz, Ullman]

Grounded question answering

To address one of the central concerns of CBMM, answering questions about visual stimuli, we have been extending our video event recognition approach to produce such answers. Answering questions goes beyond just producing a sentence that is true of a video. It requires generating an answer that conveys specific information in a succinct form while taking into account the shared understanding of the questioner and responder. Given a video, our approach can generate sentences while evaluating how discriminative they are. In other words, it determines if an answer is generally true of a video or if it specifically answers the question. This approach works in two interleaved stages. In the first stage, a candidate answer is generated and a model is produced which checks that the answer is indeed true of the video. In the second, samples are drawn from that model to check the specificity of the answer. We developed an initial algorithm for drawing samples from our models capable of answering simple questions about videos with a limited amount of clutter. [Barbu, Berzak, Harari, Katz, Ullman]

Objects and hands in context

Answering questions requires a detailed analysis of visual stimuli beyond just object recognition. To this end, we have integrated human pose recognition with our joint tracking and sentence recognition approach. Pose recognition is particularly challenging because limbs tend to be very small in the field of view, and they are deformable and largely untextured. Even state-of-the-art pose reconstruction methods are largely unable to produce acceptable results, even in controlled environments. By integrating pose recognition into our algorithm, it is able to automatically choose among many candidate poses. This allows the algorithm to use knowledge about events, such as the fact that picking up an object means that you are likely grasping it with your hand, to constrain the location of the limbs and produce better pose tracking. The current approach uses existing human pose detectors which are ill-suited for producing multiple candidate detections, a challenge that will be addressed by developing new algorithms to post-process their output. [Barbu, Harari, Katz, Ullman]

Vision and language in the brain

Another aim of our thrust is to understand how humans perform combined vision-language tasks. To address this problem we are employing in collaboration with Thrust 2 neuroimaging and electrophysiology techniques to explore the interaction between language and vision in the brain and to understand the structure of representations used for processing language. In previous work, we have developed an approach to recover, from fMRI scans alone, events being viewed by subjects. Recently, we have extended this work across subjects to recover events from fMRI scans of new subjects. This work has shown that events seem to be represented independently of the objects participating in these events. We are currently working with members of Thrust 2 to analyze human neural recordings acquired while subjects were watching Hollywood movies. As part of this effort we have developed a tool which allows us to annotate large numbers of spectrograms, comprising hours of audio, on Amazon Mechanical Turk. [Barbu, Berzak, Katz, Kreiman, Singer]

Intentions and goals in human interactions

The first phase of this project addresses the problem of detecting direction of gaze in 3D images. Our goal is to come close to human-level performance on this task. The required visual and 3D data is acquired using a Kinect sensor, while human performance is measured via psychophysics evaluation. In the second phase of the study we plan to model human–object interactions, as well as interactions between two or more humans, by means of intentions and goals, utilizing cues such as direction of gaze, head pose, body pose, and spatial relations.

We have completed the acquisition of the experimental video (Kinect) corpus. Furthermore, we extracted appearance and depth information from Kinect's video corpus, which is required for the gaze direction detector. In particular, we extracted the head location and pose, as well as the appearance of the face and eyes. We trained and tested both 2D and 3D gaze direction models (NN approach) in a leave-one-actor-out setting. We are currently also formulating a probabilistic generative model for gaze direction and goals. [Gao, Harari, Kanwisher, Tenenbaum, Ullman]

Visual processing with MiRC images

We studied human visual processing using image patches that are minimal recognizable configurations (MiRC) of larger images. Although they represent only a small portion of the image, the majority of people (on Amazon Mechanical Turk) can recognize MiRC images, while any cropping or down-sampling severely hurts recognition performance. We found that with a brief presentation time (100 ms) without a following masking image, psychophysical performance matched that of mTurk with MiRC images recognized by the majority of subjects, while their sub-MiRC children were very poorly recognized. When we followed the 100ms image presentation with a masking image (which is matched for low-level image statistics and believed to disrupt top-down processing), the recognition performance of MiRC images significantly decreased, suggesting that late stage, top-down processing is important for their recognition.

We also compared MiRC images to other image patches that a computer vision algorithm thought belonged to the same image category, but were drawn from a different image category (false positive image patches). Without any mask, humans are very good at distinguishing the MiRCs belonging to the class versus the non-class patches. When a mask is presented, human performance drops close to chance, suggesting that without top-down processing, human feedforward vision is very similar to common computer vision algorithms. [Isik, Poggio, Ullman]

Milestones, near term, 1–3 years:

- Develop (with Thrust 5) algorithms to learn basic invariances and pose variations and show that they can answer the “what,” “which,” and “who” questions on image and video databases of objects such as cars, people, and common animals.

Milestones, mid to long term, 4–10 years:

- Demonstrate models that can recognize objects and their parts, integrating vision and language in the domain of objects, and their properties and spatial relations, and develop algorithms for action recognition involving one agent and one object, answering the question “What is happening?” Interact with Thrust 1’s work on modeling developmental trajectories leading to these capabilities.
- Demonstrate bidirectional cooperation between vision and language in answering questions about, for example, social interactions involving multiple agents via questions such as “What is the person doing?” and “Who is doing what to whom?” and “What will happen next?” and “What do the people think of each other?” Interact with Thrust 4’s work on modeling of brain mechanisms involved in inferring information about social interactions from visual perception. Interact with Thrust 2’s work on understanding the representations and algorithms used by the brain in these tasks.

Thrust 4: Social Intelligence

The goal of Thrust 4 is to understand a fundamental component (and source) of human intelligence: high-level social perception. We aim to discover what social information we can extract from short silent video clips of individuals or social groups, the perceptual cues that underlie this ability, the functionally distinct components of this ability, the computations entailed in each component, and the brain basis of each. In particular, we aim to understand how the human mind and brain extracts the following kinds of information from a dynamic scene identified in the CBMM Challenge:

- What are these people doing?
- Are they friends or enemies?
- Do they trust each other?
- How, in detail, are they performing actions?
- Have you seen anything like this before?
- Why are they there? What will they do next?

The import of this thrust to the overall goals of the project is to establish the specific components of the CBMM Challenge by precisely characterizing the abilities of the best current computational system for human social perception: the human brain.

The members of Thrust 4 are making rapid progress on both ongoing projects, and the initiation of new projects, which are now taking shape with the help of our two-hour meetings of the entire group once every two weeks. Regular attendees of these meetings are PIs Kanwisher and Nakayama, CBMM postdocs Peterson (Kanwisher lab), Gao (Kanwisher & Tenenbaim labs), and Vaziri-Pashkam (Nakayama lab), as well as other members of the Thrust not currently funded by CBMM, including postdoc Powell (Saxe Lab), grad student Deen (Kanwisher and Saxe labs), grad student Samuel Anthony (Nakayama lab), and research assistants Cain and Cormeia (Nakayama lab). These meetings have been lively and energizing and have served to maintain a coherence to the different projects being pursued within this thrust.

For two ongoing fMRI studies of social perception in the STS, data collection is nearly completed and we are now writing up the work; we hope to have two papers on this work submitted within 3–6 months. One paper (Deen, Saxe, & Kanwisher) explores the overall functional organization of the STS, and the other (Koldewyn and Kanwisher) examines in detail one region within this organization: a region apparently specialized for perceiving social interactions. Other ongoing psychophysical studies are measuring the accuracy of perception of gaze direction (Gao, Tenenbaum, Hariri, and Kanwisher), and of online perceptual predictions of action such as reaching (Vaziri-Pashkam, Nakayama, and others). We expect our group to have two papers submitted on this work within six months.

New projects being planned include the collection of a large set of short real-world social videos filmed on a city sidewalk with a Kinect sensor (Peterson, Nakayama, Gao, Kanwisher, and others). We will film individuals and pairs of people, then obtain their permission on the spot to use their clips, plus also direct them to a website where they can provide information such as the traits of individuals (e.g., salary, as a proxy for dominance), states (were they happy? preoccupied? etc.), their relationship to the person they were with (co-workers? spouses? siblings?), and the nature of their interactions that moment with that accompanying person (e.g., cooperative versus antagonistic). We will then use these short silent video clips, and reduced versions thereof (e.g., snapshots revealing form but not motion information, and dynamic stick figure versions generated from the Kinect sensor, revealing biological motion information but not detailed form information) in psychophysical experiments to determine what social information we can extract from brief dynamic visual displays and which visual cues drive these abilities. This work will serve to provide a basic characterization of nonverbal visual social perceptual abilities in humans, which in turn will serve as the foundation for studies that model these perceptual abilities and that use these tasks in fMRI and TMS studies. The stimuli generated from this work will also be made available to the other thrusts for modeling work and for neurophysiological investigations.

Milestones, near term, 1–3 years:

- Literature Review: develop a taxonomy of social perception. All members of the thrust are working on this together.
- Discover with fMRI the functional architecture of social perception in the STS, a key region for perceiving dynamic social information (relevant to outcomes 1+2: "the way humans do it" and to the CBMM challenge). This work is being conducted by Deen, Saxe, and Kanwisher.
- Quantify human ability to predict another person's behavior in real time via read-out of motor behavior of the perceiver (Nakayama's "goalie" game to be added to the CBMM challenge set of questions). Vaziri-Pashkam and Nakayama.
- Generation of Kinect data on interactions to be modeled via goal-directed action, also to be incorporated in the CBMM challenge questions. In particular, create the following stimulus sets: 100s of movie clips, rated on many social dimensions (nature of the relationship, nature of interaction); Kinect videos of actors performing various actions, and pairs of individuals in various relationships to each other and interacting socially in various ways, "deconstructed" in various ways (stills, dynamic stick figures, etc.). Design and pilot behavioral tasks tapping a wide range of nonverbal social perception (NVSP) tasks. Peterson, Nakayama, & Kanwisher.
- Psychophysics: rich characterization of two visual social judgment domains (e.g., lying discrimination) and cues therein. Peterson, Nakayama, & Kanwisher.
- Machine learning of one key high-level social perceptual discrimination. (Outcome 1 and CBMM challenge). Peterson, Hariri, Gao, Tenenbaum, and others.

Milestones, mid to long term, 4–10 years:

- Functional organization of NVSP: cognitive and neural (critical to measure consistency of models with psychophysics and physiology in the CBMM challenge)
- Discover the cues, algorithms, and representations that enable high-level social perception.
- Evaluate models of NVSP for questions such as "*What is the person doing?*" and *"Who is doing what to whom?"* against fMRI and behavioral data in humans and fMRI data in monkeys.
- Discover homologies between human and monkey brain areas engaged in social perception so that underlying neural circuits can be studied at a finer grain than fMRI.

Thrust 5: Theory for Intelligence

Understanding intelligence and the brain requires theories at different levels, ranging from the biophysics of single neurons, to algorithms and circuits, to overall computations and behavior, to a theory of learning. In the past few decades, advances have been made at all levels including: statistical learning theory, machine learning, probabilistic inference, and the biophysics of computation. These theoretical foundations provide a common framework for fields as diverse as computer science, cognitive science, and neuroscience. Recent successes in intelligent systems applications—from Google to Watson—would not have been possible without these developments. For the first time, we have the beginnings of a unifying and useful mathematics of brains, minds, and machines—one with rigorous foundations, demonstrated applicability in almost every area of cognitive and neural science, and real practical value for building intelligent systems.

Theories provide understanding, guide computer implementations, and inform and interpret experiments. For this reason, Thrust 5 is not so much an independent thrust but an enabling platform, as we had originally described, common to all thrusts (we renamed it *thrust* instead of *Enabling Platform* for simplicity).

We do not expect that there will be a single theory for intelligence—it is not even clear there will be one, though we hope that our Center will make a major case for it! Our conceptual framework for vision—and for the CBMM Challenge—is presently the following working hypothesis:

- The first 100ms of vision in the ventral stream are mostly feedforward. The main computation goal is to generate image representations that are invariant or quasi-invariant to transformations experienced during development and at maturity, such as scaling, translation, and pose changes. The representation is used to answer basic questions about what kind of image and which object or person may be there.
- The answers will often have low confidence requiring an additional "verification step" which may often involve shifts of fixation and attention and actually do much more than verification. This "verification" step may rely on generative models and probabilistic inference and/or on top-down visual routines. Routines that can be synthesized on demand as a function of the visual task (think about questions in the open set of the CBMM Challenge) are needed in any case to go beyond object classification which—we argue—is only a part of vision. Running a routine may actually correspond to the process usually known as attention.

Following this conceptual framework we are working on:

- a theory of invariance computation in the ventral stream for the feedforward step
- generative models of images, probabilistic in nature and similar to computer graphics, that can be used to solve the inverse problem of vision by synthesis for the verification stages
- implementations of specific visual routines and of how they may be learned also for the verification stages

In each of these three approaches, our research is at different levels:

- at the computational and algorithmic level
- at the level of neural circuits

The main results so far are described in publications of the Center, especially in the series of technical reports called *CBMM memos,* see http://cbmm.mit.edu/publications-code-data/. Here we list a few notable ones:

- Theory of invariance
    - In collaboration with Thrust 2 (Neural Circuits) we have connected the invariance theory to specific cortical computation as measured electrophysiologically by W. Freiwald (Leibo, Freiwald, Poggio).
    - We have been able to provide a new computational model of eccentricity-dependence of receptive field sizes which makes predictions about anatomy and psychophysics (Poggio, Mutch, Isik).

  - We have shown that the invariance theory leads to computer systems with state-of-the-art recognition performance in face identification (Qianli, Leibo, Poggio).
  - We have obtained formal proofs that invariant representations can decrease the sample complexity of a classification task in vision and speech (Poggio, Rosasco, Anselmi).

- Generative models
  - Probabilistic programs provide a powerful framework for building expressive generative models; in particular, they provide the computational substrate for our generative models for physically and psychologically based scene understanding in Thrust 1, and related models in Thrusts 2, 3, and 4. Unfortunately, their great representational expressiveness comes at the cost of very challenging probabilistic inference. This year we have focused on developing new classes of inference algorithms, including a method of learning stochastic inverses from large datasets sampled from the model (Stuhlmuller et al., NIPS 2013), methods for exploiting the structure of probabilistic programs to generate more efficient MCMC kernels (Lang, Hanrahan and Goodman, 2014; Mansighka, Selsam and Perov, 2014) and more advanced methods for approximate inference based on particle MCMC and Hamiltonian Monte Carlo (papers in preparation).
  - Inspired by both the phenomena of infants' physical object perception and the needs of the CBMM visual Turing test, we have begun work on a new approach to modeling object perception that performs Bayesian inference to invert a graphics renderer, looking for three-dimensional object parses that best explain an observed two-dimensional image. This approach can naturally incorporate physical constraints (e.g., objects cannot float unsuspended in mid-air) that even 8-month-old infants have been shown to be sensitive to. A preliminary version was presented at the NIPS 2013 conference as an oral presentation (Mansinghka, Kulkarni et al), and the next step in this project is in preparation for NIPS 2014.
  - As a means to implement fast robust approximate inference in the probabilistic programs for object perception described above, we have begun exploring schemes that integrate bottom-up discriminatively trained proposals with the top-down Metropolis-Hastings proposals typical of probabilistic programs, which resample proposals from a partial evaluation of the prior. These schemes also have a very brain-like character to them: they could be a way to make sense of the interactions between top-down and bottom-up connections in the visual system. We have just begun discussions with other members of Thrust 5 to explore these connections.

- Visual routines
  - We formulated two related areas in which we focus our studies of top-down visual routines: one is answering queries about images (the 'Turing test for vision' that is the CBMM Challenge), and the other is top-down verification in visual recognition.
  - We have started to develop a two-state approach, which will be applied to both domains. The first stage will construct, in a bottom-up manner, an initial recognition and interpretation of the scene, and the second will generate and apply verification and interpretation in a task-dependent manner.
  - We started to formulate methods for applying top-down analysis to images guided by queries [Barbu, Berzak, Harari, Katz, Ullman]. We conducted studies to identify the limits of the bottom-up stage in recognition, using a combination of computational and psychophysical studies [Ullman, Harari].

- We performed an initial study that combined psychophysical and MEG experiments, using the same stimuli used in our computational studies, to identify components in the MEG signals that correspond to the bottom-up recognition and the top-down verification stages [Isik, Ullman, Harari, Poggio].

Milestones, near term, 1–3 years:

- Develop a theory of invariant recognition in hierarchical architectures, and develop associated neural model of the ventral stream. Test theory (with Thrust 3) with respect to the “what” and “who” CBMM challenge questions on various image databases. Test theory (with Thrusts 2, 4 and 1) with respect to physiology and psychophysical constraints.
- Develop a theoretical and computational framework for learning from very few labeled examples via unsupervised/weakly supervised learning of symmetries and other constraints from the environment. In particular, study online learning algorithms and connections with biophysical properties of neurons and their dynamics. Generate open-source code and different types of shared data, that can be used to integrate algorithms, methods, and theories across thrusts.
- Characterize the space of neural models for representation of concepts, their relationship / compositionalities and hierarchies.

Milestones, mid to long term, 4–10 years:

- Develop theory of visual understanding incorporating the physiology of attention and the anatomy of the back projections. Such theory will be based on formal frameworks such as Bayesian reasoning and visual routines. Demonstrate the feasibility of the theory to answer question such as “What happens next?”
- Characterize and test neural models of inference and reasoning and of models capable of representing intentional agents and their interactions.

2b. Research Progress

The Center understands the significant challenges that a large and diverse research team poses. From the beginning of the Center’s life, we have focused on developing communication among the researchers:

- We have emphasized the development of our Web site as a tool to communicate research progress, news and teaching both within the Center and to external audiences.
- We have started a series of internal technical reports (CBMM memos).
- We hold weekly research meetings every Friday, alternating between Harvard and MIT (where the large majority of PIs are). Each research meeting is focused on the discussion of progress on the CBMM challenge—typically discussing work on one of the questions of the challenge, which is presented as an informal talk by one or more of the team members. Titles, abstracts, and presenters can be found on our website at http://cbmm.mit.edu/events/. We recognize the importance of including all PIs and trainees in these sessions and will soon provide video or teleconference options for non-local PIs.
- The CBMM Postdoc group holds regular meetings, typically before the research meeting; these are separate and independently planned meetings in which postdocs

discuss ongoing research, potential collaborations between postdocs, and other postdoc-specific issues.

- We keep our CBMM Postdoc Group in contact with our industrial partners, encouraging a dialogue between the two groups regarding research opportunities and career advancement discussions. Recently, the group hosted a visit from the GE Research lab.
- We hold biweekly meetings of the Management Team which includes the thrust leaders, the research coordinator, the leaders of Education, Outreach and Knowledge Transfer activities, and the external evaluator. These meetings allow for coordination and monitoring of implementation. The research thrust teams are also meeting on a continuous basis, physically and electronically. Minutes of all management meetings are distributed to committee members and are archived in a central repository.
- We have hosted visits by external partners such as Alan Yuille (UCLA), Michael Buice (Allen Institute), Noah Goodman (Stanford), Winrich Freiwald (Rockefeller), Dr. Joel Z. Leibo (DeepMind), and Lorenzo Rosasco (U. Genoa, IIT).
- We have encouraged center staff to visit our partner institutions: Ethan Meyers visited the Allen Inst.; Joel Leibo (DeepMind) visited Rockefeller U.
- We have invited international partners to participate at the first CBMM Summer Course: Cheston Tan (A*Star) invited lecturer.
- We have organized a high quality technical workshop on Learning Data Representations.
- We have begun to systematically evaluate all education, outreach and diversity activities and to use formative evaluation to make improvements. Plans are in place to begin to assess the impact of CBMM on trainees, faculty, institutions, and the field.
- Faculty, trainees, and staff are required to submit their first Annual Report using the on-line system (ARS) in May 2014. ARS documents individuals' productivity, collaboration, engagement in CBMM activities, and serves as a basis for resource allocation and management decisions. Based on the response to the first report request, the ARS system will be refined.
- To promote communication and collaboration, the Center has also established more formal and high profile talks with the goal of informing the larger academic community about work relevant to the Center. These talks are an important part of our Knowledge Transfer mission. An example is the talk by Gary Marcus (see http://cbmm.mit.edu/event/special-seminar-computational-diversity-mesoscale-organization-neocortex/).

2c. Research Plans

Over the next 12 months, Center research will be conducted according to our specific aims and milestones.

III. Education

1a. Education Goals and Objectives

The overall education mission of the Center is to train a new generation of researchers and education leaders in the science and engineering of intelligence, with integrated knowledge and skills in computation, neuroscience, and cognitive science. Toward this mission, three primary goals of the education program are to

1. establish a model framework for education in the science and engineering of intelligence, including curriculum frameworks for interdisciplinary graduate and undergraduate training that are disseminated and adopted at a range of educational institutions
2. develop interdisciplinary graduate and undergraduate courses that integrate multiple approaches to the study of intelligence that are available to students at all CBMM partner institutions, and ultimately to the broader academic community
3. create educational opportunities that lead to the broader participation of women, underrepresented minorities, and other underserved groups in the field of intelligence science

To accomplish these goals, we need a dedicated community of educators across CBMM partner institutions engaged in the collaborative development of courses, educational experiences, learning materials, and curricula related to the study of intelligence. The overall mission and specific goals of the CBMM education program are embodied in five outcomes provided in the Strategic Plan for the Center.

1b. Education Performance and Management Indicators

Goal 1 will ultimately be measured in the number and types of institutions that adopt the curriculum frameworks that are created, the numbers and demographics of students who enroll, and the progression of these students into majors or degree programs. Adoption could take the form of a formal interdisciplinary major or degree program, or a less formal advising structure that provides guidance on academic training and research opportunities related to intelligence science. The number of faculty and trainees engaged in the process of developing curricular frameworks and communication of this work through the CBMM website, professional workshops or online curriculum resources at partner schools, are more proximal measures of performance. A first step in this process is to understand the existing graduate and undergraduate academic structures at CBMM partner institutions, including the relevant academic departments and programs, and current courses at each school that could be incorporated into these curriculum frameworks. Toward this end, a workshop was held in January, 2014, in which faculty from the Centers diversity partners shared information about the academic programs at their schools with CBMM faculty from MIT, Harvard, and Cornell. Plans were discussed for moving forward on the creation of curriculum frameworks, through collaboration between faculty across CBMM partner institutions. This workshop is described in Section III.2a.

Goal 2 can be measured by the number of interdisciplinary courses developed by CBMM faculty, the number of partner schools involved in development, the number and demographics of students taking these courses, and the extent to which curricular materials are disseminated to and adopted by other partner schools or the broader academic community. The original CBMM proposal described several courses that incorporate material that is central to the core research thrusts of the Center, expose students to important research methods and results, and integrate the perspectives of multiple disciplines. Section III.2a describes several courses taught by CBMM faculty this past academic year at MIT, Harvard, and Stanford. The summer course on Brains, Minds and Machines to be taught at the Marine Biological Laboratory at Woods Hole is described in Section III.2c. Finally, in Section III.2f, we discuss plans for the development of future graduate and undergraduate courses on the science of intelligence and a new undergraduate version of a graduate course that integrates vision and learning.

Goal 3 was directly addressed by the January workshop described above, for which progress will be measured by the number of faculty and students from schools that serve minorities, women, and other underrepresented groups, who participate in the education and research programs of the Center. Several CBMM outreach activities aimed at broadening the participation of faculty and students are described in Section VI of this report. In the context of education, we describe plans for the development of an introductory undergraduate course on the science of intelligence through a collaboration of faculty across the CBMM diversity partners (Section III.2a), and plans to offer a training workshop on MATLAB programming next January (Section III.2f).

Finally, an essential part of training students and postdocs to be the next generation of leaders in research and education is participation in professional development activities that include training in the ethical conduct of research and written and oral communication of research, and opportunities for teaching and mentoring. The participation of CBMM graduate students and postdocs in these activities is described in Section III.2b.

To manage the CBMM education programs effectively, we established an Education Committee whose members represent a range of institutions and enable close integration with other CBMM programs: Ellen Hildreth (Co-Coordinator for Education, Wellesley College), Haym Hirsh (Co-Coordinator for Education, Cornell), Matt Wilson (Associate Director, MIT), L. Mahadevan (Associate Director, Coordinator for Knowledge Transfer, and Co-Director of the Summer Course, Harvard), Mandana Sassanfar (Diversity Coordinator, MIT), and Lizanne DeStefano (External Evaluator, Illinois). This committee will oversee all aspects of the CBMM education program, including course and curriculum development, graduate and undergraduate student training, and evaluation of the education program.

1c. Education Challenges

Given the proximity of the primary partners, MIT and Harvard, a vibrant community of CBMM faculty, students, and postdocs has developed in the Cambridge area. Research collaborations are fostered by weekly research meetings that alternate between MIT and Harvard, and courses taught by Center faculty are available to students at both schools, creating a shared educational experience for CBMM students at MIT and Harvard. One of the challenges for the CBMM education program is to provide the means to ensure that faculty and students from all partner institutions can participate in these research meetings and courses. Faculty and students from other partner schools have expressed interest, for example, in viewing a live stream of the weekly research meetings, participating in a Journal Club that connects faculty and students across institutions through a virtual online communication tool, and enabling students to take CBMM courses through an online course delivery platform such as edX. The implementation of mechanisms to enable easy and regular communication across all partner institutions is in the planning phase for implementation in Year 2.

2a. Internal Education Activities

(1) CBMM Courses

We first describe five CBMM courses that were offered at MIT, Harvard, and Stanford in the fall of 2013. These courses were described in the original proposal and existed in some form prior to the granting period. They establish a foundation for CBMM course development and will continue to evolve over the coming years, enhanced by evaluation feedback and findings of the CBMM research program in ways that contribute to the core training of CBMM graduate and undergraduate students for interdisciplinary research on intelligence. The evaluation of these five courses in this first year was based on institutional mechanisms for soliciting student feedback and interviews with instructors. We also describe a sixth course that was offered for the first time this spring, jointly between MIT and Harvard, as a direct outgrowth of the research collaboration between Josh Tenenbaum (MIT) and Elizabeth Spelke (Harvard) on computational and empirical work on cognitive development. For this course, we elicited support of our external evaluation team to independently survey students and interview instructors about the quality and impact of the course and suggestions for future development. Finally, we describe a mini-course on methods for analyzing neural data offered by one of the CBMM postdocs in January, 2014.

Course 1	*Computational Cognitive Science*
Dates	Fall 2013 semester
Taught by	Josh Tenenbaum (MIT)
Participants	Graduate and undergraduate students at MIT and Harvard (Brain & Cognitive Sciences, Electrical Engineering & Computer Science, Mathematics, Physics)
Number of Participants	41 registered (27 undergraduates, 14 graduate students) (37 MIT, 4 Harvard)

Course 2	*Computation and Cognition: The Probabilistic Approach*
Dates	Fall 2013 semester
Taught by	Noah Goodman (Stanford)
Participants	Graduate and undergraduate students at Stanford (Psychology, Philosophy, Symbolic Systems)
Number of Participants	23 registered (15 undergraduates, 8 graduate students), 5–10 auditors

Josh Tenenbaum (MIT) and Noah Goodman (Stanford) collaborated on the development of a course on computational theories of human cognition that emphasizes probabilistic approaches to cognitive modeling. Topics include the principles of inductive learning and inference, representation of knowledge, and probabilistic models of cognition, with examples from concept learning, causal reasoning, language understanding, and social inference. Two versions of this course were taught at MIT and Stanford. Both had a strong interdisciplinary focus, but the MIT version had a greater emphasis on modeling methods while the Stanford version had a greater emphasis on human cognition, in part reflecting the different backgrounds of the students enrolled in these courses.

Much of the material in these courses is presented in an eBook, *Probabilistic Models of Cognition* (https://probmods.org), co-authored by Noah Goodman and Josh Tenenbaum, who made substantial revisions in preparation for the courses last fall. This eBook introduces the Church programming language for probabilistic modeling and includes interactive programming exercises throughout the book that are directly linked to an interpreter that can run model simulations and execute code entered by the user. This embedded programming environment has made the eBook a valuable experimental tool for researchers to perform simulation experiments and explore new modeling ideas. Viewers are asked to create an account to access the eBook, and to date, there are 235 registered users of the book and there have been over 5,000 unique visitors.

Extensive materials that include lecture slides and notes, readings, assignments, and project ideas, are available through course management systems at MIT and Stanford, for which faculty from other CBMM partner institutions can be given access. Josh Tenenbaum received an undergraduate teaching award from the MIT Brain and Cognitive Sciences department for his Fall 2013 offering of this course, which students found to be informative, interesting and inspiring. With the eBook in place to introduce much of the lecture content of the course, Noah Goodman is planning to replace some of the lectures with lab-style classes in which students engage in hands-on mini-projects.

Course 3	*Vision and Learning: Computers and Brains*
Dates	Fall 2013 semester
Taught by	Tomaso Poggio (MIT), Shimon Ullman (MIT, Weizmann), Ethan Meyers (postdoc, MIT)
Participants	Graduate and undergraduate students at MIT and Harvard (Brain & Cognitive Sciences, Electrical Engineering & Computer Science)
Number of Participants	15 registered (2 undergraduates, 13 graduate students) (14 MIT, 1 Harvard)

This course presents current research on the problem of learning to understand the world and interact with it using sensory information, integrating computational models of learning in both neural systems and machine vision systems and recent advances on problems from low-level learning in synapses to high-level processes such as the analysis of faces. This curriculum includes material from Thrusts 2 (Circuits for Intelligence) and 3 (Visual Understanding) of the CBMM research program. A publically accessible webpage linked from the CBMM website includes lectures slides and notes, readings, and videos of all of the lectures, presented by several CBMM faculty and guest lecturers (http://web.mit.edu/course/other/i2course/www/vision_and_learning_fall_2013.html). One student commented, *"The lecturers that Poggio & Ullman brought in were spectacular and help me get a broader understanding of some computational approaches to different parts of neuroscience."* Next fall, Tomaso Poggio and Shimon Ullman will offer a new undergraduate version of this course, in which they will be the primary lecturers. In the context of the CBMM curriculum, this course will provide advanced training in computational neuroscience for students interested in pursuing graduate work in the field of intelligence science.

Course 4	*Visual Object Recognition: Computational and Biological Mechanisms*
Dates	Fall 2013 semester
Taught by	Gabriel Kreiman (Harvard)
Participants	Graduate and undergraduate students at Harvard (Neurobiology, Computer Science, Molecular & Cellular Biology, Economics)
Number of Participants	10 registered (8 undergraduates, 2 graduate students), 6 auditors

With a focus on recognition, this course examines how circuits of neurons in visual cortex represent and transform visual information. Topics include the functional architecture of visual cortex, lesion studies, physiological studies in humans and animals, visual consciousness, computational models of visual object recognition, and computer vision algorithms for recognition. Classes combine a one-hour lecture on background material with a one-hour discussion of assigned readings from the research literature. In their evaluations, students cited the guided discussions of literature as one of the most valuable aspects of the course. All of the course materials are available at a public website (http://klab.tch.harvard.edu/academia/classes/hms_neuro300_vision/hms_neuro300_vision.html), which includes lecture slides and notes from a draft textbook on visual object recognition being written by Gabriel Kreiman. One of the challenges for this course is the diversity of student backgrounds, which can make it difficult to find research articles that are didactic, interesting, related to the lecture material, and understood both by CS and Neuroscience students. This will be a challenge for other CBMM faculty as well, as we develop courses that are suitable for an undergraduate audience with less prerequisite knowledge that spans multiple disciplines. One of the students from the course last fall will join the Kreiman Lab this coming summer.

Course 5	*Statistical Learning Theory and Applications*
Dates	Fall 2013 semester
Taught by	Tomaso Poggio (MIT), Lorenzo Rosasco (visiting faculty (U. Genoa), Team Leader LCSL (MIT-IIT))
Participants	Graduate students at MIT and Harvard (Brain & Cognitive Sciences, Electrical Engineering & Computer Science, Operations Research, Civil Engineering, Aeronautics & Astronautics)
Number of Participants	23 registered (all graduate students) (19 MIT, 4 Harvard), 13 auditors

This advanced graduate course covers theoretical foundations and recent advances in machine learning, with an emphasis on statistical learning theory. Topics include Regularization Networks, Support Vector Machines, and other supervised learning methods, as well as unsupervised methods for learning data representations with a focus on hierarchical deep architectures. The later part of the course presents a new theory of hierarchical architectures known as M-Theory, which is being explored as a possible model of visual cortex in Thrust 5 (Theory of Intelligence) of the CBMM research program. This course develops the theoretical tools needed for students to pursue research in this thrust. A syllabus with readings is available at the course website (http://www.mit.edu/~9.520/).

Course 6	*Computational Models and Cognitive Development*
Dates	Spring 2014 semester
Taught by	Josh Tenenbaum (MIT), Elizabeth Spelke (Harvard), Suzanne Corkin (MIT)
Participants	Graduate students and postdocs at MIT and Harvard (Psychology, Brain & Cognitive Sciences)
Number of Participants	Approximately 25, evenly divided between MIT and Harvard

A direct outgrowth of collaborative research in Thrust 1 (Development of Intelligence), this advanced reading seminar explores the prospects for “reverse engineering” infant and early childhood cognition over the first three years of life. Faculty and students from different disciplinary backgrounds discuss current computational and empirical research on cognitive development, including computational accounts of early-emerging core knowledge systems for intuitive physics, psychology, sociology, space and number, and the learning mechanisms that extend, enrich, and transform these systems as children grow. Students play an active leadership role in this course, working closely with the instructors to plan readings and formulate questions for discussion, and collaborating on the presentation of background material and leading class discussions. The course emphasizes key research questions and critical thinking about the relationship between empirical studies and computational models. The instructors planned an innovative class project to develop a wiki that serves as a resource for the research literature on cognitive development, a searchable database of research results and readings on various topics related to the early stages of cognitive development. CBMMs external evaluation team led by Lizanne DeStefano, will conduct a post-survey of students taking the course this spring, and interviews of faculty and students. Anecdotal feedback from participating faculty suggests that this course may provide a valuable model for other research-based seminars to help graduate students develop interdisciplinary thinking skills.

Course 7	*Methods for Analyzing Neural Data*
Dates	January 21–28, 2014
Taught by	Ethan Meyers (postdoc, MIT)
Participants	Primarily graduate students and postdocs from MIT, Harvard, Brandeis
Number of Participants	40–50 (80% MIT, 10% Harvard, 5% Brandeis, 5% other)

This short course was developed and taught at MIT by CBMM postdoc, Ethan Meyers, during the MIT “Independent Activities Period” last January. Class sessions included three one-hour lectures and two one-hour labs, and covered several methods for analyzing neural data including conventional statistics, mutual information, principal components analysis and decoding analyses. The examples focused on the analysis of neural spike activity, MEG signals, and local field potentials. The course provided some useful technical training for a broad audience, as well as a valuable teaching opportunity for a CBMM postdoc.

(2) Workshop on Broadening Participation in the Science of Intelligence

The primary aims of this workshop, which was held at MIT, were to share information about relevant education and research programs at the diversity partner schools, discuss current plans and future ideas for broadening the participation of faculty and students from these schools in CBMM education and research efforts, and provide opportunities to explore future collaborations between the visiting faculty and those at the primary CBMM partner institutions. Groundwork for the workshop was laid during the fall of 2013, through visits by Mandana Sassanfar to UPR, UCC, Hunter, and Queens, and a visit by Ellen Hildreth to Howard University. We summarize here and in Section VI Diversity, the main outcomes of the meeting. A full report is included with this annual report, and is available at http://cbmm.mit.edu/wp-content/uploads/2014/01/CBMM_Report_Jan14_Workshop.pdf.

Activity Name	*Workshop on Broadening Participation in the Science of Intelligence*
Dates	January 8–11, 2014
Led by	Ellen Hildreth and Mandana Sassanfar
Participants	Faculty from MIT, Harvard, Cornell, Howard University, Hunter and Queens Colleges, University of Puerto Rico at Rio Piedras, Central University of the Caribbean, and Wellesley College
Number of Participants	26 total: MIT (10), Harvard (2), Cornell (1), Howard (2), Hunter (3), Queens (1), UPR (3), UCC (1), Wellesley College (3)

Given the important role that our diversity partners will play in the training of students for graduate and postdoctoral work in this field, one focus of the discussions was the creation of curriculum frameworks. Visiting faculty stressed the need to identify the core concepts, knowledge base, and skill sets needed to work successfully in this field. From this core, it will be easier to identify existing programs and courses within their institutions that contribute to this core training, and to see where gaps exist that could be addressed through CBMM programs and resources, including participation in the summer research program, summer school, and training workshops. Some partner institutions, such as Hunter, Queens, and Wellesley, already have strong interdisciplinary programs that could provide a model for creating an undergraduate or graduate program based on intelligence science. Others, such as UPR and Howard, are still in a formative stage of building bridges between relevant disciplines, but have begun laying the foundations for interdepartmental collaborations.

A second focus of the education discussions at the workshop was the creation of an introductory undergraduate course on the science of intelligence, aimed at drawing young students into the field, exposing them to the big ideas and fundamental questions addressed and some of the compelling successes in intelligence research, and introducing some of the computational and empirical methods used in this research. There was a strong commitment among faculty across partner institutions, to collaborate on the development of materials for such a course, which could be shaped to fit the interests and needs of faculty and students in different programs and academic contexts. Patricia Ordóñez created a Moodle site at UPR to provide an online resource to facilitate file sharing, joint editing of shared documents, and a forum for ongoing online discussions of best practices, experiences, and ideas. The kinds of materials to be shared through this resource include the core concepts, knowledge and skills to be learned, reading materials, lecture slides, notes and videos, lab modules, software tools, project ideas, and data sets to support learning activities and research projects. Course development will proceed in Year 2, with an initial offering planned for Year 3.

2b. Trainee Professional Development Activities

CBMM graduate students and postdocs participate in professional development activities that include training in the ethical conduct of research, written and oral presentation of research to the CBMM community and at workshops and conferences, and opportunities for teaching and mentoring of students. With regard to ethics training, all graduate students and postdocs at each CBMM partner institution are required to participate in a program of instruction on the Responsible Conduct of Research (RCR) and Intellectual Property Rights. Matt Wilson organized and taught an RCR course for graduate students and postdocs in January, 2014. A table summarizing the participation of CBMM funded students in RCR training is provided below. The weekly research meetings held at MIT and Harvard frequently feature postdocs and graduate students presenting work in progress for discussion with the broader Center community. Workshops hosted by the Center, such as the *Workshop on Learning Data Representation: Hierarchies and Invariance* last November that was organized by Tomaso Poggio and Lorenzo Rosasco (visiting faculty U. Genoa), provide additional opportunities for trainees to present their work that complement ongoing participation in scientific conferences.

Activity Name	Departmental RCR course for graduate students and postdocs MIT IAP Course: 9.S911 Special Subject in Brain and Cognitive Sciences Responsible Conduct in Science
Dates	January 27-31, 2014
Led by	Matt Wilson (MIT)
Intended audience	MIT Brain and Cognitive Sciences researchers and trainees; also open to all CBMM grad students and postdocs
Approximate Number of Attendees	28, of this approximately 4 CBMM trainees

A CBMM Postdocs Group that currently has 16 members from MIT and Harvard, has met twice every month since its formation last December, to engage in networking, discussions of ongoing research, planning of activities such as workshops and meetings with industry partners of the Center. The postdocs organized a workshop with 8 visitors from GE Global Research that was held last April. At this workshop, postdocs and GE visitors presented and discussed work on many research problems of common interest in areas such as visual intelligence, human-like image analysis, anomaly detection on big grid data, the design of models and devices based on work in computational neuroscience, knowledge modeling and inference, and robotics and collaborative agents. The mentoring event with GE is summarized below.

Activity Name	CBMM Postdoctoral Mentoring Events: GE visit
Dates	April 2014
Led by	Matt Wilson (MIT)
Intended audience	CBMM grad students and postdocs, GE Global Research
Approximate Number of Attendees	37 (8 from GE, 29 CBMM postdocs, graduate students, and faculty/investigators)

CBMM postdocs and graduate students are also encouraged to pursue teaching and mentoring opportunities. This past year, postdocs helped to teach two of the academic year courses described in Section III.2a (*Vision and Learning: Computers and Brains* and *Statistical Learning Theory and Applications*) and the short course on *Methods for Analyzing Neural Data*. A CBMM graduate student taught a session on applications of MATLAB programming in neuroscience in an MIT workshop on Quantitative Biology. This session was integrated into the CBMM *Workshop on Broadening Participation in the Science of Intelligence* held last January, serving as a model for the MATLAB workshop planned for January, 2015. Several CBMM graduate students and postdocs will be teaching assistants for the CBMM summer school and help to supervise research projects conducted by small groups of participants over the two weeks of the course. CBMM graduate students also served as teaching assistants for the academic year courses. In the *Computational Cognitive Science* course, for example, the TAs helped to run tutorials on the programming tools used for modeling work in assignments and projects. Postdocs will also assist faculty with mentoring the SRP students conducting research in CBMM labs this summer.

2c. External Education Activities

Beginning on May 28, 2014, we will launch our first annual summer course on Brains, Minds and Machines, held on the campus of the Marine Biological Laboratory in Woods Hole, MA. This intensive two-week course will give advanced students a “deep end” introduction to the problem of intelligence—how the brain produces intelligent behavior and how we may be able to replicate intelligence in machines. Like CBMM itself, the philosophy of the course is that the synergistic combination of cognitive science, neurobiology, engineering, mathematics, and computer science holds the promise to build much more robust and sophisticated algorithms implemented in intelligent machines. The first of the two weeks will focus on general theoretical foundations and methods; the second week will examine four key areas of research for understanding intelligence, associated with the different research thrusts in the Center. The archived content of the course will serve as the foundation for new graduate and undergraduate courses on intelligence science. Further details of the course, including a schedule and topics, can be found in the Knowledge Transfer Section of this report.

2d. Integrated Research and Education

All of the CBMM courses described in Section II.2a have a strong research component. Through lecture content and assigned readings, students are exposed directly to the research of CBMM faculty, including past research that led to the formulation of the current research program and recently published results. These courses also introduce theoretical and empirical methods used in this research and in some cases, provide practice with the application of these methods through assignment work involving problem solving and critical analysis of research articles. Finally, these courses incorporate research-like final projects that are completed individually or by pairs of students.

The research of the Center has also been infused into other courses taught by CBMM faculty, creating an opportunity to draw new students into intelligence research. Patrick Winston’s MIT course, *Introduction to Artificial Intelligence*, which enrolls about 300 students every fall (https://ai6034.mit.edu/wiki/index.php?title=Main_Page), incorporates weekly “Right-Now” talks by guest faculty about their current research. The program has included talks by CBMM faculty Matt Wilson, Ed Boyden, Nancy Kanwisher, Shimon Ullman, and Boris Katz. These right-now talks were featured in an article in the April 2013 *MIT Newsletter* (http://web.mit.edu/fnl/volume/254/winston.html), which noted that 90% of students surveyed found these talks to be a good idea. The research of CBMM faculty Tomaso Poggio, Shimon Ullman, Joshua Tenenbaum, and Elizabeth Spelke was examined this spring in Patrick Winston’s MIT course, *The Human Intelligence Enterprise* (http://courses.csail.mit.edu/6.803/index.html), a research seminar for graduate students and advanced undergraduates focused on understanding human intelligence from a computational perspective.

As mentioned earlier, the eBook, *Probabilistic Models of Cognition*, by Noah Goodman and Josh Tenenbaum, which is broadly available to the academic community, was developed in an educational context, but has become a valuable research tool, enabling researchers to use the interactive programming examples to explore new modeling ideas and run simulation experiments. Other Center activities that integrate education and research include the summer research program described in Section VI and trainee professional development activities described in Section III.2b.

2e. Education Progress

The CBMM education program is aimed at training graduate and undergraduate students broadly in computational and empirical approaches to brain science that are essential to understanding intelligence. In the Strategic Plan, six near-term education milestones were enumerated to measure progress toward this ultimate goal. In this section, we briefly summarize our progress relative to these milestones.

Milestone #1: Develop graduate and undergraduate versions of an introductory course on the interdisciplinary science of intelligence, using material drawn from the annual summer course at the Marine Biological Laboratory in Woods Hole.

Progress toward Milestone #1: Plans are in place to offer a graduate version of this course at MIT, entitled *Aspects of a Computational Theory of Intelligence*, to be taught by Tomaso Poggio, Patrick Winston, and Ellen Hildreth in the fall of 2014 (see Section III.2f). At the January workshop on *Broadening Participation in the Science of Intelligence*, faculty from the CBMM diversity partners agreed to collaborate on the development of materials for an undergraduate version of this course and discussed the kinds of materials that are needed. This effort will begin after the first CBMM summer school ends in June, 2014.

Milestone #2: Establish mechanisms to support close collaborations between faculty at CBMM partner institutions on the development of new interdisciplinary courses and learning materials to be integrated into existing courses.

Progress on Milestone #2: An online site has been created at UPR for gathering curricular materials for an introductory course on the science of intelligence, which includes a forum for online discussions between faculty at partner institutions and a facility for sharing documents and other materials. All of the courses described in Section III.2a have materials posted online that are accessible to faculty at CBMM partner institutions.

Milestone #3: Offer online versions of courses on computational cognition and the science of intelligence. Establish an online teaching consortium based on the edX platform, to offer interdisciplinary CBMM courses to students across partner institutions.

Progress on Milestone #3: At this point, extensive materials for the computational cognition courses at MIT and Stanford are available through online course management systems that are accessible to faculty at other CBMM partner schools. The eBook for these courses is broadly available. Graduate and undergraduate introductory courses on the science of intelligence are under development. A decision has been made to hire a web developer for CBMM who can also support the transition of course materials to edX.

Milestone #4: Offer short training workshops on an annual basis, for students and faculty from the minority serving partner school, on core skills needed to conduct integrated computational and empirical research on intelligence. Offer a workshop on MATLAB programming and its application to work in areas such as neural modeling, image analysis, and machine learning.

Progress on Milestone #4: The need for a workshop on MATLAB programming was identified at the January workshop, and we plan to offer this workshop in January 2015. Substantial hands-on computer work is being planned for the CBMM summer school, some of which will be adopted for the workshop.

Milestone #5: Develop curricular frameworks for interdisciplinary undergraduate and graduate education. Identify core concepts, knowledge, and skills needed for advanced work in the science of intelligence. At each partner institution, identify courses that contribute to this core knowledge, and opportunities to expand their curriculum and integrate intelligence science into existing disciplinary programs.

Progress on Milestone #5: We are still at a formative stage in this process, but have taken an important first step in developing a deeper understanding of the relevant academic programs that currently exist at CBMM partner schools that may adopt these curricula. We have begun to collect information about the content of existing courses that clearly provide relevant training for research in this field, and plan to integrate this effort with an examination of ongoing research within the five thrusts, to identify common background elements that should form the core of an undergraduate or graduate training program for interdisciplinary research on intelligence.

Milestone #6: Establish professional development activities for students and postdocs in the areas of written and oral communication, ethics, leadership, teaching, and mentoring skills.

Progress on Milestone #6: These activities were described in Section III.2b and include ethics training, communication of research through oral and written presentations in Center venues as well as scientific workshops and conferences, organization of professional development activities such as meetings with representatives from the CBMM industrial partners, and teaching and mentoring activities.

2f. Education Plans

In this section, we briefly summarize the continuing and new education activities planned for Year 2.

(1) CBMM courses

(a) The following CBMM courses will be offered again in Year 2:

- (MIT) *Computational Cognitive Science*, for graduate students and advanced undergraduates (Josh Tenenbaum) (Fall 2014)
- (Stanford) *Computation and Cognition: The Probabilistic Approach*, for graduate students and advanced undergraduates (Noah Goodman, Stanford) (Spring 2014)

- (Harvard) *Visual Object Recognition: Computational and Biological Mechanisms*, for graduate students and advanced undergraduates (Gabriel Kreiman) (Fall 2014)
- (MIT) *Statistical Learning Theory and Applications*, for graduate students (Tomaso Poggio, Lorenzo Rosasco) (Fall 2014)

(b) The following two new courses will be developed, taught and evaluated in Year 2:

- (MIT) *Aspects of a Computational Theory of Intelligence*, for graduate students (Tomaso Poggio, Patrick Winston, Ellen Hildreth) (Fall 2014)

  This course will introduce new graduate students to the research of the Center and will incorporate two kinds of class sessions: (1) team taught seminar classes led primarily by CBMM faculty that are intended to stimulate discussion of the integration of research methods and results across disciplines, and (2) recitation classes led by a graduate teaching assistant that introduce key background material and include student-led presentations and discussions of research questions. The course will draw upon material developed for the CBMM summer course on Brains, Minds, and Machines.

- (MIT) *Computational Aspects of Biological Learning*, for advanced undergraduates (Tomaso Poggio, Shimon Ullman) (Fall 2014)

  This is a new undergraduate version of the course entitled, *Vision and Learning: Computers and Brains*, that was described in Section III.2a. This course examines supervised and unsupervised learning methods with emphasis on biologically plausible mechanisms for learning in the brain by neurons and synapses.

(c) We will form a working group of faculty from the CBMM diversity partner institutions to collaborate on the development of materials for an introductory undergraduate course on intelligence.

(d) The CBMM Summer School on Brains, Minds and Machines will be offered again in the summer of 2015.

(e) A training workshop on MATLAB programming will be offered in January 2015, for students and faculty from institutions that serve women, underrepresented minorities, and other underserved groups.

Formative evaluation of all courses will continue in Year 2 with emphasis on assessing student learning across disciplines, contribution to building a new discipline of the science and engineering of intelligence, and transportability to partner institutions and beyond.

(2) Development of curriculum frameworks

To advance the development of curriculum frameworks for graduate and undergraduate training, we will do the following in Year 2:

(a) Collect information about existing graduate and undergraduate academic programs at partner institutions, including relevant majors for students interested in pursuing graduate or postdoctoral work in intelligence science, and existing courses that provide important training for research in this field.

(b) Work on identifying the core concepts, knowledge, and skill sets needed to conduct various types of research in this field, guided in part by the broad scope of the CBMM research program.

(c) For the CBMM website, create a web-brochure describing the core elements of academic training for research in intelligence science, and organize information about existing academic programs and courses in a way that is informative to students and faculty advisors.

(3) Professional development activities

A graduate student leadership council will be formed, to foster communication among CBMM graduate students across all CBMM partners, on research, education, and professional development. Building upon the experience of the CBMM Postdoc Group this past year, the graduate student group will help to design and organize future activities that target leadership skills beyond those provided by the ongoing training that they receive through their research labs and courses.

IV. KNOWLEDGE TRANSFER

1a. Knowledge Transfer Goals and Objectives

The overall set of goals and objectives for Knowledge Transfer was laid out in our Strategic Plan as three Outcomes.

Outcome 1:

A cohesive Center drawing together neuroscientists, cognitive scientists, and computer scientists from academia and industry to tackle the new field of Science and Engineering of Intelligence

Outcome 2:

A global community of scientists and engineers dedicated to this new field

Outcome 3:

An active program of activities aimed at increasing public understanding and awareness of our goals, our accomplishments, and potential benefits of our research for society

We believe we are well on our way towards the first goal, have taken first steps towards the second goal, and have concrete plans for the third goal.

1b. Knowledge Transfer Performance and Management Indicators

The Knowledge Transfer Performance and Management Indicators were laid out in the Milestones section of our Strategic Plan.

Milestones, near term, 1-3 years:

1) Introduce 25 young scientists each year to the science of intelligence
    a. Host a summer course at MBL every year
    b. Host 1–2 scientific workshops every year
2) Establish relationships with 2–3 industry partners with an AI focus
    a. Host one workshop with a set of companies/groups with an AI focus to explore significance and direction of AI in their industry
3) Deepen relationships with 2 significant AI industrial partners
    a. Host two workshops (one with a big company, one with a small company), with partners who have significant AI focus, to explore deeper relationships
4) Strengthen centereness of CBMM

    a. Host one retreat per year
        i. Year 1: up to 60 participants
        ii. Year 2–3: up to 80 participants
    b. CBMM Weekly Research Meetings alternating between MIT and Harvard
5) Strengthen academic exchange with CBMM partner institutions
    a. Year 3: Each CBMM faculty member will have contributed to the Outreach program.
6) Website fully functional by end of Year 1
7) 1–2 public talks per year, beginning Year 2

1c. Knowledge Transfer Challenges

The path towards the first two Outcomes outlined in our Strategic Plan is clear. We have made less progress towards our goal of increasing public awareness and understanding of our mission and achievements. We believe that we will need to hire at least a part-time staff member to enhance our web-based knowledge transfer to the interested public, with a heavy emphasis on video.

2a. Knowledge Transfer Activities

Knowledge Transfer Activity Name: Summer Course on Brains, Minds, and Machines		
Led by		L. Mahadevan and Tomaso Poggio
Organizations Involved:		
1	Marine Biological Laboratories	Woods Hole, MA
2	Harvard University	Cambridge, MA
3	Massachusetts Institute of Technology	Cambridge, MA
4	Children's Hospital Boston	Boston, MA
5	Rockefeller University	New York, NY
6	UCLA	Los Angeles, CA
7	Istituto Italiano di Tecnologia	Genoa, Italy
8	Allen Insitute for Brain Science	Seattle, WA
9	A*Star	Singapore

This two-week course from May 29 through June 12, 2014 will give 25 advanced students an intensive introduction to how the brain produces intelligent behavior and how we may be able to replicate intelligence in machines.

Knowledge Transfer Activity Name: CBMM Workshop on Learning Data Representation: Hierarchies and Invariance		
Led by		Lorenzo Rosasco and Tomaso Poggio
Organizations Involved:		
1	Istituto Italiano di Tecnologia	Genoa, Italy
2	Massachusetts Institute of Technology	Cambridge, MA
3	California Institute of Technology	Pasadena, CA
4	Salk Institute	La Jolla, CA
5	New York University	New York, NY
6	UCLA	Los Angeles, CA
7	University of Toronto	Toronto, Canada
8	Hunter College	New York, NY
9	Siemens Corporate Technology	Princeton, NJ

The goal of the meeting, held 22–24 November 2013, was to advance one of the key topics of CBMM: learning invariant and hierarchical representations of information about the world (i.e., data).

Knowledge Transfer Activity Name: CBMM Workshop on Probabilistic Language of Thought		
Led by		Joshua Tenenbaum and Noah Goodman
Organizations Involved		
1	Massachusetts Institute of Technology	Cambridge, MA
2	Stanford University	Palo Alto, CA

This meeting is planned for August 2014.

Additional Outcomes

We have two new Industrial Partners since the Center got underway: Schlumberger Ltd, the world leader in oil-field services, joined as a partner in January 2014 and General Electric, one of the largest companies in the world also joined as a partner and was hosted by CBMM members at MIT to exchange ideas about problems of mutual interest via a workshop on April 25, 2014.

2c. Knowledge Transfer Progress

In the Strategic Plan, seven near-term (1–3 year) Knowledge Transfer Milestones were set. Here, we summarize our progress towards the following milestones (see also above).

Milestone #1: Introduce 20–25 young scientists each year to the science of intelligence

- Host 2-week summer course at MBL every year
- Host 1–2 scientific workshops every year

Progress toward Milestone #1:

- CBMM Summer Course

  Beginning on the evening of May 28 of this year, we will launch our first annual summer course. This intensive two-week course will be held on the campus of the Marine Biological Laboratory in Woods Hole, MA. It will give advanced students a “deep end” introduction to the problem of intelligence—how the brain produces intelligent behavior and how we may be able to replicate intelligence in machines. The premise is that today’s AI technologies, such as Watson and Siri, are impressive, but their domain specificity and reliance on vast numbers of labeled examples are obvious limitations; few view this as brain-like or human intelligence. Like CBMM itself, the philosophy of the course is that the synergistic combination of cognitive science, neurobiology, engineering, mathematics, and computer science holds the promise for building much more robust and sophisticated algorithms implemented in intelligent machines. The goal of this course is to help produce a community of leaders that is equally knowledgeable in neuroscience, cognitive science, and computer science.

  Even in this, the first year of the course, and despite the constraint of a late spring course time, interest in the course was noteworthy. We received and reviewed 140 applications from around the country and around the world. These were overwhelmingly serious, competitive applications. We limited the course to 25 students from outside of the core of CBMM researchers, which meant that the students appear to be a very strong group, indeed.

Course faculty will be CBMM faculty, for the most part, with a smaller number of guest lectures.

The first of the two weeks will focus on general theoretical foundations and methods; the second week will examine four key areas of research for understanding intelligence.

The theoretical foundations discussed in the first week will consist of:

- Inverse problems & well-posedness as a unifying theme;
- Signal processing;
- Machine Learning;
- Bayesian inference;
- Computational vision;
- Planning and motor control; and
- Neuroscience: neurons and models.

These topics will be complemented in the first week by MathCamps and NeuroCamps, to refresh the necessary background for some of the students.

The four areas of research examined in the second week will be:

1. *Development of Intelligence* – Understanding intelligence requires discovering how it develops from the interplay of learning and innate structure.
2. *Circuits for Intelligence* – Understanding the physical machinery of intelligence requires analyzing brains across multiple levels of analysis, from neural circuits to large-scale brain architecture.
3. *Visual Intelligence* – Visual intelligence goes beyond the narrow domains of face recognition or detecting pedestrians crossing the street to detailed scene understanding, including context, actions, inferences, predictions, linguistic associations, and narrative.
4. *Social Intelligence* – Intelligence emerges from the social interactions among individuals.

Core presentations will be given jointly by neuroscientists, cognitive scientists, and computer scientists who have worked together. In each of the two weeks, the first two days of intensive lectures will be followed by three days of morning lectures and afternoons of computational labs, with some additional evening research seminars. To reinforce the theme of collaboration between (computer science + math) and (neuroscience + cognitive science), exercises and projects often will be performed in teams that combine students with both backgrounds.

The last two days will be reserved for student presentations of their projects. These projects provide the opportunity for students to work closely with the resident faculty, to develop ideas that grew out of the lectures and seminars, and to connect these ideas with problems from the students' own research at their home institutions.

This summer course will not only be a way to transfer knowledge to the broader community of young and aspiring scientists, it will be a way to bond our Center community even more firmly together. The evaluation will track participants over time to assess the impact of the Summer School on scholarship, professional networking careers, and the development of the new discipline.

- CBMM Workshop on Learning Data Representation: Hierarchies and Invariance

The goal of the meeting, held 22–24 November 2013, was to investigate advances and challenges in learning "good representations" from data, in particular representations that can reduce the complexity of later supervised learning stages. The meeting gathered experts in the field to discuss current and future challenges for the theory and applications of learning representations. The key topics were

- Early Features in Vision
- Learning Features and Representations
- Learning Invariances and Hierarchies
- Beyond Feedforward Architectures

Milestone #2: Establish relationships with 2–3 industry partners with AI focus

- Host one workshop with a set of companies/groups with an AI focus to explore significance and direction of AI in their industry

Progress toward Milestone #2:

We have established relationships with Schlumberger Ltd and General Electric Ltd. These companies have committed to involvement with CBMM. We are still in discussions about the particular forms that the relationship will take, and a first meeting with GE managers occurred on April 25, 2014. A similar meeting with Schlumberger is planned for Fall 2014. We have a gift fund of $100K from Schlumberger to explore potential collaborative projects and partial funding for graduate student and postdoctoral fellowships.

We are planning a series of industrial workshops, when research leaders from these companies will describe their AI goals and CBMM faculty will describe our ongoing work.

Milestone #3: Deepen relationships with 2 significant AI industrial partners

- Host two workshops (one with a big company, one with a small company), with partners who have significant AI focus, to explore deeper relationships

Progress toward Milestone #3:

This will be evaluated in Years 2-3 after the planned industry workshops commence in Year 2.

Milestone #4: Strengthen centerness of CBMM

- CBMM Weekly Research Meetings alternating between MIT and Harvard
- Host one retreat per year
    - Year 1: up to 60 participants
    - Year 2–3: up to 80 participants

Progress toward Milestone #4:

Activity Name	CBMM Weekly Research Meeting
Dates:	Sept. 27, 2013; Oct. 4, 2013; Oct. 11, 201; Oct. 18, 2013; Nov. 1, 2013; Nov. 8, 2013; Nov. 15, 2013; Dec. 13, 2013; Feb. 7, 2014; Feb. 14, 2014; Feb. 28, 2014; March 14, 2014; March 21, 2104; April 4, 2014; April 11, 2014; April 18, 2014; April 25, 2014; May 2, 104; May 9, 2014; May 16, 2014
Led by	Tomaso Poggio (MIT) and Kenny Blum (Harvard)
Participants	All CBMM faculty, postdocs and graduate students
Approx. Number of Attendees	Attendance varies, approximately 40–45 attendees per week

First retreat is planned for January 2015. Annually thereafter.

Milestone #5: Strengthen academic exchange with BPWM partner institutions

- Year 3: Each CBMM faculty member will have contributed to the Outreach program.

Progress toward Milestone #5:

All of the MIT and Harvard CBMM faculty participated actively in the Workshop for Broadening Participation led by Sassanfar and Hildreth in January 2014. Further academic exchange will be evaluated in Years 2-3 after the CBMM faculty mentor summer students from our diversity partner institutions.

Milestone #6: Website fully functional by end of Year 1

Progress toward Milestone #6:

We have a live CBMM website (http://cbmm.mit.edu), which is fully functional from a practical standpoint, with both a public-facing, and a data-collecting side.

Milestone #7: 1–2 public talks per year, beginning Year 2

Progress toward Milestone #7:

Activity Name	Introduction to the new Center for Brains, Minds and Machines
Date:	Oct. 25, 2013
Hosted by:	McGovern Institute for Brain Research at MIT
Speaker(s):	Dean Marc Kastner, Bob Desimone, Tomaso Poggio, Patrick Winston

Activity Name	Special Seminar: Understanding the building blocks of neural computation: Insights from connectomics and theory
Date:	Oct. 20, 2013
Hosted by:	Tomaso Poggio
Speaker	Dmitri “Mitya” Chklovskii, Janelia Farm, HHMI

Activity Name	Special Seminar: What is the information content of an algorithm?
Date:	Nov. 7, 2013
Hosted by:	Tomaso Poggio, Lorenzo Rosasco
Speaker	Joachim M. Buhman, ETH

Activity Name	Special Seminar: Constructing space: how a naive agent can learn spatial relationships by observing sensorimotor contingencies
Date:	March 6, 2014
Hosted by:	Joshua Tenenbaum
Speaker	Alexander V. Terekhov Postdoc, Institute for Intelligent Systems and Robotics, Paris Descartes University

Activity Name	Special Seminar: Making Collective Intelligence Work: Learning, Liquidity, and Manipulation in Markets
Date:	April 17, 2014
Hosted by:	Tomaso Poggio
Speaker	Sanmay Das, Washington University in St. Louis

Activity Name	Special Seminar: Computational diversity and the mesoscale organization of the neocortex
Date:	April 22, 2014
Hosted by:	Joshua Tenenbaum
Speaker	Gary Marcus Professor, NYU

Activity Name	Special Seminar: Parsing Objects and Scenes in Two- and Three-Dimensions
Date:	May 16, 2014
Hosted by:	Tomaso Poggio
Speaker	Alan Yuille, UCLA, CBMM

2d. Knowledge Transfer Plans

More workshops: We are in the planning stages of our workshops for Year 2.

International workshop: Our first CBMM Workshop, held in November 2013 at MIT, was in close collaboration and co-sponsorship with IIT Genova. We are beginning to plan our second international workshop, at the location of one of our international collaborators.

More industrial partners: We are in active discussions with several major AI companies as well as a number of smaller AI startups. We note that two of our original startup partners were acquired by our major partners within just the last year.

Hold CBMM retreat in Year 2: We are planning our first annual retreat in January of 2015.

Increase CBMM visits and talks at partner institutions: In Years 2-3, after the CBMM faculty mentor summer students from our diversity partners, we will begin a regular series of visits and talks at their institutions.

Increase video content on website: We have hired a local videographer, Giro Studio, to record video content for the CBMM website.

Faculty public lectures: We have had seven lectures in the Brains, Minds, & Machines Seminar series. The latest was by Gary Marcus, Professor of Psychology at NYU and Visiting Cognitive Scientist at the Allen Institute for Brain Science. He is a New York Times bestselling author and blogs for the New Yorker. A screen shot of the video for his lecture can be seen below.

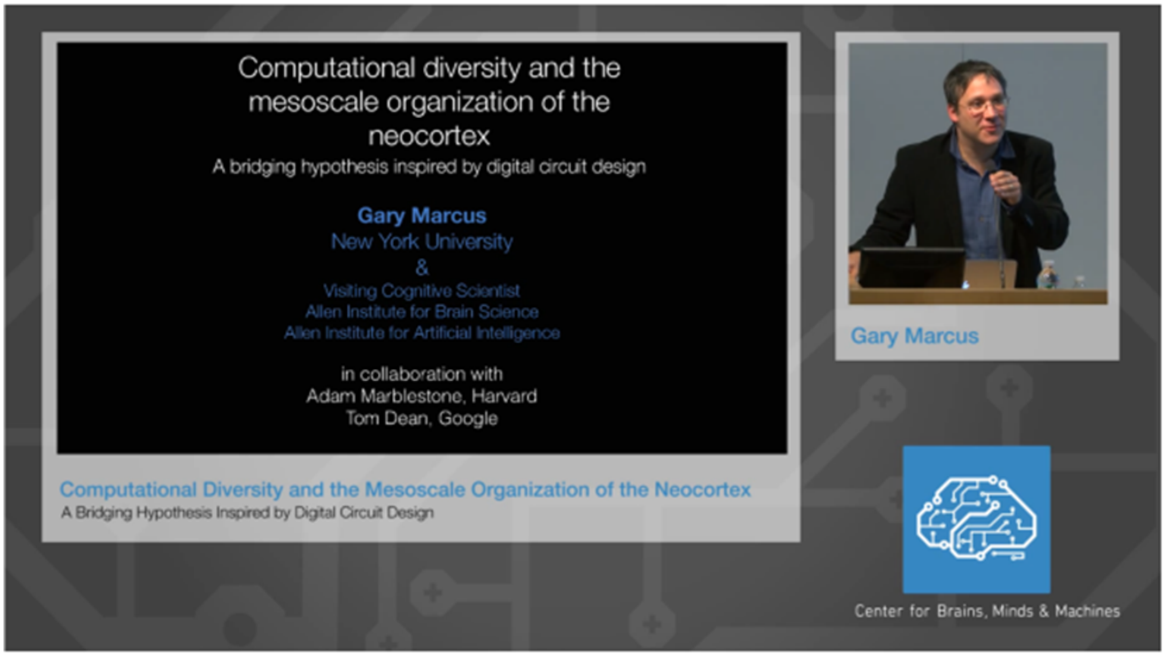
Computational diversity and the mesoscale organization of the neocortex
A bridging hypothesis inspired by digital circuit design
Gary Marcus
New York University
&
Visiting Cognitive Scientist
Allen Institute for Brain Science
Allen Institute for Artificial Intelligence
in collaboration with
Adam Marblestone, Harvard
Tom Dean, Google

Computational Diversity and the Mesoscale Organization of the Neocortex
A Bridging Hypothesis Inspired by Digital Circuit Design

Figure 1 Screenshot of video lecture, Gary Marcus – NYU, Allen Institute for Brain Science, pictured

In our efforts to fulfill our Education, Diversity, Knowledge Transfer, and Research goals to widely distribute the progress of the Center, we have hired a local videographer, Giro Studio, to record classes, lecture series, seminars, and Center-related events. Our videos will serve as a science education resource for our Partners and the public through our webpage (http://www.cbmm.mit.edu).

V. EXTERNAL PARTNERSHIPS

1a. External Partnership Goals and Objectives

CBMM aims to discover how intelligence is grounded in computation, how those computations develop in childhood, how those computations are implemented in neural systems, and how social interaction amplifies the power of those computations. Our own belief in what intelligence is and how it evolves means that we believe in communicating ideas and know-how since this is how human intelligence—genes and memes—increases in history. Thus we put a high priority in our ideas sharing activities: Education, Outreach, and Knowledge Transfer. These three activities are complemented by our Industrial Partnerships and our International Partnerships (which are in fact part of the Knowledge Transfer area). More specifically, the goals and objectives of our Partnerships are to extend our Educational efforts beyond our Center, in creating a generation of leaders that is equally knowledgeable in computer science and neurocognitive science and to widely disseminate research findings. The latter goal will be accomplished by developing collaborative relationships with commercial industries and international institutions devoted to the science and engineering of intelligent machines.

A main goal of the Center is the establishment of an emerging field, the Science and Engineering of Intelligence. This new field will leverage the progress in computer science, neuroscience, and cognitive science to generate a new discipline between science and engineering that addresses a growing interest among incoming graduate students. The ability to develop and build intelligent machines will influence a technology-based economy in the long term. The Center's outreach and knowledge transfer program will attract young students into this exciting new interdisciplinary field, improve the participation of underrepresented minorities, and strengthen US competitiveness in a global knowledge and intelligence economy. As a consequence, we want to package our new knowledge in accessible ways, including model courses at the graduate and undergraduate levels, that should be available beyond our Center member Institutions. Similarly, we want to ensure that our new knowledge is quickly and broadly disseminated and brought to bear on the great challenges of the 21st century, so as to serve the people of the nation and the world. Partnerships with non-Center organizations are needed to fully realize its mission. Therefore, we have established external partnerships with national and international organizations in industry and academia. This goal was articulated in our Strategic Plan and has not changed substantially since that time.

1b. External Partnership Performance and Management Indicators

The Center plans to foster collaborative research efforts, to organize sessions at scientific meetings, to communicate Center research to the larger scientific community, to identify and set policies to deal with ethical concerns related to the work of the Center, provide graduate educational and training opportunities, and to partner with industry organizations to develop applications based on Center research. All of the related indicators can be found in the sections dedicated to Education, Outreach and Knowledge Transfer.

1c. External Partnership Problems

The Center accomplished its early external partnership goals and did not encounter any major problems. A minor problem is that our smaller industrial partners have decreased in number—mainly because Google bought two of them (DeepMind and Boston Dynamics) since our Center started. This is however a good sign that our area of research and our choices of partners find an approval of sorts from what we think is the largest commercial player in the market of intelligence.

2a. External Partnership Activities

The Center has established a number of External Partnerships in the first reporting period. The industrial partnerships, the international partnerships and the broadening participation partnerships are listed at http://cbmm.mit.edu/about/partners/.

The industrial partners are General Electric, Siemens Corporation, Schlumberger, DeepMind-Google, Google, IBM, Microsoft Research, Rethink Robotics, Mobileye, Orcam

The international partners are A*STAR (Singapore), Weizmann Institute of Science (Israel), Universita' di Genova (Italy), National Center for Biological Sciences (India), Max Planck Institute for Biological Cybernetics (Germany), Italian Institute of Technology (Italy), Hebrew University (Israel), City University (Hong Kong)

The broadening participation partners are Wellesley College, University of Puerto Rico (UPRRP), Universidad Central del Caribe (UCC), The City University of New York: Hunter College and Queen's College, Howard University.

Partnership Activity	Collaboration on joint project: Reverse-Engineering Visual Intelligence for CognitiVe Enhancement (REVIVE)		
Led By	Tomaso Poggio		
Organizations Involved			
	Name of Organization	Shared Resources	Use of Resources
1	A*Star		

Partnership Activity	Program Committee, International Conference on Image Analysis and Processing (ICIAP). Naples Italy. September 9, 2013		
Led By	Shimon Ullman		
Organizations Involved			
	Name of Organization	Shared Resources	Use of Resources
1	Italian Group of Researchers in Pattern Recognition (GIRPR)		
2	Italian International Association for Pattern Recognition (IAPR) Member Society		

Partnership Activity	Organized workshop: Learning Data Representation: Hierarchies and Invariance, Massachusetts Institute of Technology. Cambridge, MA. November 22, 2013		
Led By	Lorenzo Rosasco		
Organizations Involved			
	Name of Organization	Shared Resources	Use of Resources
1	MIT		
2	Italian Institute of Technology (IIT)		

<table>
<tr><td colspan="2">Partnership Activity</td><td colspan="2">Coordinated Workshop on Broadening Participation in the Science of Intelligence-Quantitative Biology Workshop: An introduction to MATLAB: “Introduction to Visual Neuroscience - Using MATLAB to Analyze Imaging Data, Machine Learning Applied to Neuroscience of Vision, Data Simulation and Analysis using R and Python.” January 8–11, 2014</td></tr>
<tr><td colspan="2">Led By</td><td colspan="2">Mandana Sassanfar</td></tr>
<tr><td colspan="4">Organizations Involved</td></tr>
<tr><td></td><td>Name of Organization</td><td>Shared Resources</td><td>Use of Resources</td></tr>
<tr><td>1</td><td>MIT</td><td></td><td></td></tr>
</table>

<table>
<tr><td colspan="2">Partnership Activity</td><td colspan="2">Organizing Committee, Israel Machine Vision Conference (IMVC). Tel Aviv Israel. April 1, 2014</td></tr>
<tr><td colspan="2">Led By</td><td colspan="2">Shimon Ullman</td></tr>
<tr><td colspan="4">Organizations Involved</td></tr>
<tr><td></td><td>Name of Organization</td><td>Shared Resources</td><td>Use of Resources</td></tr>
<tr><td>1</td><td>Weizman Institute of Science</td><td></td><td></td></tr>
</table>

<table>
<tr><td colspan="2">Partnership Activity</td><td colspan="2">Co-Organized Conference: National Institute of Informatics (NII) Shonan Meeting: Deep Learning: Theory, Algorithms, and Applications
Shonan Village Center, Japan. May 19-22, 2014</td></tr>
<tr><td colspan="2">Led By</td><td colspan="2">Tomaso Poggio</td></tr>
<tr><td colspan="4">Organizations Involved</td></tr>
<tr><td></td><td>Name of Organization</td><td>Shared Resources</td><td>Use of Resources</td></tr>
<tr><td>1</td><td>National Institute of Informatics (NII)</td><td></td><td></td></tr>
</table>

2b. Describe any other outcomes or impacts of partnership activities not listed elsewhere.

None

2c. Describe how the Center is doing with respect to the indicators/metrics listed above. Include any data that have been collected on the indicators/metrics.

The Center accomplished its external partnership goals in the first reporting period. Center members are currently organizing a number of symposia/workshops at international scientific meetings, and will continue to do so throughout the next reporting period. The details of these events are provided in the Knowledge Transfer section of this report.

2d. Describe your plans for partnership activities for the next reporting period with attention to any major changes in direction or level of activity.

We do not anticipate any major changes in our partnership plans in the upcoming reporting period. In the next reporting period the Center has identified a number of potential outreach partners and events including the Cambridge Science Festival and the MIT Museum.
The nature of our partnership plans and our vision for the value the CBMM can bring to these organizations and events are detailed in the Diversity portion of this report.

VI. DIVERSITY AND BROADENING PARTICIPATION

1a. Diversity Objectives and Goals

One important goal of CBMM is to train the next generation of scientific leaders in the field of neuroscience, specifically in computational neuroscience and intelligence science. Therefore CBMM must recruit talented individuals from a broad and diverse background and ensure that women, under-represented minorities, disadvantaged students and students with disabilities have full access and opportunity to learn about this new emerging field. In order to increase the number of individuals from these groups who choose to pursue an advanced degree in cognitive or computational neuroscience, CBMM has implemented a number of outreach programs and initiatives aimed at increasing diversity.

These programs are listed below:

I. A 6-day intensive CBMM Workshop on Broadening Participation in the Science of Intelligence with lectures and hands-on activities to introduce students and faculty to important experimental and programming tools used to generate and analyze experimental data in neuroscience research, and explore models of intelligence processes. This workshop is offered during the Winter break to facilitate scheduling and to allow students and faculty to participate without missing classes. The class limit is set at 25 but may have to expand to accommodate more participants if there is a high demand.

II. A 10-week summer research internship program (CBMM-SRP) for undergraduates and post-baccalaureate students majoring in various STEM disciplines to expose them to CBMM research. This program is essential to build a pipeline of students who would be well prepared to apply to graduate programs in neuroscience.

III. A series of videotaped introductory lectures geared at an audience of talented undergraduates who have little or no knowledge of the field. These lectures are meant to expose students to the various research thrusts of CBMM, and arouse their interest in the subject matter. The lectures will be videotaped in High and posted on the CBMM website for wide accessibility.

IV. Summer sabbatical for faculty from minority serving institutions and their students. This is an important program to help faculty from our diversity partners to build strong research collaborations and design new curriculum for their institution.

V. Seminars by CBMM faculty at women and minority-serving institutions. This series will help reach out to undergraduate and graduate audiences and introduce them to the field of intelligence.

VI. Travel to minority-serving institutions and presence at minority conferences.

1b. Progress Assessment

Lizanne Destefano, our external reviewer, will be reviewing all aspects of the diversity program.

For the 10-week summer internship program participants (students and faculty) will complete an online evaluation form and attend a focus group. The results of the performance evaluation will provide metrics and an objective assessment of the immediate benefits of the various programs. We also expect that a number of the summer students will present their research at National meetings. We will report these numbers along with the outcomes (travel awards, prizes for presentations, active recruitment by graduate programs).

Most summer students are expected to enroll in PhD programs in neuroscience or computer science. Participants will be followed for at least 5 years to determine the long-term effectiveness of the program. A similar evaluation will be conducted for the 6-day workshop. The benefits will be assessed by the number of participants who are accepted into competitive summer program at MIT and elsewhere, and how faculty participants are able to integrate some of the material into their curriculum. For the summer sabbatical, success will be measured by the number of proposals, or publications, or courses developed by the visiting faculty in collaboration with the CBMM host faculty. If the visiting faculty is a junior faculty, the role of the summer sabbatical in obtaining tenure would be assessed.

The number of viewers for each lecture posted on the website will be monitored and reported. The number of invitations received by the various CBMM faculty to speak at minority serving institutions will also be monitored.

1c. Problems and challenges

One issue we have encountered is logistical. For example, the 2014 CBMM summer course held at Woods Hole overlaps with the CBMM summer internship program. We have tried to work around the issue this year by delaying the start of the summer program for some of the students. In the future we plan to schedule the summer course at a later time to avoid an overlap with the summer research internship program.

We also realize that to expand the number of participants in the summer program we will need to recruit faculty hosts that are outside of CBMM. These faculty will not receive funding from CBMM but their research will be relevant to the CBMM research goals, and the research experience and mentorship they will offer the summer students will be invaluable. Every student in the CBMM summer program will be exposed to the CBMM research goals through faculty seminars given by CBMM members.

We also need to advertise our programs more broadly to ensure that we reach out to a very diverse population. For example we will work with MIT's Student Disabilities Services office to advertise the summer internship program to students with disabilities at other institutions.

Finally planning a sabbatical requires time. Although a summer sabbatical eliminates the need to require release time from teaching and is easier to plan, it still requires advanced preparation and careful planning. It may be necessary to host fewer faculty in the first two years and increase the number of faculty hosted by the Center in later years.

2a. Diversity Activities

I – Workshop

The first CBMM Workshop on Broadening Participation in the Science of Intelligence was held at MIT from January 8 to 11, 2014. Thirteen faculty from all six CBMM diversity partner institutions for the broader participation of women and minorities (BPWM) attended the workshop, including Howard University, Wellesley College, University of Puerto Rico at Rio Piedras (UPR), Hunter College and Queen's Colleges, and the Universidad Central del Caribe, as well as thirteen CBMM faculty from MIT, Harvard and Cornell.

Activity Name	Discussion Panel on Applying to graduate school
Dates	January 2014
Led by	Mandana Sassanfar (MIT)
Participants	Howard University, Wellesley College, University of Puerto Rico at Rio Piedras (UPR), Hunter College and Queen's Colleges, and the Universidad Central del Caribe and CBMM faculty from MIT, Harvard and Cornell
Number of Participants	26 (13 From Diversity Partner Institutions, 13 from Primary Academic Institutions)

One of the main purposes of this first workshop was to start an in-depth conversation about the specific needs for novel courses and curriculum and challenges at each partner institution, to identify areas for educational and research collaborations across institutions, discuss the content and format of future workshops, the training of students, and to plan the summer sabbaticals and seminar series. Five thrust leaders (the theory of intelligence was not included) gave research presentations and faculty from each partner institution talked about their curriculum and current research programs, and some of the challenges they face in bridging the gap between neuroscience and computer science. Within the workshop, a Discussion Panel on Applying to Graduate School was incorporated to assist partner institutions and their students who wish to pursue graduate education.

The content and format of the next workshop was discussed at length. The future January workshops will be named "Quantitative and Computational NeuroScience Workshop" (QCNS workshop) and introduce students to experimental tools such as fMRI, and MEG and computational tools such as MATLAB and Python to prepare them for research internships. The faculty from the diversity partner institutions visited the fMRI facilities and attended a hands-on four-hour session on the use of MATLAB and machine learning in Neuroscience. Faculty from Puerto Rico, Howard University, and Hunter joined their students who were enrolled in the 6-day Quantitative Biology workshop and participated in the hands-on activities. Three of the students who attended the workshop have been selected to participate in the 2014 CBMM summer internship program.

II – Summer internship program

The CBMM summer program has 12 funded slots, covering travel, stipend, housing, meals, social activities, a GRE prep course, poster session, and a lecture series. Each diversity partner institution was given two slots (8 in total-2 for each of CUNY, Puerto Rico, Wellesley College and Howard). The students were selected by faculty members from our diversity partner institutions. The remaining 4 slots were filled through a competitive online application process. Competitive applicants had a minimum GPA of 3.5, three strong letters of references and some prior research experience.

Due to cost sharing opportunities the 2014 CBMM summer program is hosting 16 students. 13 are fully or partially funded by NSF and 3 are international students who are funded by separate sources. All of the students have been placed with faculty who are either CBMM PIs or whose research is relevant to the CBMM research goals. Participants in the CBMM summer program meet as a group twice a week to attend lectures by CBMM faculty, give presentations about their research to their peers, or attend academic seminars on applying to graduate school, applying for fellowships, etc.

Activity Name	CBMM Summer Research Program Mentors
Dates	June – August 2014
Led by	Boris Katz (MIT) Tomaso A. Poggio (MIT) Rebecca Saxe (MIT) Elizabeth Spelke (Harvard) Josh Tenenbaum (MIT) Patrick Winston (MIT) Nancy Kanwisher (MIT) Matt Wilson (MIT) Gabriel Krieman (Harvard/Children's Hospital Boston)
Participants	Students accepted to the CBMM Summer Research Program
Number of Participants	16

Students in the 2014 CBMM summer program come from very diverse backgrounds and are majoring in neuroscience, psychology, computer science, music, electrical engineering, physics, biology, biochemistry or art history.

The demographics of the 13 CBMM funded students is as follows: 8 females and 6 males; 5 Hispanics, 2 African-American, 1 Asian, 6 Caucasian, 2 non-traditional students including a single mother. At least seven are receiving 100% financial aid from their undergraduate institution.12 students will work in MIT research labs and 2 in Harvard labs. The host faculty include 9 CBMM PIs including three thrust leaders. Some of the students have been placed in labs that are not funded by CBMM, but whose research is relevant to the CBMM mission.

The students will spend 9 to 10 weeks conducting full-time supervised research and participate in activities specifically designed to prepare them for advanced undergraduate and graduate level classes and research. Students are asked to view the research focus of the various thrusts by visiting the CBMM website and are placed in labs according to their interests, where they are closely mentored by a graduate student or postdoc, in addition to their faculty advisor. The students will also attend a GRE prep course (four 4-hour session on two separate week-ends)

The summer students will present their research at a public poster session at the end of the program (August 7). Faculty from our diversity partner institutions have been invited to attend the poster session and the program will cover up to $250 per faculty member to help with travel costs. In addition, Lizanne Destefano, our external reviewer, will also attend and will hold a focus group with the CBMM summer students on August 8.

Summer students are also encouraged to present their posters at ABRCMS, SACNAS or at the Society for Neuroscience meeting. We plan to follow the career path of all CBMM summer students for at least 5 years after they leave the summer program. We plan to create a short video to advertise the summer program. The video will contain interviews with the summer students and their PIs.

III – Lecture series

Seven lectures have been scheduled to be videotaped this summer. These introductory lectures will be given in front of the 2014 CBMM summer students and will be videotaped for broader accessibility and posted on the CBMM website. Professors Ed Boyden, Nancy Kanwisher (Research Thrust Leader), Rebecca Saxe, Matt Wilson, L. Mahadevan, Patrick Winston and Gabriel Kreiman (Research Thrust Leader) will be speakers. Their lectures will give an overview of their field as well as examples of their own research. These lectures will be accessible to undergraduate audiences and will later be complemented by more advance lectures as the CBMM curriculum is developed.

IV – Faculty Summer sabbatical

The summer faculty sabbatical program provides faculty from our diversity partner institutions an opportunity to spend up to 12 weeks at MIT or Harvard in the lab of a CBMM faculty member. During this time faculty can learn new techniques, use the state of the art facilities to advance their own research, and explore opportunities for future collaborations and joint grant proposals. The visiting faculty can bring one or two students from their lab (graduate or undergraduate students). These students will be funded by the CBMM summer internship program and fully included in all of the activities offered through the CBMM summer program.

For Year 1 Professor Maria Bykhovskaia and one of her graduate students from the Neuroscience department at the Universidad Central del Caribe are spending 8 weeks in the lab of Matt Wilson.

V – Travel to minority-serving institutions and Presence at minority conferences

Between September 2013 and May 2014 Dr. Sassanfar traveled to the University of Puerto Rico at Rio Piedras, the Universidad Central del Caribe in Bayamon, Hunter College, Queen's College, Howard University, North Carolina Central University, and North Carolina A&T, to meet with faculty and students, and to talk about opportunities at CBMM. At each institution she met with faculty in neuroscience and computer science and with students interested in summer internships at CBMM. At Howard which she visited in May she was also met with the two students who have been selected to participate in the 2014 CBMM summer program and spoke with their faculty mentors. She also attended the annual meeting of the Society for the Advancement of Chicanos and Native Americans in Science (SACNAS) and the Annual Biomedical Conference for Minority Students (ABRCMS). In addition she was a panelist at the Biomedical Research Conference in Boston.

CBMM will also have representation at the 2014 SACNAS and ABRCMS annual meetings. CBMM will share a booth with the MIT Department of Brain and Cognitive science. Dr, Sassanfar and a grad student will be at hand to talk to students and faculty about the various CBMM programs and provide brochures. In addition we have initiated discussions with the organizers of SACNAS and ABRCMS about the possibility of scheduling a lecture by a CBMM faculty at their 2015 or 2016 annual conference. Dr. Sassanfar will also continue to travel to minority serving institutions and diversity partner institutions to talk to interested students about educational and research opportunities offered by CBMM.

2b. Impact of Diversity Activities

I – Workshop

The impact of the CBMM Workshop on Broadening Participation in the Science of Intelligence remains to be seen. Due to this being the first workshop, the Center expects to gauge the long-term impact in the follow-up with students who move on to graduate education programs in the field of neuroscience and computer science.

II – Summer internship program

Students in the 2014 CBMM summer program come from very diverse backgrounds—we expect to achieve a number of goals through their participation in the summer internships. In the short-term, students will be prepared for graduate education and research in the field of neuroscience and computer science. Students will have greater knowledge on how to undertake independent research projects, how to present research results to investigators, peers, and the public. Long-term impacts of the summer research program will be evaluated in two or more years, when we can evaluate the percentage of students accepted into graduate programs.

III – Lecture series

The lecture series planned for the CBMM summer students will be utilized in the long term to provide partner institutions with a foundation for understanding the basics of neuroscience, computer science, and the field of intelligence. These introductory lectures will also help future applicants to the CBMM summer program select research topics and host laboratories. Faculty at Howard University and the University of Puerto Rico plan to use these lectures to introduce their students to the various CBMM research activities, as well as require students to view these lectures before the CBMM faculty travel to these universities to give research seminars.

IV – Faculty Summer sabbatical

The impact of the sabbatical program will be further evaluated in future reporting periods. The Center expects that the participating faculty members will successfully apply for grants and increase the number of publications related to their research in addition to providing students from partner institutions with summer internships.

V – Presence at Minority conferences

CBMM will evaluate the impact of conferences in future reporting periods but expect increases in the number of applicants and participants in the outreach workshop and summer research internship.

2c. Diversity Progress

The CBMM diversity program seeks to create a diverse community of scientists that includes the broader participation of women, minorities, and individuals with disabilities interested in the study of human intelligence. In the strategic plan, 10 of milestones were created for years 1–3 and will be assessed to meet this goal.

Milestone #1: Establish an annual summer program for 10 undergraduates from minority serving institutions (MSI), and non-research intensive institutions, to help prepare them for graduate school.

Progress toward Milestone #1: The Center has exceeded the number of expected participants attending the summer program and expects to increase the number of students in the upcoming years. CBMM is hosting 16 students, 13 are fully or partially funded by NSF and 3 are international students who are funded by separate sources.

Milestone #2: Post up to eight videotaped lectures given by CBMM faculty to undergraduates from institution for BPWM on the CBMM website for broad access.

Progress toward Milestone #2: Seven lectures during the Summer of 2014 will be videotaped and made accessible to diversity partner institutions. Plans to tape additional lectures are expected for Years 2–3.

Milestone #3: Establish an annual 6-day workshop for 25 students and faculty from MSI and non-research intensive institutions.

Progress toward Milestone #3: The CBMM Workshop on Broadening Participation in the Science of Intelligence was held in January 2014. The outcomes of the workshop will be further evaluated in the upcoming years of the Center, however three of the students who attended the 2014 workshop have been selected to participate in the 2014 CBMM summer internship program. Future workshops will be named "Quantitative and Computational NeuroScience Workshop" (QCNS workshop) and, as previously mentioned, will introduce students to experimental tools and computational tools to prepare them for research internships and graduate education.

Milestone #4: Establish on average one research collaboration per year between CBMM PIs and CBMM faculty at BPWM partner institutions.

Progress toward Milestone #4: The Center has one faculty member and student from a partner institution working with a CBMM faculty member for Summer 2014.

Milestone #5: Work towards submitting two grants to support collaborative research or educational endeavors between CBMM faculty and faculty at BPWM partner institutions.

Progress toward Milestone #5: This cannot be evaluated for the first year, however we plan to evaluate the submission of grants from faculty and partner institutions in Years 2–3.

Milestone #6: Submit three co-authored publications that include summer students or faculty from BPWM partner institutions.

Progress toward Milestone #6: This cannot be evaluated for the first year, however we plan to evaluate co-authored publications by BPWM institutions in Years 2–3.

Milestone #7: Summer students will have presented up to 10 posters at national (for example ABRCMS, SACNAS, AAAS) or international meetings (for example Society for Neuroscience).

Progress toward Milestone #7: This will be evaluated in Years 2–3 after the summer students have completed their research internships.

Milestone #8: At least three CBMM faculty members per year will visit institutions for BPWM and meet with undergraduate and graduate students.

Progress toward Milestone #8: This will be evaluated in Years 2–3 after the CBMM faculty mentor summer students from our diversity partner institutions. Additionally, we plan to have faculty participating in the lecture series visit with our diversity partners.

Milestone #9: Two to three faculty members from BPWM partner institutions will be invited to give seminars to CBMM faculty at MIT or Harvard each year (up to nine by year 2017).

Progress toward Milestone #9: For Year 1, we have Prof. Bevil Conway (Wellesley) and Prof. Maria Bykhovskaia (Universidad Del Caribe) scheduled to speak at CBMM. We plan to increase the number of seminars given by faculty members from BPWM partner institutions in Years 2–3.

Milestone #10: Three faculty members from BPWM partner institutions will have spent a summer sabbatical in a CBMM PI's lab.

Progress toward Milestone #10: We have one faculty member: Prof. Maria Bykhovskaia (Universidad Del Caribe) doing a summer sabbatical with Matt Wilson (Associate Director). We plan to have faculty from partner institutions spend summer sabbaticals with CBMM investigators in Years 2–3.

2d. Future Plans

We plan to continue to strengthen and expand our programs to increase diversity in the field of the science of intelligence.

We are working with the office for usability and readability to ensure that our online application for the summer program is accessible to everyone including visually impaired students. We have also decided that at this early stage of the emergence of a new complex interdisciplinary field for the study of intelligence it will be important to include graduating seniors and post-baccalaureate students in the summer internship program. Therefore the summer program will not be strictly limited to undergraduates and will not be called a Research Experience for Undergraduates (REU) program.

We also plan to fund a few high school students to carry on research projects in CBMM labs and encourage them to present their research at regional, state and national science fairs. This will allow us to recruit undergraduates to CBMM very early in their scientific career. We have already identified a high school student to work in the lab of Matt Wilson.

We plan to take advantage of the offer by our diversity partner institutions to host the January workshops. Next year's workshop will be hosted by MIT to facilitate the development and design of the hands-on activities. In following years the workshop could be hosted at CUNY and Puerto Rico.

Dr. Sassanfar the Center's Director of Diversity and Outreach has been invited to be a member of the External Advisory committee on a new NIH-funded MBRS RISE proposal by Howard University. She has also written letters of support for separate proposals by Prof. Irving Vega from the University of Puerto Rico, and Prof. Kebreten Manaye from Howard University for the NIH-funded BP-ENDURE program which aims at increasing the number of URMs in neuroscience graduate programs. Profs. Vega and Manaye are both active CBMM diversity faculty partners. If their proposals are funded it will provide even more access to CBMM activities by a broader and larger group of students from these institutions, and will foster a deeper engagement and cooperation between CBMM research faculty and faculty at these institutions.

New initiatives

In the future we would like to create a visiting scholar's program: We hope to identify each year one or two very talented undergraduates who are either URMs or have disabilities to spend the spring semester studying computational neuroscience at MIT or Wellesley College, attend the January workshop and participate in the summer research program in a CBMM-affiliated lab at Harvard or MIT. This would provide an exceptionally integrated approach to training and preparing these students for graduate-level courses and research in the complex field of the science of intelligence.

VII. MANAGEMENT

1a. Management Strategy

Leadership and Management Mission Statement: Envision and enable the Center's mission through inclusive and transparent decision-making; inspire Center participants to work as a team with a focus on making progress on the CBMM challenge; test management and organizational strategies—beyond usual scientific funding practice—for effective collaborative research of sizable teams on deep scientific problems.

The CBMM is based at MIT with close organizational and infrastructural connections to Harvard. The CBMM activities are organized by a Director, who is assisted by an Associate Director from each of the two primary institutions. The Center also has leaders in Education, Research, Diversity, and Knowledge Transfer and a Research Coordinator and an External Evaluator. In addition, we have a full-time staff including a Managing Director (at Harvard) and a Center Manager (at MIT.) Other administrative staff are a staff assistant (MIT), and a half-time coordinator (Harvard); we are seeking a web developer cum videographer for administration and Education.

The Center leadership is largely unchanged since the submission of the Strategic Plan. However, responsibilities have been optimized: in particular Education is now co-led by Ellen Hildreth (Wellesley) and Haym Hirsh (Cornell), Knowledge Transfer will continue to be led by Mahadevan (Harvard) with Boris Katz (MIT) sharing the leadership. Thrust 3 is now co-led by Shimon Ullman and Boris Katz. The updated organizational chart below shows the full management team.

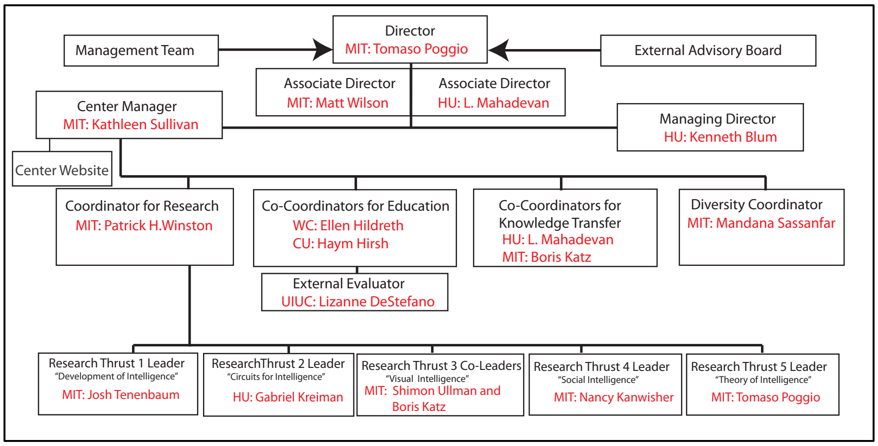

1b. Management Expectation and Evaluation Indicators

We already described the content of this section in the part of the Context statement titled Performance and management indicators.

1c. Management Progress

The Center has accomplished early management goals. Referring to the ***Expectation*** indicators cited in section 1b:

- The primary direction of the CBMM is unchanged since the proposal stage, but areas of emphasis have evolved. First, we have focused on building the cohesion of the management team through a series of meetings to discuss larger principles, detailed management mechanisms, and the myriad organizational matters that our new STC confronts. Second, we have emphasized developing the cohesion of the Boston-area cluster of scientists, through weekly scientific discussions of the 5 research thrusts and the many specific, collaborative projects within each thrust. Third, we have concentrated our initial efforts on a small number of high-impact initiatives, such as research conferences and our upcoming summer course at Woods Hole MBL. Finally, having established our well-knit core, we are beginning to welcome interested students, postdocs, and faculty into our research discussions.
- Decision-making has been distributed to the thrust leaders, the leaders of education, knowledge transfer, outreach, and the institutional associate directors at Harvard and MIT.
- The management team has had open discussions of complicated issues such as selecting summer course participants, the composition of the external advisory board, and opening the CBMM to new participants.
- The management team has made flexible arrangements to ensure persistent future support for graduate students.

Referring to the ***Evaluation*** indicators cited in section 1b:

- All of the Center's projects fit well within the research thrusts we have organized. The Knowledge Transfer, Education, and Outreach efforts contribute directly to the Center's objectives.
- The spirit of collaboration is strong in our Center. The CBMM projects link at least two research groups focused on some aspect of intelligence.
- A major accomplishment in this first year was a stage of intensive communication within the Center, comprised of weekly meetings of PIs, students and postdocs (from Harvard, MIT and Wellesley) and regular meetings of the management team. We also generated a Strategic Plan for the Center during a February Strategic Planning Workshop at MIT, which leveraged material developed during the first 4 months of intensive discussions and interactions within the Center.
- The growth of our community has been promoted by two important mechanisms. First, two workshops have already been organized by the Center, one which took place in November, and one which is forthcoming. Second, CBMM will host a 2-week summer course at Woods Hole on precisely the topic of our STC. In its inaugural year, this course had 140 competitive applicants from around the world, from whom we selected 25. The course lectures will be streamed and archived, so we anticipate that the total impact will be much wider.

- In January we had our first Diversity workshop, which brought thirteen faculty from all six CBMM diversity partner institutions to MIT. This summer we will host 12 undergraduates, from our Diversity Partner institutions, in laboratories at MIT and Harvard.

The Management Team meets approximately twice a month, depending upon academic calendar. The table below lists the Management Team meeting dates and members. Minutes of all meetings are distributed to committee members and archived in a central file repository.

Activity Name	Management Team Meetings
Dates	April 26, 2013; June 20, 2013; Oct. 9, 2013; Nov. 15, 2013; Jan. 7, 2014; Feb. 14, 2014; March 14, 2014; March 28, 2014; April 4, 2014; April 11, 2014; May 2, 2014; May 23, 2014
Led by	Tomaso Poggio (MIT) and Kenneth Blum (Harvard)
Participants	Center Director, Managing Director, Associate Directors, Coordinators for Education, Coordinator for Research, Coordinator for Diversity, Research Thrust Leaders, External Evaluator, and the Center Manager
Number of Attendees (varies)	MIT: Tomaso Poggio, Matt Wilson, Patrick H. Winston, Boris Katz, Mandana Sassanfar, Josh Tenenbaum, Nancy Kanwisher, Shimon Ullman, and Kathleen Sullivan Harvard: Kenneth Blum, L. Mahadevan Cornell U: Haym Hirsh Wellesley: Ellen Hildreth UIUC: Lizanne DeStefano

1d. Management Problems

The most significant management challenge for CBMM has been our geographic distribution. We have addressed this challenge successfully by three mechanisms: (1) alternating management meeting locations between MIT and Harvard, (2) using conference calls and videoconferencing to link to Haym Hirsh (Cornell), and Lizanne DeStefano (Illinois), and (3) appointing Ellen Hildreth (Wellesley) to the management team as co-director of Education.

2. Management Communications Systems

As already described in our Strategic Plan, there has been an evolution from the time of the proposal about the organization of our Center. One of its key principles is now “centerness”: our CBMM will fund and nurture collaborative projects that cannot be done in a single lab with typical single investigator grants but only in a Center like ours. This means that no single PI has students or postdocs funded by CBMM or is funded directly. Instead, collaborative projects have priority for funding (projects have to be collaborations between two or more PIs).

For the same reason a key endeavor of the Center is to develop a set of databases (mainly images and videos) that will be used across different labs and techniques to measure performance of the mind and of the brain for recognition/perception of objects, of people, of interactions between people and objects, of people’s actions and of people’s social interactions; the same data will be used to measure how well our models and our computer systems perform in absolute and relative terms (see earlier section).

As a general policy thrust leaders are in charge of hiring postdoc/students for the cooperative projects in their thrust -- of course with the help of the PIs in the thrust, the Research Coordinator, and the director. This policy gives a significant responsibility to the thrust leaders, while bypassing PIs in direct funding decisions. We believe this is a small price to pay for an effective organization of a relatively large collaborative effort in basic research. Of course, large collaborations on projects of a more engineering flavor—such as the Manhattan project and the construction of large particle accelerators—have been successfully achieved in the past

We have established an Ethics Task Force to address ethics issues and ethics training, and a Postdoc Training Faculty Advisory group to organize a variety of meetings to facilitate communication among Center researchers; both of these groups have been meeting on an as needed basis. In the next reporting period the Center hopes to extend communications systems to Center graduate students through the formation of our Student Leadership Council. This group will provide a mechanism for Center students to interface with the Management Team and will allow students to provide input on the education and outreach initiatives of the Center.

The CBMM Annual Retreat is planned for January 2015 and will provide an opportunity for all Center members including students, postdoctoral fellows, researchers, and staff to gather, discuss research and other accomplishments of the past year, to plan initiatives for the upcoming year, and to engage in professional development and student leadership activities. This three-day retreat will also give the Center an opportunity to receive input from our External Advisory Committee on all aspects of the Center including research, management, education, knowledge transfer, and diversity.

Our current Management meetings have been effective at maintaining good communication throughout the Center and we have not encountered any significant problems in achieving Center integration in the first reporting period.

3. Center External Advisory Committee Members

We have prepared a list of potential members and submitted it to NSF for approval.

4. Center Strategic Plan

We have a version of the strategic plan ready to be finalized.

VIII. CENTER-WIDE OUTPUTS AND ISSUES

1a. Publications

Peer Reviewed Publications

Anselmi F, Leibo JZ, Rosasco L, Mutch J, Tacchetti A, Poggio T. Unsupervised Learning of Invariant Representations in Hierarchical Architectures. Nature. Submitted.

Baldauf D, Desimone R. Neural mechanisms of object-based attention. Science. 2014;344(6182):424-7.

Gerstenberg T, Goodman ND, Lagnado DA, Tenenbaum JB. From counterfactual simulation to causal judgment. Cognitive Science Conference Proceedings. 2014.

Jara-Ettinger J, Gweon H, Tenenbaum, JB, Schulz LE (2014). I'd do anything for a cookie (but I won't do that): Children's understanding of the costs and rewards underlying rational action. 36th Annual Proceedings of the Cognitive Science Society.

Jara-Ettinger J, Schulz LE. Running to do evil: Costs incurred by the perpetrator affect moral judgment. Kim N. 36th Annual Proceedings of the Cognitive Science Society. Awaiting publication.

Krompass D, Jiang X, Nickel M, Tresp V. Factorizing Probabilistic Databases. Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining. Submitted.

Li Y, Hou X, Koch C, Rehg J, Yuille AL. The Secrets of Salient Object Segmentation. Proceedings of Computer Vision and Pattern Recognition (CVPR). Submitted.

Magid R, Schulz LE. Imagination and the generation of new ideas. Sheskin M. Cognitive Development. Awaiting publication.

Nickel M, Jiang X, Tresp V. Learning from Latent and Observable Patterns on Multi-Relational Data. Advances in Neural Information Processing Systems 27 (NIPS 2014). Submitted.

Nickel M, Ciliberto C, Rosasco L. Learning Vector-Valued Functions via Sparse Multilinear Operators. Advances in Neural Information Processing Systems 27 (NIPS 2014). Submitted.

Scott K, Schulz LE. Interhemispheric integration of visual concepts in infancy. 36th Annual Proceedings of the Cognitive Science Society. Awaiting publication.

Siegel MH, Magid R, Tenenbaum JB, Schulz LE. Black boxes: Hypothesis testing via indirect perceptual evidence. Proceedings of the Cognitive Science Society. 2014.

Siegel M, Tenenbaum J, Schulz LE. Black boxes: Hypothesis testing via indirect perceptual evidence. Magid R. 36th Annual Proceedings of the Cognitive Science Society. Awaiting publication.

Tsividis P., Gershman S., Tenenbaum JB, Schulz LE (2014). Information selection in noisy environments with large action spaces.36th Annual Proceedings of the Cognitive Science Society. Awaiting publication.

Tsividis P., Tenenbaum J, Schulz LE. Information selection in noisy environments with large action spaces. Gershman S. 36th Annual Proceedings of the Cognitive Science Society. Awaiting publication.

Wu Y, Baker CL, Tenenbaum JB, Schulz LE. Joint inferences of belief and desire from facial expressions. 36th Annual Proceedings of the Cognitive Science Society. Awaiting publication.

Yan P, Schulz LE. Preschoolers expect others to learn rationally from evidence. Magid R. 36th Annual Proceedings of the Cognitive Science Society. Awaiting publication.

Zhang C, Evangelopoulos G, Voinea S, Rosasco L, Poggio T. A Deep Representation for Invariance And Music Classification. Proc. IEEE 2014 International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2014) [Internet]. 2014. http://arxiv.org/abs/1404.0400v1

Book and Book Chapters

Gerstenberg T, Ullman T., Kleiman-Weiner M, Lagnado DA, Tenenbaum JB. Wins above replacement: Responsibility attributions as counterfactual replacements. Cognitive Science Conference Proceedings. 2014.

Hildreth E, Sassanfar M. Workshop on Broadening Participation in the Science of Intelligence, Summary Report. In 2014.

Kreiman G, Rutishauser U, Cerf M, Fried I. “The next ten years and beyond.” In: Single Neuron Studies of the Human Brain. MIT Press: June, 2014.

Kreiman G. “Principles of neural coding.” In: Computational Models of Visual Object Recognition. CRC Press: October, 2013.

Kreiman G. “Neural correlates of consciousness: perception and volition.” In: Cognitive Neuroscience. MIT Press: June, 2014.

Mormann F, Ison M, Quiroga R, Koch C, Fried I and Kreiman G. “Visual cognitive adventures of single neurons in the human medial temporal lobe.” In: Single Neuron Studies of the Human Brain. MIT Press: April, 2014.

Rutishauser U, Cerf M, Kreiman G. “Data analysis techniques for human microwire recordings: spike detection and sorting, decoding, relation between units and local field potentials.” In: Single Neuron Studies of the Human Brain. MIT Press: June, 2014.

Siddharth N, Barbu A, Siskind JMark. Seeing What You're Told: Sentence-Guided Activity Recognition In Video. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition [Internet]. 2014;:8. http://0xab.com/papers/cvpr2014a.pdf

Non-Peer Reviewed Publications

Leibo J.Z, Liao Q, Anselmi F, Poggio T. The invariance hypothesis implies domain-specific regions in visual cortex. BioRxiv, [Internet]. 2014.http://biorxiv.org/content/early/2014/04/24/004473.full-text.pdf+html

1b. Conference and Poster Presentations

Berzak, Y., Reichart, R., Katz, B. Reconstructing Native Language Typology from Foreign Language Usage. Eighteenth Conference on Computational Natural Language Learning, Baltimore, MD. June 26–27, 2014.

Boyden, E. Invited Presentation. “Optical Tools for Mapping and Engineering the Brain.” Lester Wolfe Workshop on Laser Biomedicine, Massachusetts General Hospital, Boston, MA. April 8, 2014

Boyden, E. Invited Presentation. “Tools for recording and controlling neural activity.” Neuronal Circuits Meeting, Cold Spring Harbor, NY. April 2–5, 2014

Boyden, E. Invited Presentation. “Optogenetics: Tools for Analyzing and Controlling Brain Circuits with Light.” Royal Swedish Academy of Sciences Symposium: Optogenetics, Stockholm, Sweden. December 13, 2013.

Boyden, E. Invited Presentation. “Optical, Molecular, and Robotic Tools for Integrative Single Cell Analysis.” Society for Neuroscience Symposium, San Diego, CA. November 9–13, 2013

Boyden, E. Invited Presentation. “Optogenetics: Tools for Controlling Brain Circuits with Light.” Brain Prize Meeting, Hindsgavl Castle, Denmark. October 21, 2013

Boyden, E. Invited Presentation. “Engineering the Brain.” EmTech 2013 Conference, Cambridge, MA. October 9, 2013.

Boyden, E. Invited Presentation. “Technologies for Analyzing and Engineering Brain Computations.” MIT/MGH Initiative Symposium in Neuroscience, Cambridge, MA. September 6, 2013

DeStefano, L., Tillman, A., Lessons learned from evaluating a multi-site National Science Foundation Science Technology Center. American Evaluation Association Annual Meeting, Washington, D.C. October 18, 2013

DeStefano, L., Tillman, A., Panel: Content, Pedagogy, and Diversity: Evaluating STEM (Science, Technology, Engineering, and Mathematics) Programs Using the Values-engaged, Educative Approach. American Evaluation Association Annual Meeting, Washington, D.C. October 18, 2013

DeStefano, L., Tillman, A., Attending to culture and diversity evaluations of undergraduate and graduate Science, Technology, Engineering, and Mathematics (STEM) research traineeships. American Evaluation Association Annual Meeting, Washington, D.C. October 19, 2013

Evangelopoulos, G., Rosasco, L., Poggio, T.A., A Deep Representation for Invariance and Music Classification. 2014 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Florence, Italy. May 9, 2014

Gao, T., Harari, D., Tenenbaum, J., Kanwisher, N., What are you looking at?: The acuity of joint attention. Vision Science Society, St. Pete Beach, FL. May 20, 2014

Isik, L., Meyers, M., Leibo, J.Z., Poggio, T.A., A spatiotemporal profile of invariant object recognition in the human visual system. Society for Neuroscience, San Diego, CA. November 9–11, 2013.

Isik, L., Han, Y., Poggio, T.A., Decoding invariant visual information with MEG sensor and source data. Neural Information Processing Systems (NIPS) Workshop on Machine Learning and Interpretation in Neuroimaging, Lake Tahoe, NV. December 9–10, 2013.

Kreiman, G., Nassi, J.J., Gomez-Laberge, C., Born R., Increasing the visuotopic extent of normalization through cortico-cortical feedback. Society for Neuroscience, San Diego, CA. November 1, 2013

Kreiman, G., Miconi, T., Singer, J., Normalized Hebbian learning develops both simple and complex receptive fields from naturalistic video. Computational and Systems Neuroscience (COSYNE), Salt Lake City, UT. March 1, 2014

Linderman, S., Jonas, E., Kording, K., Adams, R., Workshop on Discovering Structure in Neural Data. Computational and Systems Neuroscience (COSYNE), Salt Lake City, UT. March 2, 2014

Pashkam, M.V., Nakayama K., Cormiea, S., Fast mirroring of an opponent's action in a competitive interaction. Vision Sciences Society, St. Pete Beach, FL. May 17, 2014.

Poggio, T.A., Invited Presentation. “The Center for Minds, Brains, and Machines.” MIT Research & Development Conference 2013, Cambridge, MA. November 13, 2013

Poggio, T.A., Invited Presentation. “Brains, Minds, and Machines.” Rome Science Festival: Human Language and Machine Learning Language, Rome, Italy. January 23, 2014.

Scott, K., Interhemispheric integration in infancy: split-brain babies? Towards a Science of Consciousness, Tucson, AZ. April 22, 2014

Spelke, E., Invited Presentation. “Core Social Cognition.” Ecole Normale Supérieure, Institute Jean Nicod, Paris, France. September 16, 2013

Spelke, E., Invited Presentation. “Core social cognition: Insights from infants.” Simons Foundation Autism Research Initiative Annual Meeting, New York, NY. September 29, 2013

Spelke, E., Invited Presentation. “What makes humans different?” Mind Reading: Human Origins and Theory of Mind, La Jolla, CA. October 19, 2013

Spelke, E., Invited Presentation. “Core cognition: Enhancing early education by building on children.” Latin American School for Education, Cognitive and Neural Sciences Punta del Este, Uruguay. March 14, 2014

Ullman, T., Learning physics from dynamical scenes. 36th Annual Conference of the Cognitive Science Society, Quebec City, Canada. July 24, 2014

Ullman, S., Invited Presentation. “From simple innate to complex visual concepts.” The Annual Meeting of The Israeli Center of Research Excellence (iCORE) in the Cognitive Sciences: Challenges and Debates in the Frontiers of Brain and Cognition Research. The Weizmann Institute of Science, Rehovot, Israel. December 25–26, 2013.

Upcoming conference presentations

Tacchetti, A., Rosasco, L., Villa, S., Regularization by Early Stopping for Online Learning Algorithms. Neural Information Processing Systems (NIPS), Montreal, Quebec, Canada. December 8–12, 2014

Poster Presentations

Boyden, E., Moore-Kochlacs, C., Principles of high–fidelity, high–density 3–d neural recording. Cognitive Neuroscience Society (CNS) 2014.

Deen, B., Kanwisher, N., Saxe, R., Exploring superior temporal sulcus responses and patterns with a broad set of naturalistic stimuli. Society for Neuroscience. November 10, 2013.

Harari, D., Gao, T., Tenenbaum, J., Kanwisher, N., What are you looking at?: The acuity of joint attention. Vision Sciences Society Annual Meeting. May 20, 2014.

Linderman, S., Nemati, S., Chen, Z. A Probabilistic Modeling Approach for Uncovering Neural Population Rotational Dynamics. Computational and Systems Neuroscience (COSYNE). February 27, 2014.

Marantan, A., Catching on to How to Catch a Ball. Center for Brain Science - Harvard University - Annual Retreat. May 16, 2014.

Meyers, E., Borzello, M., Freiwald, W., Tsao, D., Decoding what types of information are in the macaque face patch system. Computational Systems Neuroscience. February 28, 2014

Nickel, M., Tensor, V., Factorization for Large-Scale Relational Learning. New England Machine Learning Day 2014. May 13, 2014

Poggio, T.A, Tan, C., Singer, J., Serre, T., Sheinberg, D., Invited Presentation. "Neural representation of action sequences: How far can a simple snippet-matching model take us." Neural Information Processing Systems (NIPS) Foundation 2013. December 6, 2013.

Spelke, E., Spokes, A., Children's Expectations and Understanding of Kinship as a Social Category. 8th Biennial Meeting of the Cognitive Development Society. October 19, 2013.

Spelke, E., Powell, L., Young infants use imitation to infer the social preferences of imitators but not targets. 8th Biennial Meeting of the Cognitive Development Society. October 19, 2013.

Spelke, E., Dillon, M., Functional and spatial dissociation in the brain systems encoding object shape and direction. Meeting of the Cognitive Development Society. April 8, 2014.

Upcoming Conference Presentations:

Meyers, E., Schafer, R., Zhang, Y., Poggio, T.A, Desimone, R., Visual selectivity and attentional modulation in V4, IT and the posterior lateral pulvinar. Society for Neuroscience. November 12, 2014.

Tacchetti, A., Isik, L., Invariant representations for action recognition in the human visual system. Society for Neuroscience. November 15–19, 2014.

1c. Other Dissemination Activities

Other Presentations and Lectures

Boyden, E., Invited Seminar. “Tools for Mapping and Engineering Brain Computations.” Honors Program, NYU. New York, NY. February 10, 2014.

Boyden, E., Invited Lecture. “Tools for Mapping and Engineering Brain Computations.” Yale, New Haven, CT.

Boyden, E., Invited Lecture. “Optogenetics: Tools for Mapping and Controlling Brain Dynamics.” CURE the Epilepsies: Frontiers in Research Seminar Series Albert Einstein College of Medicine, Bronx, NY. December 20, 2013.

Boyden, E., Invited Lecture. “Tools for Mapping and Engineering Brain Computations.” Herman P. Schwan Distinguished Lecture, University of Pennsylvania. Philadelphia, PA. September, 12, 2013.

Boyden, E., Invited Lecture. “Optogenetics.” Gabbay Award Lecture, Brandeis University. Waltham, MA. October 10, 2013.

Boyden, E., Invited Lecture. “Tools for Mapping Brain Computations.” Industry-Academy Symposium in CNS, Tel Aviv University, Israel. October 13, 2012.

Boyden, E., Invited Lecture. “Optogenetics and Other Tools for Controlling and Analyzing Neural Circuits.” Accelerating Translational Neurotechnology: Fourth Annual Aspen Brain Forum, Aspen, CO. September 18–20, 2013.

Evangelopoulos, G., Meeting. CBMM Postdoc Group Meetings, Cambridge, MA. February 10, 2014–May 31, 2014.

Evangelopoulos, G., Presentation. Invariant Representation Learning. CBMM Postdoc Group Meeting, Cambridge, MA. February 21, 2014.

Freiwald, W., Liebo, J., Invited Seminar. “On the neural mechanisms of face recognition: from experiments to theory.” CBMM Weekly Research Meeting, Cambridge, MA. April 18, 2014.

Gao, T., Invited Seminar. “Visual Roots of Social Cognition.” Brown Social Lunch Series, Providence, RI. December 5, 2013.

Katz, B., Invited Lecture. “Telling Machines about the World." CSAIL Alliance Program (CAP) 7th Annual Meeting, Cambridge CA. May 30, 2014.

Katz, B., Invited Presentation. “START Question Answering System." Bloomberg R&D Machine Learning Group, New York NY. April 29, 2014.

Katz, B., Invited Presentation. "Information Access using Natural Language." Advanced Technologies Centre, Infocomm Development Authority of Singapore (IDA), Cambridge MA. April 22, 2014

Katz, B., Lecture. Combining language and Vision Processing. Workshop on Broadening Participation in the Science of Intelligence. CBMM, Cambridge, MA. January 10, 2014.

Marantan, A., Presentation. Catching on to How to Catch a Ball. Harvard University Class: Neural Control of Movement, Cambridge, MA. May 15, 2014.

Nickel, M., Session Chair. New England Machine Learning Day 2014, Cambridge, MA. May 13, 2014.

Nickel, M., Lewis, O., Presentation. Towards Relational Scene Understanding. Machine Learning Group at University of Genoa, Genoa, Italy. April 5, 2014.

Penagos, H., Hale, G., Workshop. Real time neural decoding of hippocampal spiking activity. GE Analytics and CBMM collaboration meeting, Cambridge, MA. April 25, 2014.

Penagos, H., Hale, G., Seminar. Coordinated activity between retrosplenial cortex and hippocampus during awake replay. GE Analytics and CBMM collaboration meeting, Cambridge, MA. April 25, 2014.

Penagos, H., Hale, G., Seminar. Coordinated activity between retrosplenial cortex and hippocampus during awake replay. CBMM Postdoc Group Meeting, Cambridge, MA. April 18, 2014.

Penagos, H., Hale, G., Seminar. Hippocampus as a target structure for physiological examples of neural circuits solving problems. CBMM Weekly Research Meeting, Cambridge, MA. March 14, 2014.

Poggio, T.A., Invited Lecture. “Computational Approaches to mind and brain.” Johns Hopkins University, Baltimore, MD. May 8, 2014.

Poggio, T.A., Invited Lecture. “The Computational Magic of the Ventral Stream: a theory.” Neurological Institute of Columbia University Speaker Series, New York, NY. April 7, 2014.

Poggio, T.A., Invited Lecture. “Brains, Minds, and Machines.” Siemens Distinguished Speakers Series. April 8, 2014.

Poggio, T.A., Invited Workshop. “M-Theory.” Workshop on Learning Data and Representation: Hierarchies and Invariance, Cambridge, MA. November 22–24, 2013.

Poggio, T.A., Invited Lecture. “Object Recognition by Hierarchical Learning Machine.” Max Planck Institute for Biological Cybernetics, Tübingen, Baden-Württemberg, Germany. May 13, 2014

Poggio, T.A., Invited Lecture. University of California, Berkeley Mathematics Department, Berkeley, CA. August 29, 2013.

Poggio, T.A., Invited Lecture. “The Computational Magic of the Ventral Stream: sketch of a theory.” Meeting with Google, Mountain View, CA. August 27, 2013.

Rosasco, L., Seminar. Early Stopping for Online Learning Algorithms. Neural Information Processing System (NIPS) Workshop on Nonparametric Methods for Machine Learning, Lake Tahoe, NV. December 9, 2013.

Rosasco, L., Invited Seminar. “Piecewise Approximation for Learning Data Representation.” 8th International Conference CURVES and SURFACES, Paris, France. June 12–18, 2014.

Rosasco, L., Invited Seminar. “Piecewise Approximation for Learning Data Representation.” UCLA Computer Science Department Seminar Series, Los Angeles, CA. April 22, 2014.

Sadagopan, S., Seminar. Cortical processing of voices and faces: linking neural computations to mechanisms. Department of Biological Sciences, Purdue University, West Lafayette, IN. February 13, 2014.

Sadagopan, S., Seminar. Cortical processing of voices and faces: linking neural computations to mechanisms. Dept. of Otolaryngology, University of Pittsburgh, Pittsburgh, PA. January 23, 2014.

Schulz, L., Invited Workshop. “Goal-oriented hypothesis generation and imagination.” Eighth Biennial Meeting of the Cognitive Development Society: Computational Models of Cognitive Development, Memphis, TN. October 17–18, 2013.

Schulz, L., Invited Workshop. “The Origins of Inquiry: Inference and exploration in early childhood.” Max Planck Institute Developmental Workshop, Leipzig, Germany. March 26–28, 2014.

Schulz, L., Invited Presentation. “Inferential economics: Children's sensitivity to the cost and value of information NYU Cognition and Perception Colloquium, New York, NY. April 10–11, 2014.

Spokes, A., Elizabeth, S., Invited Seminar. "Children.” Boston Area Moral Cognition Group. Boston, MA. February 25, 2014.

Tang, H., Seminar. Recognition of occluded objects in human visual cortex. Kirby Neurobiology Center Lab Results Talk, Boston, MA. April 30, 2014.

Winston, P., Lecture. Introduction to Center for Brains, Minds, and Machines (CBMM). Workshop on Broadening Participation in the Science of Intelligence. CBMM, Cambridge, MA. January 10, 2014.

Winston, P., Invited Lecture. “Reflections on Intelligence, Creativity, and Design.” Design Class, MIT Department of Architecture. Cambridge, MA. April 29, 2014.

Winston, P., Invited Lecture. “Story Understanding and Persuasion.” Army Asymmetric Warfare Center, Ft. Meade, MD. April 25, 2014.

Winston, P., Invited Lecture. “The Center for Brains Minds and Machines.” MIT Alumni Association Talk, Washington, D.C. April 24, 2014.

Winston, P., Invited Lecture. “What's Next after Next.” MIT Industrial Liaison Program Symposium, Cambridge, MA. November 13, 2013.

2. Awards and Honors

	Recipient	Reason for Award	Award Name and Sponsor	Date	Award Type
1	Bergen, Leon		Doctoral Dissertation Improvement Grant: National Science Foundation (NSF)	May 2014	Scientific
2	Boyden, Ed		Herman P. Schwan Distinguished Lecture: University of Pennsylvania	September 2013	Scientific
3	Boyden, Ed		Theodore Koppanyi Lecturer: Georgetown University	March 2014	Scientific
4	Boyden, Ed		Duncan Lecturer: Northwestern University	May 2014	Scientific
5	Boyden, Ed		Schuetze Award in Neuroscience	June 2014	Scientific
6	Isik, Leyla		MIT McGovern Institute Gorenberg Fellowship	September 2013	Fellowship
7	Saxe, Rebecca		National Academy of Science Troland Award: National Academy of Sciences (NAS)	May 2014	Scientific
8	Schulz, Laura		Distinguished Scientific Award for Early Career Contribution to Psychology: American Psychological Association	January 2014	Scientific
9	Sompolinsky, Haim		Elected to the European Molecular Biology Organization: European Molecular Biology Organization (EBMO)	May 2014	Scientific
10	Sompolinsky, Haim		Mathematical Neuroscience Prize: Israel Brain Technology	October 2013	Scientific
11	Spelke, Elizabeth		Inaugural Prize in Cognitive and Psychological Sciences: National Academy of Sciences (NAS)	April 2014	Scientific
12	Tenenbaum, Josh		Elected Fellow of the Cognitive Science Society:	2013	Scientific
13	Ullman, Shimon		The Ian P Howard Memorial Talk: Center for Vision Research (CVR), University of Toronto.	October 2013	Scientific

3. Graduating Students

There were no graduating students during the reported period.

	Postdoctoral Fellow Name	Placement Type	Placement Institution/Organization
1	Chen, Zhe	Academia	New York University
2	Meyers, Ethan	Academia	Hampshire College

4a. General Outputs of Knowledge Transfer

	Patent Name and Inventor(s)/Author(s)	Inventor(s)/Author(s)	Number	Application Date	Receipt Date
1	Methods and Apparatus for Learning Representations:	Poggio, Tomaso A. Leibo, Joel Z.	14/231,503	March 31, 2014	

4b. Other Outputs of Knowledge Transfer

No other outputs to report.

5. CBMM Participants

First Annual Reporting Period: September 1, 2013 – May 31, 2014

Category: (a) undergraduate students, (b) graduate students, (c) faculty, (d) visiting faculty, (e) other research scientists, (f) postdoctorates, (g) pre-college students, (H) teachers, (i) educators and (j) other participants

Affiliates: those individuals who spend less than 160 hours, over a twelve month period, involved with Center initiatives

Department: academic department for participant, if applicable

Gender: Female, Male

Disability: Hearing Impairment, Visual Impairment, Mobility/Orthopedic Impairment, Other, None

Ethnicity: Hispanic or Latino, Not Hispanic or Latino

Race: American Indian or Alaskan Native, Asian, Black or African American, Native Hawaiian or Other Pacific Islander, White

Citizenship: U.S. Citizen, Permanent Resident, Other non-U.S. Citizen

* Indicates Affiliates

- - - Indicated that information has been redacted

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
1	Afshordi, Narges*	b	Harvard University	Psychology	- - -	- - -	- - -	- - -	- - -
2	Anselmi, Fabio	f	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
3	Baker, Chris*	f	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
4	Barbu, Andrei	f	Massachusetts Institute of Technology	Computer Science and Artificial Intelligence Laboratory	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
5	Bergen, Leon	b	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
6	Berzak, Yevgeni	b	Massachusetts Institute of Technology	Electrical Engineering & Computer Science	- - -	- - -	- - -	- - -	- - -
7	Blum, Kenneth	j-Managing Director	Harvard University	Center for Brain Sciences	- - -	- - -	- - -	- - -	- - -
8	Boyden, Ed	c	Massachusetts Institute of Technology	Media Lab, McGovern Institute, Biological Engineering, and Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
9	Buice, Michael	e	Allen Institute for Brain Science		- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
10	Chen, Zhe (Sage)	f	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
11	Cooper, Conisha	j	Harvard University	Psychology	- - -	- - -	- - -	- - -	- - -
12	Deen, Ben	b	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
13	Desimone, Robert	c	Massachusetts Institute of Technology	McGovern Institute for Brain Research	- - -	- - -	- - -	- - -	- - -
14	Destefano, Lizanne	c	University of Illinois at Urbana-Champaign	Educational Psychology	- - -	- - -	- - -	- - -	- - -
15	Dillon, Moira	b	Harvard University	Psychology	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
16	Evangelopoulos, Georgios	f	Massachusetts Institute of Technology	McGovern Institute for Brain Research	- - -	- - -	- - -	- - -	- - -
17	Freiwald, Winrich	c	The Rockefeller University		- - -	- - -	- - -	- - -	- - -
18	Gao, Tao	f	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
19	Gerstenberg, Tobias	f	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
20	Goodman , Noah*	c	Stanford University	Psychology	- - -	- - -	- - -	- - -	- - -
21	Harari, Daniel	f	Massachusetts Institute of Technology	BCS	- - -	- - -	- - -	- - -	- - -
22	Hildreth, Ellen C.	c	Wellesley College	Computer Science	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
23	Hirsh, Haym	c	Cornell University	Computer and Information Science	- - -	- - -	- - -	- - -	- - -
24	Isakov, Alex	b	Harvard University	Engineering and Applied Sciences	- - -	- - -	- - -	- - -	- - -
25	Isik, Leyla	b	Massachusetts Institute of Technology	Computational and Systems Biology	- - -	- - -	- - -	- - -	- - -
26	Jara-Ettinger, Julian	b	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
27	Kanwisher , Nancy	c	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
28	Katz, Boris	c	Massachusetts Institute of Technology	Computer Science and Artificial Intelligence Laboratory	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
29	Koch, Christof*	e	Allen Institute for Brain Science		- - -	- - -	- - -	- - -	- - -
30	Kreiman, Gabriel	c	Harvard University	Ophthalmology; Neurobiology	- - -	- - -	- - -	- - -	- - -
31	Lewis, Owen	b	Massachusetts Institute of Technology	Brain and Cognitive Science	- - -	- - -	- - -	- - -	- - -
32	Linderman, Scott	b	Harvard University	School of Engineering and Applied Sciences	- - -	- - -	- - -	- - -	- - -
33	Magid, Rachel	j	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
34	Mahadevan, L.	c	Harvard University	School of Engineering and Applied Sciences, Physics, Organismic and Evolutionary Biology	- - -	- - -	- - -	- - -	- - -
35	Marantan, Andrew	b	Harvard University	Physics	- - -	- - -	- - -	- - -	- - -
36	Meroz, Yasmine	f	Harvard University	School of Engineering and Applied Science	- - -	- - -	- - -	- - -	- - -
37	Meyers, Ethan	f	Massachusetts Institute of Technology	BCS	- - -	- - -	- - -	- - -	- - -
38	Mutch, Jim	b	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
39	Nakayama, Ken*	c	Harvard University		- - -	- - -	- - -	- - -	- - -
40	Newman, Jonathan P.	f	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
41	Nickel, Maximilian	f	Massachusetts Institute of Technology	McGovern Institute for Brain Research	- - -	- - -	- - -	- - -	- - -
42	Payer, Kristofor	j	Massachusetts Institute of Technology	MTL	- - -	- - -	- - -	- - -	- - -
43	Peleg, Orit	f	Harvard University	School of Engineering and Applied Sciences	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
44	Penagos, Hector	f	Massachusetts Institute of Technology	Picower Institute for Learning and Memory	- - -	- - -	- - -	- - -	- - -
45	Peterson, Matthew Ferris	f	Massachusetts Institute of Technology	McGovern Institute for Brain Research	- - -	- - -	- - -	- - -	- - -
46	Poggio, Tomaso A.	c	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
47	Powell, Lindsey*	f	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
48	Robertson, Caroline*	f	Massachusetts Institute of Technology	McGovern Institute for Brain Research	- - -	- - -	- - -	- - -	- - -
49	Rosasco, Lorenzo	d/e	Massachusetts Institute of Technology	McGovern Institute for Brain Research	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
50	Sadagopan, Srivatsun	f	The Rockefeller University	Laboratory of Neural Systems	- - -	- - -	- - -	- - -	- - -
51	Sassanfar, Mandana	c	Massachusetts Institute of Technology	Biology	- - -	- - -	- - -	- - -	- - -
52	Saxe, Rebecca R.*	c	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
53	Schmidt, Ellyn	j	Harvard University	Psychology	- - -	- - -	- - -	- - -	- - -
54	Schulz, Laura*	c	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
55	Scott, Kimberly	b	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
56	Siegel, Max	b	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
57	Singer, Jedediah	f	Boston Children's Hospital	Ophthalmology	- - -	- - -	- - -	- - -	- - -
58	Skerry, Amy	b	Harvard University	Department of Psychology	- - -	- - -	- - -	- - -	- - -
59	Sompolinsky, Haim	c	Harvard University	Center for Brain Sciences	- - -	- - -	- - -	- - -	- - -
60	Spelke, Elizabeth	c	Harvard University	Psychology	- - -	- - -	- - -	- - -	- - -
61	Spokes, Annie*	b	Harvard University	Psychology	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
62	Sullivan, Kathleen	j	Massachusetts Institute of Technology	McGovern Institute for Brain Research	- - -	- - -	- - -	- - -	- - -
63	Tacchetti, Andrea	b	Massachusetts Institute of Technology	Electrical Engineering & Computer Science	- - -	- - -	- - -	- - -	- - -
64	Tang, Hanlin	b	Harvard University	Biophysics	- - -	- - -	- - -	- - -	- - -
65	Tenenbaum, Joshua	c	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
66	Ullman, Shimon	c	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
67	Ullman, Tomer	b	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
68	Valiant, Leslie*	c	Harvard University	School of Engineering and Applied Sciences, Physics, Organismic and Evolutionary Biology	- - -	- - -	- - -	- - -	- - -
69	Vaziri Pashkam, Maryam	f	Harvard University	Psychology	- - -	- - -	- - -	- - -	- - -
70	Wilson, Matt	c	Massachusetts Institute of Technology	Brain and Cognitive Sciences	- - -	- - -	- - -	- - -	- - -
71	Winston, Patrick Henry	c	Massachusetts Institute of Technology	Electrical Engineering & Computer Science	- - -	- - -	- - -	- - -	- - -
72	Wong, Neelum	j	Massachusetts Institute of Technology	McGovern Institute for Brain Research	- - -	- - -	- - -	- - -	- - -

	Participant Name	Category	Institutional Affiliation	Department (if applicable)	Gender	Disability Status	Ethnicity	Race	Citizenship
73	Yuille, Alan*	c	UCLA	Statistics	- - -	- - -	- - -	- - -	- - -

6. Center Partners

	Organization Name	Organization Type*	Address	Contact Name	Type of Partner**	160 hours or more? (Y/N)
1	Massachusetts Institute of Technology (MIT)	Academic	77 Massachusetts Ave. Cambridge, MA 02139	Tomaso Poggio	Research Partner	Y
2	Harvard University	Academic	Massachusetts Hall Cambridge, MA 02138	L. Mahadevan	Research Partner	Y
3	Cornell University	Academic	Day Hall Lobby Ithaca, NY 14853	Haym Hirsh	Research Partner	Y
4	Rockefeller University	Academic	1230 York Avenue New York, NY 10065	Winrich Freiwald	Research Partner	Y
5	Stanford University	Academic	450 Serra Mall Stanford, CA 94305	Noah Goodman	Research Partner	Y
6	University of California, Los Angeles (UCLA)	Academic	Los Angeles, CA 90095	Alan Yuille	Research Partner	Y
7	Allen Institute for Brain Science	Research Institution	551 N 34th St. #200 Seattle, WA 98103	Christof Koch	Research Partner	Y
8	Wellesley College	Academic	106 Central St. Wellesley, MA 02481	Ellen Hildreth	Education, Diversity	Y
9	Children’s Hospital Boston	Research Institution	300 Longwood Ave. Boston, MA 02115	Gabriel Kreiman	Research Partner	Y
10	University of Puerto Rico – Río Piedras (UPRRP)	Academic	PO Box 22360 San Juan, Puerto Rico 00931-360	Irving E. Vega	Diversity	N
11	Universidad Central del Caribe (UCC)	Academic	2U6 Ave. Laurel Lomas Verdes Bayamon, Puerto Rico 00956	Maria Bykhovskaia	Diversity	N

	Organization Name	Organization Type*	Address	Contact Name	Type of Partner**	160 hours or more? (Y/N)
12	The City University of New York: Hunter College	Academic	695 Park Ave. NY, NY 10065	Susan Epstein	Diversity	N
13	The City University of New York: Queen's College	Academic	65-30 Kissena Blvd. Flushing, NY 11367	Joshua C. Brumberg	Diversity	N
14	Howard University	Academic	2400 Sixth St. NW Washington, DC 20059	Mohamed F. Chouikha	Diversity	N
15	Institute for Infocomm Research, A*STAR	Research Institution	1 Fusionopolis Way #21-01 Connexis (South Tower) Singapore 138632	Lim Joo-Hwee	International	N
16	Weizmann Institute of Science	Academic	234 Herzl Street Rehovot 7610001 Israel	Shimon Ullman	International	Y
17	University of Genoa	Academic	Via Balbi, 5, 16126 Genova, Italy	Alessandro Verri	International	Y
18	City University, Hong Kong	Academic	Tat Chee Avenue Kowloon, Hong Kong	Stephen Smale	International	N
19	Hebrew University of Jerusalem	Academic	Har ha-Tsofim, Jerusalem, Israel	Haim Sompolinsky, Amnon Shashua	International	N
20	Max Planck Institute for Biological Cybernetics, Tübingen	Research Institution	Max Planck Institute for Biological Cybernetics P.O. Box: 21 69 72012 Tübingen, Germany	Heinrich Buelthoff	International	N

	Organization Name	Organization Type*	Address	Contact Name	Type of Partner**	160 hours or more? (Y/N)
21	Istituto Italiano di Tecnologia (IIT)	Research Institution	Via Morego, 30 16163 Genova, Italy	Giorgio Metta, Lorenzo Rosasco	International	N
22	National Center for Biological Sciences, Bangalore, India	Research Institution	GKVK, Bellary Rd. Bangalore 560065, India	K. Vijay Raghavan	International	N
23	DeepMind Technologies Ltd	Company	5 New Street Square, London EC4A 3TW, UK	Demis Hassabis	Industrial	N
24	GE	Company	1 Research Circle Niskayuna, NY 12309	Mark Grabb	Industrial	N
25	Google	Company	1600 Amphitheatre Parkway Mountain View, CA 94043	Peter Norvig,	Industrial	N
26	IBM Research	Corporate Research	1 New Orchard Rd. Armonk, New York 10504-1722	Zachary Lemnios	Industrial	N
27	Microsoft Research	Corporate Research	One Memorial Drive Cambridge, MA 02142	Andrew Blake	Industrial	N
28	Mobileye	Company	Har Hotzvim, 13 Hartom Street, PO Box 45157 Jerusalem 9777513, Israel	Amnon Shashua	Industrial	N
29	Orcam	Company	Jerusalem, Israel	Amnon Shashua	Industrial	N
30	Rethink Robotics, Inc.	Company	27 Wormwood St. Boston, MA 02210	Rodney Brooks	Industrial	N
31	Schlumberger	Company	One Hampshire St. B253 Cambridge, MA 02139-1578	Tarek Habashy	Industrial	N

	Organization Name	Organization Type*	Address	Contact Name	Type of Partner**	160 hours or more? (Y/N)
32	Siemens Corp.	Company	755 College Road East Princeton, N.J. 08540	Dorin Comaniciu	Industrial	N

7. NSF Summary Table

1	Number of participating institutions (all academic)	16
2	Number of institutional partners (non-academic)	16
3	Total leveraged support for the current year	$100,000 (Schlumberger)
4	Number of participants	61 (excluding affiliates) 73 (including affiliates)

8. Center Media Publicity

Materials included in Appendix D:

Artificial-intelligence research revives its old ambitions
MIT News Office
September 9, 2013
http://newsoffice.mit.edu/2013/center-for-brains-minds-and-machines-0909

MIT center receives $25 million to unravel the mysteries of human intelligence
Boston.com
September 9, 2013
http://www.boston.com/news/science/blogs/science-in-mind/2013/09/09/mit-center-receives-million-unravel-the-mysteries-human-intelligence/SKNU4umjtPy5pN7nFnYcXL/blog.html

NSF pumps $25M into MIT artificial intelligence center
Boston Business Journal
September 10, 2013
http://www.bizjournals.com/boston/blog/techflash/2013/09/nsf-pumps-25m-into-mit-artificial.html

$25 million NSF grant to team including 6 CBS faculty
Center for Brain Science, News Blog, Harvard U.
September 10, 2013
http://cbs.fas.harvard.edu/resources/news/25-million-nsf-grant-team-including-6-cbs-faculty

Making artificial intelligence more human
Boston Globe
October 7, 2013
http://www.bostonglobe.com/news/science/2013/10/06/mit-artificial-intelligence-center-backed-federal-grant-learning-from-infant-brain-research/MdPnWBnGv7KA1N3CVssKEO/story.html

CBMM co-hosts workshop on data representation
MIT News Office
November 14, 2013
http://newsoffice.mit.edu/2013/cbmm-co-hosts-workshop-on-data-representation

Brainlike Computers, Learning From Experience
The New York Times
December 28, 2013
http://www.nytimes.com/2013/12/29/science/brainlike-computers-learning-from-experience.html

Obama Administration Proposes Doubling Support for The BRAIN Initiative
White House Press Release
March 4, 2014
http://www.whitehouse.gov/sites/default/files/microsites/ostp/FY%202015%20BRAIN.pdf

The brain is a Swiss Army knife: Nancy Kanwisher at TED2014
TEDBlog
March 19, 2014
http://blog.ted.com/2014/03/19/the-brain-is-a-swiss-army-knife-nancy-kanwisher-at-ted2014/

How the brain pays attention
MIT News Office
April 10, 2014
http://newsoffice.mit.edu/2014/how-brain-pays-attention

Mens et Apparata. Creating artificial intelligence turns out to be far more challenging than anyone expected. But the new Center for Brains, Minds, and Machines is ready to try again. This time, computer scientists, biologists, and neuroscientists will be tackling the problem together.
MIT Technology Review
Apr 22, 2014
http://www.technologyreview.com/article/526376/mens-et-apparata/

Illuminating neuron activity in 3-D
MIT News Office
May 18, 2014
https://newsoffice.mit.edu/2014/illuminating-neuron-activity-3-d-0518

IX. INDIRECT/OTHER IMPACTS

No indirect/other impacts to report.

Appendix D.

Please find included the Center Media Publicity documents:

Artificial-intelligence research revives its old ambitions, MIT News Office, September 9, 2013

MIT center receives $25 million to unravel the mysteries of human intelligence, Boston.com, September 9, 2013

NSF pumps $25M into MIT artificial intelligence center, Boston Business Journal, September 10, 2013

$25 million NSF grant to team including 6 CBS faculty, Center for Brain Science, News Blog, Harvard U., September 10, 2013

Making artificial intelligence more human, Boston Globe, October 7, 2013

CBMM co-hosts workshop on data representation, MIT News Office, November 14, 2013

Brainlike Computers, Learning From Experience, The New York Times, December 28, 2013

Obama Administration Proposes Doubling Support for The BRAIN Initiative, White House Press Release, March 4, 2014

The brain is a Swiss Army knife: Nancy Kanwisher at TED2014, **TED**Blog, March 19, 2014

How the brain pays attention, MIT News Office, April 10, 2014

Mens et Apparata. MIT Technology Review, Apr 22, 2014

Illuminating neuron activity in 3-D , MIT News Office, May 18, 2014

MIT News
ON CAMPUS AND AROUND THE WORLD

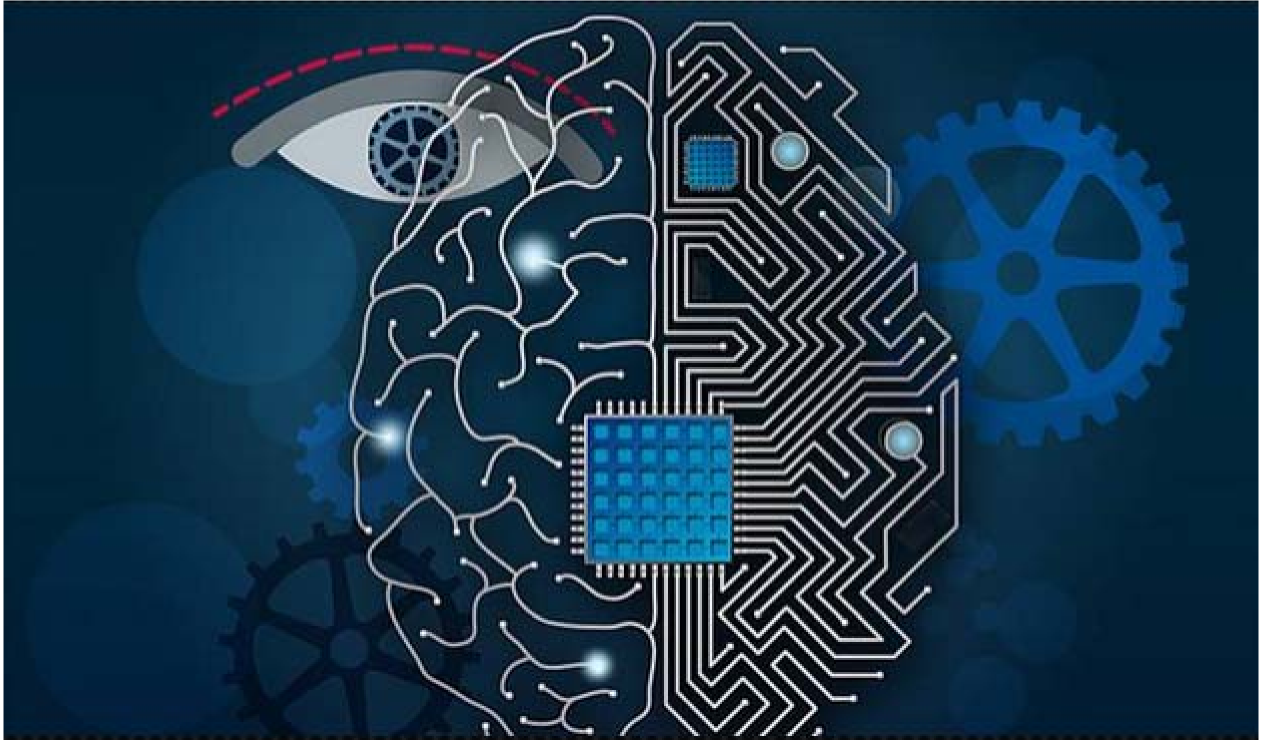

Artificial-intelligence research revives its old ambitions

A new interdisciplinary research center at MIT, funded by the National Science Foundation, aims at nothing less than unraveling the mystery of intelligence.

Larry Hardesty, MIT News Office

September 9, 2013

The birth of artificial-intelligence research as an autonomous discipline is generally thought to have been the month long Dartmouth Summer Research Project on Artificial Intelligence in 1956, which convened 10 leading electrical engineers — including MIT’s Marvin Minsky and Claude Shannon — to discuss “how to make machines use language” and “form abstractions and concepts.” A decade later, impressed by rapid advances in the design of digital computers, Minsky was emboldened to declare that “within a generation ... the problem of creating ‘artificial intelligence’ will substantially be solved.”

The problem, of course, turned out to be much more difficult than AI’s pioneers had imagined. In recent years, by exploiting machine learning — in which computers learn to perform tasks from sets of training examples — artificial-intelligence researchers have built special-purpose systems that can do things like interpret spoken language or play Jeopardy with great success.

But according to Tomaso Poggio, the Eugene McDermott Professor of Brain Sciences and Human Behavior at MIT, “These recent achievements have, ironically, underscored the limitations of computer science and artificial intelligence. We do not yet understand how the brain gives rise to intelligence, nor do we know how to build machines that are as broadly intelligent as we are.”

Poggio thinks that AI research needs to revive its early ambitions. “It’s time to try again,” he says. “We know much more than we did before about biological brains and how they produce intelligent behavior. We’re now at the point where we can start applying that understanding from neuroscience, cognitive science and computer science to the design of intelligent machines.”

The National Science Foundation (NSF) appears to agree: Today, it announced that one of three new research centers funded through its Science and Technology Centers Integrative Partnerships program will be the Center for Brains, Minds and Machines (CBMM), based at MIT and headed by Poggio. Like all the centers funded through the program, CBMM will initially receive $25 million over five years.

Homegrown initiative

CBMM grew out of the MIT Intelligence Initiative, an interdisciplinary program aimed at understanding how intelligence arises in the human brain and how it could be replicated in machines.

“[MIT President] Rafael Reif, when he was provost, came to speak to the faculty and challenged us to come up with new visions, new ideas,” Poggio says. He and MIT’s Joshua Tenenbaum, also a professor in the Department of Brain and Cognitive Sciences (BCS) and a principal investigator in the Computer Science and Artificial Intelligence Laboratory, responded by proposing a program that would integrate research at BCS and the Department of Electrical Engineering and Computer Science. “With a system as complicated as the brain, there is a point where you need to get people to work together across different disciplines and techniques,” Poggio says. Funded by MIT’s School of Science, the initiative was formally launched, in 2011, at a symposium during MIT’s 150th anniversary.

Headquartered at MIT, CBMM will be, like all the NSF centers, a multi-institution collaboration. Of the 20 faculty members currently affiliated with the center, 10 are from MIT, five are from Harvard University, and the rest are from Cornell University, Rockefeller University, the University of California at Los Angeles, Stanford University and the Allen Institute for Brain Science. The center’s international partners are the Italian Institute of Technology; the Max Planck Institute in Germany; City University of Hong Kong; the National Centre for Biological Sciences in India; and Israel’s Weizmann Institute and Hebrew University. Its industrial partners are Google, Microsoft, IBM, Mobileye, Orcam, Boston Dynamics, Willow Garage, DeepMind and Rethink Robotics. Also affiliated with center are Howard University; Hunter College; Universidad Central del Caribe, Puerto Rico; the University of Puerto Rico, Río Piedras; and Wellesley College.

CBMM aims to foster collaboration not just between institutions but also across disciplinary boundaries. Graduate students and postdocs funded through the center will have joint advisors, preferably drawn from different research areas.

Research themes

The center’s four main research themes are also intrinsically interdisciplinary. They are the integration of intelligence, including vision, language and motor skills; circuits for intelligence, which will span research in neurobiology and electrical engineering; the development of intelligence in children; and social intelligence. Poggio will also lead the development of a theoretical platform intended to undergird the work in all four areas.

“Those four thrusts really do fit together, in the sense that they cover what we think are the biggest challenges facing us when we try to develop a computational understanding of what intelligence is all about,” says Patrick Winston, the Ford Foundation Professor of Engineering at MIT and research coordinator for CBMM.

For instance, he explains, in human cognition, vision, language and motor skills are inextricably linked, even though they’ve been treated as separate problems in most recent AI research. One of Winston’s favorite examples is that of image labeling: A human subject will identify an image of a man holding a glass to his lips as that of a man drinking. If the man is holding the glass a few inches further forward, it’s an instance of a different activity — toasting. But a human will also identify an image of a cat turning its head up to catch a few drops of water from a faucet as an instance of drinking. “You have to be thinking about what you see there as a story,” Winston says. “They get the same label because it’s the same story, not because it looks the same.”

Similarly, Winston explains, development is its own research thrust because intelligence is fundamentally shaped through interaction with the environment. There’s evidence, Winston says, that mammals that receive inadequate visual stimulation in the first few weeks of life never develop functional eyesight, even though their eyes are otherwise unimpaired. “You need to stimulate the neural mechanisms in order for them to assemble themselves into a functioning system,” Winston says. “We think that that’s true generally, of our entire spectrum of capabilities. You need to have language, you need to see things, you need to have language and vision work together from the beginning to ensure that the parts develop properly to form a working whole.”

boston.com

MIT center receives $25 million to unravel the mysteries of human intelligence

By Carolyn Y. Johnson
Globe Staff | 09.09.13 | 5:19 PM

A new center headquartered at the Massachusetts Institute of Technology will focus on bringing researchers from separate fields together to try and crack one of the biggest questions facing science today: what is intelligence, and how can we engineer it?

The Center for Brains, Minds, and Machines, supported by a $25 million grant from the National Science Foundation, will include faculty from MIT, Harvard University, and a handful of other universities, as well as a slate of industrial and international partners. By tapping a broad range of expertise, including scholars who study how a baby's mind develops and others trying to understand how the brain makes sense of social situations, the researchers hope to take a definitive step forward over the next five years toward the long-held goal of understanding intelligence and building computers capable of thinking like people.

"It is a short time in terms of the size of the problem; I don't expect at all we'll solve the problem of intelligence and how the brain works and how the mind works," said Tomaso Poggio, the director of the center and an MIT professor of brain sciences and human behavior. "But we hope to make significant progress, and we hope to shape the research, not only here in Cambridge, but also around the world in how to approach this problem."

The project was proposed well before President Obama announced his BRAIN initiative to map the circuitry of the brain, but Poggio said the center represents a key piece of the National Science Foundation's work related to the initiative.

Boston Business Journal

NSF pumps $25M into MIT artificial intelligence center

Sep 10, 2013, 9:53am EDT

The National Science Foundation will pump $25 million over five years into the Center for Brains, Minds, and Machines (CBMM), which was recently launched out of a program at MIT.

The CBMM, which grew out of the MIT Intelligence Initiative, a program focused on how the human brain can be replicated by machines, was one of three new research centers funded through NSF's Science and Technology Centers Integrative Partnership program. The center, which will be led by Tomaso Poggio, the Eugene McDermott Professor of Brain Sciences and Human Behavior at MIT, will be a multi-institution collaboration.

A total of 20 faculty members are involved with the center. Ten of the faculty members are from MIT, five are from Harvard University and the remaining members are from Cornell University, Rockefeller University, the University of California at Los Angeles, Stanford University and the Allen Institute for Brain Science, will focus on four themes: intelligence, including vision, language and motor skills; circuits for intelligence, which will focus on neurobiology and electrical engineering; the development of intelligence in children; and social intelligence.

“Those four thrusts really do fit together, in the sense that they cover what we think are the biggest challenges facing us when we try to develop a computational understanding of what intelligence is all about,” says Patrick Winston, research coordinator for the CBMM.

CENTER FOR
BRAIN SCIENCE

▾ COMMUNITY ▾ SCIENCE ▾ RESOURCES ABOUT CONTACT US

$25 MILLION NSF GRANT TO TEAM INCLUDING 6 CBS FACULTY

Published Date: September 10, 2013

The National Science Foundation (NSF) announced a $25 million award over 5 years to a new Science and Technology Center (STC) called the Center for Brains, Minds and Machines. This STC will be a multi-institution collaboration, with a partnership between MIT and Harvard at its heart.

The startling triumphs of Deep Blue in chess and Watson in Jeopardy, coupled with news of driverless cars and drones making automated landings on aircraft carriers, not to mention the commercial splash made by Siri, make it clear that advances in machine learning have been profound. Nevertheless, it is clear that these specialized systems are not intelligent in the manner of humans or even mice. The new STC will combine neuroscience and cognitive science with computer science and engineering to renew an attack on this problem of intelligence in brains and machines.

The new STC will be headquartered at MIT and led by Tomaso Poggio; Mahadevan will be the associate director from Harvard, where he will establish partnerships with industry and play a prominent role in education. Ten of the Center faculty are from MIT, 6 are from CBS, and one each if from Cornell, Rockefeller, UCLA, Stanford, and the Allen Institute for Brain Science. Also affiliated with center are Howard University; Hunter College; Universidad Central del Caribe, Puerto Rico; the University of Puerto Rico, Río Piedras; and Wellesley College.

CBS faculty in the Center are Mahadevan, Gabriel Kreiman, Ken Nakayama, Haim Sompolinsky, Liz Spelke, and Les Valiant. CBS will oversee the administration of the Harvard part of this award.

The Center will have five main research themes: circuits for intelligence; the development of intelligence in children; social intelligence; the integration of visual, motor, language, and social intelligence; and theoretical aspects of intelligence. The Center aims to foster collaboration not just between institutions but also across disciplinary boundaries. Graduate students and postdocs funded through the center will have joint advisors, generally drawn from different research areas.

The Center will have industrial partnerships with Google, Microsoft, IBM, Mobileye, Orcam, Boston Dynamics, Willow Garage, Deep Minds, and Rethink Robotics. It also will have international partners: the Italian Institute of Technology in Genoa; the Max Planck Institute in Tübingen; City University of Hong Kong; the National Centre for Biological Sciences in Bangalore; the Weizmann Institute in Rehovot; and Hebrew University in Jerusalem.

The Boston Globe

Making artificial intelligence more human

MIT center, backed by $25m federal grant, learning from infant brain research

By Carolyn Y. Johnson | GLOBE STAFF OCTOBER 07, 2013

CAMBRIDGE — The bold quest to build intelligent machines has, after more than half a century, brought us to this point: Scientists can build a “Jeopardy!” champion, but a child can handily outperform a computer when it comes to deciphering social situations, learning, or pretty much any activity outside the machine’s narrow band of expertise.

To change that, a group of leading infant researchers, neurobiologists, computer scientists, and robotics and software companies are joining forces in a major effort to finally achieve and even expand the grandiose ambitions of artificial intelligence, supported by a $25 million grant from the National Science Foundation.

At a new center based at the Massachusetts Institute of Technology, researchers will seek to craft intelligence that includes not just knowledge but also an infant’s ability to intuit basic concepts of psychology or physics. Answering cleverly posed trivia questions is impressive, but perhaps more so is the ability to make sense of observations — answering seemingly simple inquiries such as “what is the object closest to the window?” or “what is the woman looking at?”

“I think this is the greatest problem in science and technology, greater than the origin of the universe or the origin of life or the nature of matter, partly because it’s a problem about who we are,” said Tomaso Poggio, the director of the new Center for Brains, Minds, and Machines, and a professor at MIT. “It’s a problem about the very tool you use to solve all other problems: your brain.”

The center will draw together 20 faculty from MIT, Harvard University, and other major universities, as well as business partners that includes Google, Microsoft, and local robotics companies Boston Dynamics and Rethink Robotics.

Recognizing the importance of advances in technology and science, President Obama announced a major national effort this year to map the activity of the brain. The ambition to build intelligent machines, scientists argue, is at a pivotal moment for similar reasons, fueled not just by advances in computers and robotics but by progress in the ability to probe the brain circuits that underlie specific behaviors and by greater understanding of how intelligence develops in the infant brain.

The past few decades have seen an enormous flood of information about the infant and child mind, overturning ideas about what babies know and how they learn. For example, developmental psychologists have delineated precisely the way children begin to grasp how the rules of gravity apply to objects over the first year of life.

“In the early days, understanding intelligence had to do with reasoning and problem-solving,” said Patrick Winston, a professor of engineering at MIT and the center’s research coordinator. “On the science side, we’ve moved away from mathematical reasoning to common sense.”

For Winston, what makes human intelligence most stand apart from machines — and from the rest of the animal world — is our ability to tell and comprehend stories. Finding sense in a fairy tale may seem trivial, but Winston has been working for years on something not far from that: a computer program called the Genesis project that can be fed a block of text and do something approximating what we do when we read a story. Given a quick summary of Shakespeare’s “Macbeth” or a story about a conflict between two countries, the program can try to draw causal links and figure out why things happened and what it all means. It can detect concepts such as revenge and assess people’s character.

At his office in the Stata Center, Winston showed the program at work. Boxes popped up on the screen and rearranged themselves as the program essentially diagrammed sentences and paragraphs. At this point, Genesis can easily be broken by a story it does not comprehend, but the goal is to work up to a library of more than 100 stories it can correctly parse within a year. In demo mode, it was sensitive to slight alterations in meaning, detecting that a politician “forcing” a country to move toward democracy was different than “asking” for the same outcome.

The machine has been taught concepts, such as “the enemy of my enemy is my friend.” These lessons can seem almost humorous when written in the explanatory language for the computer: “XX is an entity. YY is an entity. XX harms YY. . . . XX’s harming YY leads to YY’s harming XX.” But in some ways, that lesson is not all that different from how people acquire many of those concepts. People, too, have to be told what a Pyrrhic victory is.

Across the street, at a building teeming with neurobiologists and cognitive scientists, Joshua Tenenbaum, another member of the center, is taking a slightly different tack in the effort to understand intelligence: trying to build a child’s mind.

“Let’s try to reverse-engineer the early stages of cognition. What do young babies know?” Tenenbaum said. “Even young babies, 3 or 4 years old, are more intelligent than any machine has ever been. Let’s build a road map of cognitive development over the first three years of life, but let’s build it in engineering terms — the same terms I would use to build a self-driving car.”

The reason such an effort now seems plausible, Tenenbaum said, is that engineers and developmental psychologists are finally using the same powerful type of math, whether they are working to design programs or describe the developing mind.

Progress has already been made in artificial intelligence — tremendously, in some task-based areas. Poggio, the center director, said that when he started in the field three decades ago, he worked on developing a camera system that could detect pedestrians. A type of computer vision that made about 10 mistakes a second was considered a major feat. Now, he said, computer vision systems developed for driving applications make mistakes once every 30,000 hours of driving.

It’s even possible to see computers overtaking some human abilities. Researcher Joel Leibo has been working on a vision system that recognizes faces at a glance, even when given the difficult challenge of matching people’s faces shown from different angles. The system is getting more sophisticated; it did only slightly worse than a research assistant who tried the same face-matching task.

“There are systems that are doing certain things that are really difficult to do — and difficult for humans to be able to do,” Poggio said. “But none of these systems are really intelligent; you cannot have a conversation with these systems.”

These researchers want to build a different kind of intelligence. Imagine a cafeteria at MIT at lunchtime. People are doing what people do daily: talking, eating, arguing, sipping a drink through a straw. Researchers want to build intelligence smart enough to take in the scene and describe, in words, exactly what’s happening.

MIT News
ON CAMPUS AND AROUND THE WORLD

CBMM co-hosts workshop on data representation

Center for Brains, Minds and Machines
November 14, 2013

The new Center for Brains, Minds and Machines, a National Science Foundation-funded center on the interdisciplinary study of intelligence, is starting to gain momentum. Last month the McGovern Institute for Brain Research at MIT hosted an introductory seminar and reception for the new Center. Investigators from MIT and Harvard gave presentations of their research plans to the wider scientific community.

In addition to collaborative research, the CBMM plans to develop a community of researchers via programs such as an intensive summer school and technical workshops that will train the next generation of scientists and engineers in an emerging new field -- the Science and Engineering of Intelligence. This new field will catalyze cross-fertilization between computer science, math and statistics, robotics, neuroscience, and cognitive science.

The first such workshop will be held in collaboration with the Center's international partner, the Italian Institute of Technology (IIT). The workshop is titled "Learning Data Representation: Hierarchies and Invariance" and will take place at the McGovern Institute for Brain Research at MIT from Nov. 22-24.

The goal of the meeting is to investigate advances and challenges in learning "good representations" from data, in particular representations that can reduce the complexity of later supervised learning stages. The meeting will gather experts in the field to discuss current and future challenges for the theory and applications of learning representations. It is hoped that the meeting may mark the beginning of a new phase in machine learning where it is possible to develop algorithms capable of learning like humans. In this novel framework, instead of very large sets of labeled data (big data), learning needs only very small sets of labeled examples — a new small data paradigm.

The New York Times

Brainlike Computers, Learning From Experience

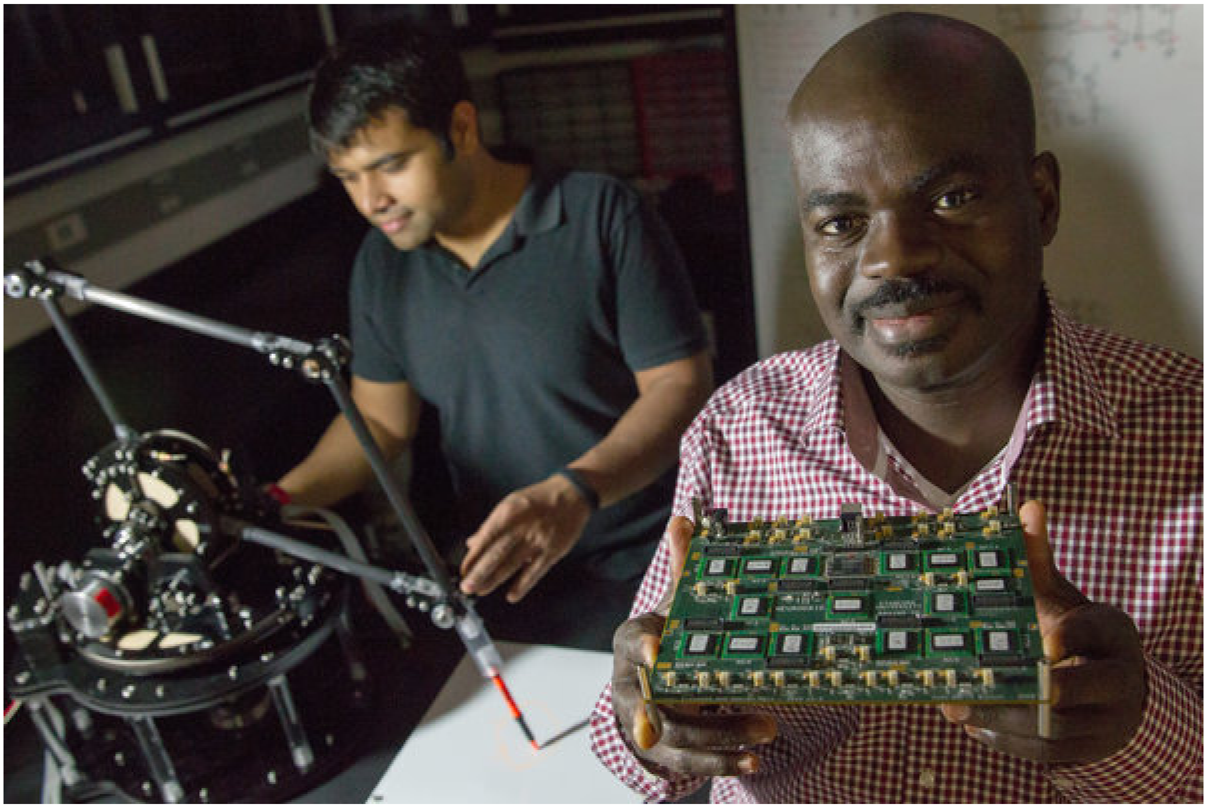

Erin Lubin/The New York Times

Kwabena Boahen holding a biologically inspired processor attached to a robotic arm in a laboratory at Stanford University.

By JOHN MARKOFF

Published: December 28, 2013

PALO ALTO, Calif. — Computers have entered the age when they are able to learn from their own mistakes, a development that is about to turn the digital world on its head.

The first commercial version of the new kind of computer chip is scheduled to be released in 2014. Not only can it automate tasks that now require painstaking programming — for example, moving a robot’s arm smoothly and efficiently — but it can also sidestep and even tolerate errors, potentially making the term “computer crash” obsolete.

The new computing approach, already in use by some large technology companies, is based on the biological nervous system, specifically on how neurons react to stimuli and connect with other neurons to interpret information. It allows computers to absorb new information while carrying out a task, and adjust what they do based on the changing signals.

In coming years, the approach will make possible a new generation of artificial intelligence systems that will perform some functions that humans do with ease: see, speak, listen, navigate, manipulate and control. That can hold enormous consequences for tasks like facial and speech recognition, navigation and planning, which are still in elementary stages and rely heavily on human programming.

Designers say the computing style can clear the way for robots that can safely walk and drive in the physical world, though a thinking or conscious computer, a staple of science fiction, is still far off on the digital horizon.

“We’re moving from engineering computing systems to something that has many of the characteristics of biological computing,” said Larry Smarr, an astrophysicist who directs the California Institute for Telecommunications and Information Technology, one of many research centers devoted to developing these new kinds of computer circuits.

Conventional computers are limited by what they have been programmed to do. Computer vision systems, for example, only “recognize” objects that can be identified by the statistics-oriented algorithms programmed into them. An algorithm is like a recipe, a set of step-by-step instructions to perform a calculation.

But last year, Google researchers were able to get a machine-learning algorithm, known as a neural network, to perform an identification task without supervision. The network scanned a database of 10 million images, and in doing so trained itself to recognize cats.

In June, the company said it had used those neural network techniques to develop a new search service to help customers find specific photos more accurately.

The new approach, used in both hardware and software, is being driven by the explosion of scientific knowledge about the brain. Kwabena Boahen, a computer scientist who leads Stanford’s Brains in Silicon research program, said that is also its limitation, as scientists are far from fully understanding how brains function.

“We have no clue,” he said. “I’m an engineer, and I build things. There are these highfalutin theories, but give me one that will let me build something.”

Until now, the design of computers was dictated by ideas originated by the mathematician John von Neumann about 65 years ago. Microprocessors perform operations at lightning speed, following instructions programmed using long strings of 1s and 0s. They generally store that information separately in what is known, colloquially, as memory, either in the processor itself, in adjacent storage chips or in higher capacity magnetic disk drives.

The data — for instance, temperatures for a climate model or letters for word processing — are shuttled in and out of the processor's short-term memory while the computer carries out the programmed action. The result is then moved to its main memory.

The new processors consist of electronic components that can be connected by wires that mimic biological synapses. Because they are based on large groups of neuron-like elements, they are known as neuromorphic processors, a term credited to the California Institute of Technology physicist Carver Mead, who pioneered the concept in the late 1980s.

They are not "programmed." Rather the connections between the circuits are "weighted" according to correlations in data that the processor has already "learned." Those weights are then altered as data flows in to the chip, causing them to change their values and to "spike." That generates a signal that travels to other components and, in reaction, changes the neural network, in essence programming the next actions much the same way that information alters human thoughts and actions.

"Instead of bringing data to computation as we do today, we can now bring computation to data," said Dharmendra Modha, an I.B.M. computer scientist who leads the company's cognitive computing research effort. "Sensors become the computer, and it opens up a new way to use computer chips that can be everywhere."

The new computers, which are still based on silicon chips, will not replace today's computers, but will augment them, at least for now. Many computer designers see them as coprocessors, meaning they can work in tandem with other circuits that can be embedded in smartphones and in the giant centralized computers that make up the cloud. Modern computers already consist of a variety of coprocessors that perform specialized tasks, like producing graphics on your cellphone and converting visual, audio and other data for your laptop.

One great advantage of the new approach is its ability to tolerate glitches. Traditional computers are precise, but they cannot work around the failure of even a single transistor. With the biological designs, the algorithms are ever changing, allowing the system to continuously adapt and work around failures to complete tasks.

Traditional computers are also remarkably energy inefficient, especially when compared to actual brains, which the new neurons are built to mimic.

I.B.M. announced last year that it had built a supercomputer simulation of the brain that encompassed roughly 10 billion neurons — more than 10 percent of a human brain. It ran about 1,500 times more slowly than an actual brain. Further, it required several megawatts of power, compared with just 20 watts of power used by the biological brain.

Running the program, known as Compass, which attempts to simulate a brain, at the speed of a human brain would require a flow of electricity in a conventional computer that is equivalent to what is needed to power both San Francisco and New York, Dr. Modha said.

I.B.M. and Qualcomm, as well as the Stanford research team, have already designed neuromorphic processors, and Qualcomm has said that it is coming out in 2014 with a commercial version, which is expected to be used largely for further development. Moreover, many universities are now focused on this new style of computing. This fall the National Science Foundation financed the Center for Brains, Minds and Machines, a new research center based at the Massachusetts Institute of Technology, with Harvard and Cornell.

The largest class on campus this fall at Stanford was a graduate level machine-learning course covering both statistical and biological approaches, taught by the computer scientist Andrew Ng. More than 760 students enrolled. “That reflects the zeitgeist,” said Terry Sejnowski, a computational neuroscientist at the Salk Institute, who pioneered early biologically inspired algorithms. “Everyone knows there is something big happening, and they’re trying find out what it is.”

Obama Administration Proposes Doubling Support for the BRAIN Initiative

EXECUTIVE OFFICE OF THE PRESIDENT OF THE UNITED STATES
OFFICE OF SCIENCE AND TECHNOLOGY POLICY

> “So there is this enormous mystery waiting to be unlocked, and the BRAIN Initiative will change that by giving scientists the tools they need to get a dynamic picture of the brain in action and better understand how we think and how we learn and how we remember. And that knowledge could be -- will be -- transformative.”
>
> \- President Barack Obama
> April 2013

On April 2, 2013, President Obama launched the Brain Research though Advancing Innovative Neurotechnologies (BRAIN) Initiative, a Grand Challenge designed to revolutionize our understanding of the human brain. Under this initiative, Federal agencies such as the Defense Advanced Research Projects Agency (DARPA), the National Institutes of Health (NIH), the National Science Foundation (NSF), and the Food and Drug Administration (FDA) are supporting the development and application of innovative new technologies that can create a dynamic understanding of brain function and its relationship to behavior. These scientific and technological advances could also lead to improvements in our ability to diagnose, treat, and even prevent diseases of the brain.

The President’s 2015 Budget proposes to double the Federal investment in The BRAIN Initiative from about $100 million in FY 2014 to approximately $200 million in FY 2015. Proposed investments by the NIH, DARPA, and NSF are described below.

Given the audacious goals of the initiative, the President has called for this to be an “all hands on deck” effort involving not only the Federal Government but also companies, health systems, patient advocacy organizations, philanthropists, state governments, research universities, private research institutes, and scientific societies. Later this year, the White House will hold an event to feature the role of these organizations in achieving the President’s bold vision.

National Institutes of Health: In FY 2015, NIH plans to expand its commitment to the success of The BRAIN Initiative, with an estimated $100 million in funding from the agency. NIH will develop and apply new tools to map the circuits of the brain, measure the dynamic patterns of activity within those circuits, and understand how they create unique cognitive and behavioral capabilities. Ultimately, this fundamental knowledge is expected to revolutionize our understanding of complex brain functions and their links to behavior and disease.

A working group comprising top neuroscience experts will continue to inform the development of NIH’s multi-year, trans-NIH scientific plan for The BRAIN initiative. This plan will outline the ultimate vision for NIH’s role in The BRAIN Initiative, including specific measurable goals and timetables. Based on this group’s preliminary recommendations, NIH initial efforts are focusing on building a new arsenal of tools and technologies for studying the brain. This state-of-the-art “toolbox” will include a systematic inventory of all the different types of cells in the brain, targeted genetic and non-genetic approaches for accessing specific cells and circuits, new and better capabilities for recording from rapidly firing collections of neurons, and interdisciplinary approaches to understanding how brain circuits produce unique human functions. NIH is also charting the course for the next generation of non-invasive imaging techniques that can be used to explore human brain functions and behaviors in real time.

The Opportunity, Growth, and Security Initiative proposed in the 2015 Budget would provide additional NIH funds for The BRAIN Initiative.

Defense Advanced Research Projects Agency: In FY 2015, DARPA plans to invest an estimated $80 million to support The BRAIN Initiative. DARPA's investments aim to leverage brain-function research to alleviate the burden of illness and injury and provide novel, neurotechnology-based capabilities for military personnel and civilians alike. In addition, DARPA is working to improve researchers' ability to understand the brain by fostering advancements in data handling, imaging, and advanced analytics.

In FY 2015, the Restoring Active Memory (RAM) effort will further develop memory prostheses as part of its larger effort to identify how memories are encoded in the brain during learning and skill acquisition with the ultimate goal of accelerating warfighter recovery after traumatic brain injury. DARPA's neuro- adaptive technology efforts, like Systems-Based Neurotechnology for Emerging Therapies (SUBNETS), aim to create closed-loop medical devices able to measure and modulate networks of neurons in research participants with intractable psychiatric illness and alleviate severe symptoms of diseases like post- traumatic stress disorder and major depression. DARPA's neuroscience technologies program will create interfaces for handling and analyzing large datasets of neural data, allowing investigators to rapidly and transparently solve complex problems of computation, generate new models, and model the brain in multiple dimensions and spatiotemporal scales. New military medical imaging efforts will provide new discovery tools capable of understanding structures of the behaving brain at high resolution in a stable manner over multiple experiments and generate tremendous amounts of data regarding the functional and structural connections between regions of the brain. Finally, the Prosthetic Hand Proprioception and Touch Interfaces (HAPTIX) effort will develop human-ready implantable electronic microsystems that monitor and modulate information in motor and sensory fibers of peripheral nerves, enabling amputees to achieve advanced and intuitive control and sensory functions with prosthetic limbs.

National Science Foundation: In FY 2015, NSF plans to invest $20 million to support The BRAIN initiative. To attain a fundamental scientific understanding of the complexity of the brain, in context and in action, NSF investments in The BRAIN Initiative are focused on generating an array of physical and conceptual tools needed to determine how healthy brains function over the lifespan of an organism, including humans. NSF will also focus on the development and use of these tools to produce a comprehensive understanding of how thoughts, memories, and actions emerge from the dynamic actions of the brain. NSF is prioritizing research in three areas where the agency's capacities are uniquely strong: integrative and interdisciplinary research; new theories, computational models, and analytical tools that will guide research questions and synthesize experimental data; and the development of innovative technologies and data infrastructure required to handle the large-scale datasets resulting from this research. Examples of investments that NSF has already made to support The BRAIN Initiative include a new $25 million Science and Technology Center on "Brains, Minds and Machines" at the Massachusetts Institute of Technology and new Research Coordination Networks (RCNs) to organize the scientific community and increase collaboration.

TEDBlog

Live From TED2014

The brain is a Swiss Army knife: Nancy Kanwisher at TED2014

Nancy Kanwisher. Photo: James Duncan Davidson

Onstage at TED, Nancy Kanwisher starts by telling us one of the most surprising results from recent neuroscience discoveries: The brain is not a general-purpose processor, but a collection of specialized components, "collectively building up who we are as human beings and thinkers."

Imagine, she says, walking into a daycare center and suddenly realizing you can't recognize any of the children, including your own. This isn't a strange fantasy. It's called prosopagnosia, and it happens to people. The really strange thing about it is, in those with that condition, only facial recognition is affected. There are many conditions like this, and Kanwisher say, "these syndromes collectively have suggested for a long time that the brain is divvied up into specific components."

The effort to identify these components has jumped with the invention of fMRI, functional magnetic resonance imaging. Imaging has been around for a while, but the real advance happened when people discovered how to map activity. When neurons fire, they need more blood. And blood flow is local. So fMRI lets us see what parts of the brain are more active than others.

So, what can you learn from this?

One of her first studies was about face recognition. It was known that prosopagnosia affected a specific region, but was there something special about that region in healthy brains too? She went into a scanner herself, looking at images of faces and objects, for two hours straight. ("As someone who has close to the world record of total number of hours spent in a scanner, I can tell you one of the most important skills for fMRI research is bladder control.")

The images were primitive by today's standards, but she found a region with higher activity. Was it a fluke? To test that, she repeated the test many times, and then scanned other people. It turns out almost everyone has a similar face-processing region in a similar part of the brain.

"But what does this region actually *do*?" Is it really face-recognition? Or does it do other things? Maybe it responds to any body part, or anything human, or anything round. She spent a lot of the next few years testing those hypotheses.

Does that nail it? Nope. Brain imaging can't tell you if the region is necessary for anything. "Brain imaging can only tell you what regions turn on and off. To tell what part of a brain is necessary for a function, you need to mess with it."

They did get one chance with an epileptic man. As part of a diagnostic procedure, electrodes were implanted to find the source of his epilepsy, and by chance two electrodes were in the face region. With his consent, they asked him what happened when they stimulated the region. When they did, he reported their face changed — into somebody he'd seen before. "This experiment finally nails the case," says Kanwisher. "This region is not only selectively involved in facial recognition but causally."

There are many other specialized parts of the brain. Kanwisher spent a lot of time in the scanner in the past month to show them to the TED audience. She takes us on a visual tour, showing the locations of regions that respond to:

- Faces.
- Color.
- Regions of space.
- Visual motion.
- Body parts.
- Hearing sounds with pitch (As opposed to sounds w/o pitch).
- Hearing sounds without pitch.
- Speech.

Are there specialized regions for complex processes? She says yes, including regions for:

- Language. A very specific part: understanding the meaning of a sentence.
- When you're understanding what another person is thinking, "the most amazing region we've found so far."

There are probably more to be discovered. "But importantly," she said, "I don't think we have specializations in the brain for every important mental function."

A few years ago a scientist in her lab thought he had found a special region for detecting food, which would be important for survival. But then he designed the crucial experiment to test that hypothesis. It turns out it wasn't about food, but colors and shapes. Again, not every process has a specific place in the brain.

So, how do we process all the other information? "In addition to these highly specialized components, we also have a lot of general-purpose machinery." They have found certain regions that seem to be engaged with any difficult task at all.

Kanwisher also points out that these regions are present in pretty much any normal brain in pretty much the same region — they are part of the fundamental machinery. It didn't have to be this way: "The brain could have been more like a kitchen knife than a Swiss Army knife." Instead we have a complex and rich picture of general-purpose components as well as highly specialized components.

It's still very much the early days of this kind of work, she says: "The most fundamental questions remain unanswered." For example:

- What do these regions do?
- Why do we need several face regions?
- How do they divide tasks?
- How are they connected?
- How does this very systematic structure get built? In an individual's development and through human evolution?

She closes by talking about the high cost of neuroscience research, and noting that many people justify it based on the promise of cures. Of course that's important, she says, but, "This is worth doing even if it never led to treatment for another disease. What could be more thrilling than to understand the fundamental mechanisms that underly human experience, who we are? This is, I think, the greatest scientific quest of all time."

MIT News
ON CAMPUS AND AROUND THE WORLD

Screen shots from a video of overlapping images of faces and houses, shown to subjects who were asked to pay attention to one or the other.

Video by Daniel Baldauf, screen shots colorized by MIT News

How the brain pays attention

Neuroscientists identify a brain circuit that's key to shifting our focus from one object to another.

Anne Trafton | MIT News Office
April 10, 2014

Picking out a face in the crowd is a complicated task: Your brain has to retrieve the memory of the face you're seeking, then hold it in place while scanning the crowd, paying special attention to finding a match.

A new study by MIT neuroscientists reveals how the brain achieves this type of focused attention on faces or other objects: A part of the prefrontal cortex known as the inferior frontal junction (IFJ) controls visual processing areas that are tuned to recognize a specific category of objects, the researchers report in the April 10 online edition of *Science*.

Scientists know much less about this type of attention, known as object-based attention, than spatial attention, which involves focusing on what's happening in a particular location. However, the new findings suggest that these two types of attention have similar mechanisms involving related brain regions, says Robert Desimone, the Doris and Don Berkey Professor of Neuroscience, director of MIT's McGovern Institute for Brain Research, and senior author of the paper.

"The interactions are surprisingly similar to those seen in spatial attention," Desimone says. "It seems like it's a parallel process involving different areas."

In both cases, the prefrontal cortex — the control center for most cognitive functions — appears to take charge of the brain's attention and control relevant parts of the visual cortex, which receives sensory input. For spatial attention, that involves regions of the visual cortex that map to a particular area within the visual field.

In the new study, the researchers found that IFJ coordinates with a brain region that processes faces, known as the fusiform face area (FFA), and a region that interprets information about places, known as the parahippocampal place area (PPA). The FFA and PPA were first identified in the human cortex by Nancy Kanwisher, the Walter A. Rosenblith Professor of Cognitive Neuroscience at MIT.

The IFJ has previously been implicated in a cognitive ability known as working memory, which is what allows us to gather and coordinate information while performing a task — such as remembering and dialing a phone number, or doing a math problem.

For this study, the researchers used magnetoencephalography (MEG) to scan human subjects as they viewed a series of overlapping images of faces and houses. Unlike functional magnetic resonance imaging (fMRI), which is commonly used to measure brain activity, MEG can reveal the precise timing of neural activity, down to the millisecond. The researchers presented the overlapping streams at two different rhythms — two images per second and 1.5 images per second — allowing them to identify brain regions responding to those stimuli.

"We wanted to frequency-tag each stimulus with different rhythms. When you look at all of the brain activity, you can tell apart signals that are engaged in processing each stimulus," says Daniel Baldauf, a postdoc at the McGovern Institute and the lead author of the paper.

Each subject was told to pay attention to either faces or houses; because the houses and faces were in the same spot, the brain could not use spatial information to distinguish them. When the subjects were told to look for faces, activity in the FFA and the IFJ became synchronized, suggesting that they were communicating with each other. When the subjects paid attention to houses, the IFJ synchronized instead with the PPA.

The researchers also found that the communication was initiated by the IFJ and the activity was staggered by 20 milliseconds — about the amount of time it would take for neurons to electrically convey information from the IFJ to either the FFA or PPA. The researchers believe that the IFJ holds onto the idea of the object that the brain is looking for and directs the correct part of the brain to look for it.

The MEG scanner, as well as the study's "elegant design," were critical to discovering this relationship, says Robert Knight, a professor of psychology and neuroscience at the University of California at Berkeley who was not part of the research team.

"Functional MRI gives hints of connectivity," Knight says, "but the time course is way too slow to show these millisecond-scale frequencies and to establish what they show, which is that the inferior frontal lobe is the prime driver."

Further bolstering this idea, the researchers used an MRI-based method to measure the white matter that connects different brain regions and found that the IFJ is highly connected with both the FFA and PPA.

Members of Desimone’s lab are now studying how the brain shifts its focus between different types of sensory input, such as vision and hearing. They are also investigating whether it might be possible to train people to better focus their attention by controlling the brain interactions involved in this process.

“You have to identify the basic neural mechanisms and do basic research studies, which sometimes generate ideas for things that could be of practical benefit,” Desimone says. “It’s too early to say whether this training is even going to work at all, but it’s something that we’re actively pursuing.”

The research was funded by the National Institutes of Health and the National Science Foundation.

MIT
Technology
Review

Mens et Apparata

Creating artificial intelligence turns out to be far more challenging than anyone expected. But the new Center for Brains, Minds, and Machines is ready to try again. This time, computer scientists, biologists, and neuroscientists will be tackling the problem together.

By Larry Hardesty on April 23, 2014

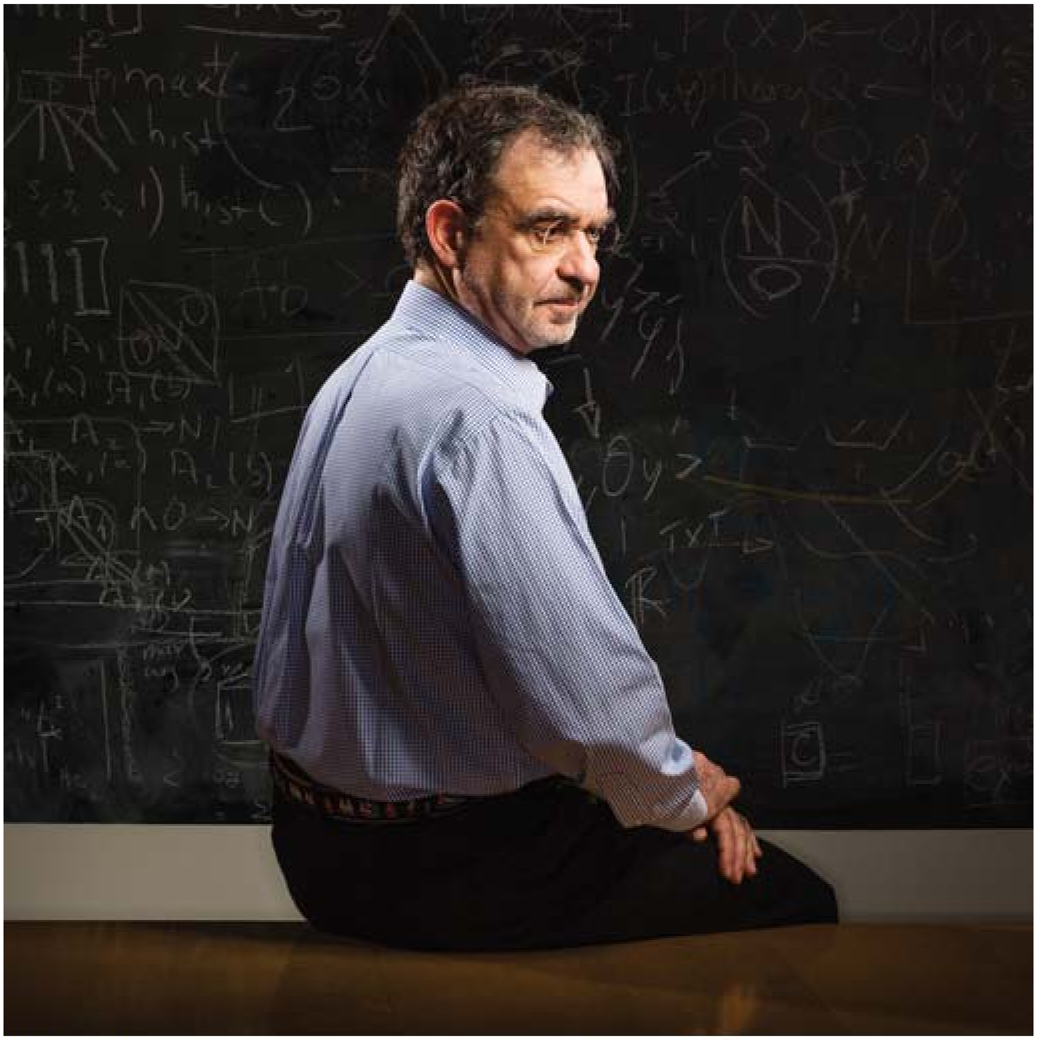

Tomaso Poggio

In the summer of 1955, a quartet of leading U.S. mathematicians—the term “computer scientist” wasn’t in use yet—proposed a conference at Dartmouth College to investigate a subject that they dubbed “artificial intelligence.” “The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it,” the proposal said.

The monthlong conference, which took place in 1956, is generally regarded as the genesis of artificial-intelligence research. Three of the proposal’s authors—LISP inventor John McCarthy, information theory pioneer Claude Shannon, SM ’40, PhD ’40, and future Turing Award winner Marvin Minsky—would later teach at MIT. McCarthy and Minsky (who remains on the faculty after 55 years) founded the MIT Artificial Intelligence Laboratory.

By 1967, progress in computing technology had been so rapid that Minsky, in his book *Computation: Finite and Infinite Machines*, was emboldened to write, “Within a generation, I am convinced, few compartments of intellect will remain outside the machine’s realm—the problems of creating ‘artificial intelligence’ will be substantially solved.”

Minsky’s prediction, of course, was overly optimistic. It turns out that winning at chess, which the early AI researchers took as the paradigmatic application of intelligence, is a much easier computational problem than, say, distinguishing spoken words or recognizing faces.

In the 1980s and ’90s, as the difficulty of replicating human intelligence became clear, AI came to mean something very different: practical, special-purpose computer systems often based on “machine learning,” which applies statistical analysis to huge numbers of training examples. That’s the approach that gave us voice-recognition systems and automatic text translators.

Now researchers at MIT believe it’s time to revive AI’s grand ambitions, in the hope of developing both better therapies for neurological disorders and computer systems that can anticipate our needs with humanlike intuition. And the National Science Foundation appears to agree. In September it announced a $25 million grant for the Center for Brains, Minds, and Machines (CBMM), which is based at MIT’s McGovern Institute for Brain Research. MIT is supplying 12 primary investigators; six others come from Harvard and five more from other institutions.

CBMM is led by Tomaso Poggio, a professor of brain sciences and human behavior and a principal investigator at both the McGovern Institute and the Computer Science and Artificial Intelligence Laboratory (CSAIL). His dual appointments illustrate the chief premise behind the new center: that we will make much faster progress toward understanding human intelligence if computational, biological, and psychological approaches are combined rather than explored in isolation.

“Instead of relying only on computer science, as they did 50 years ago, this center is really a bet that in order to replicate human intelligence, you need to understand more about the brain and about cognition,” Poggio says.

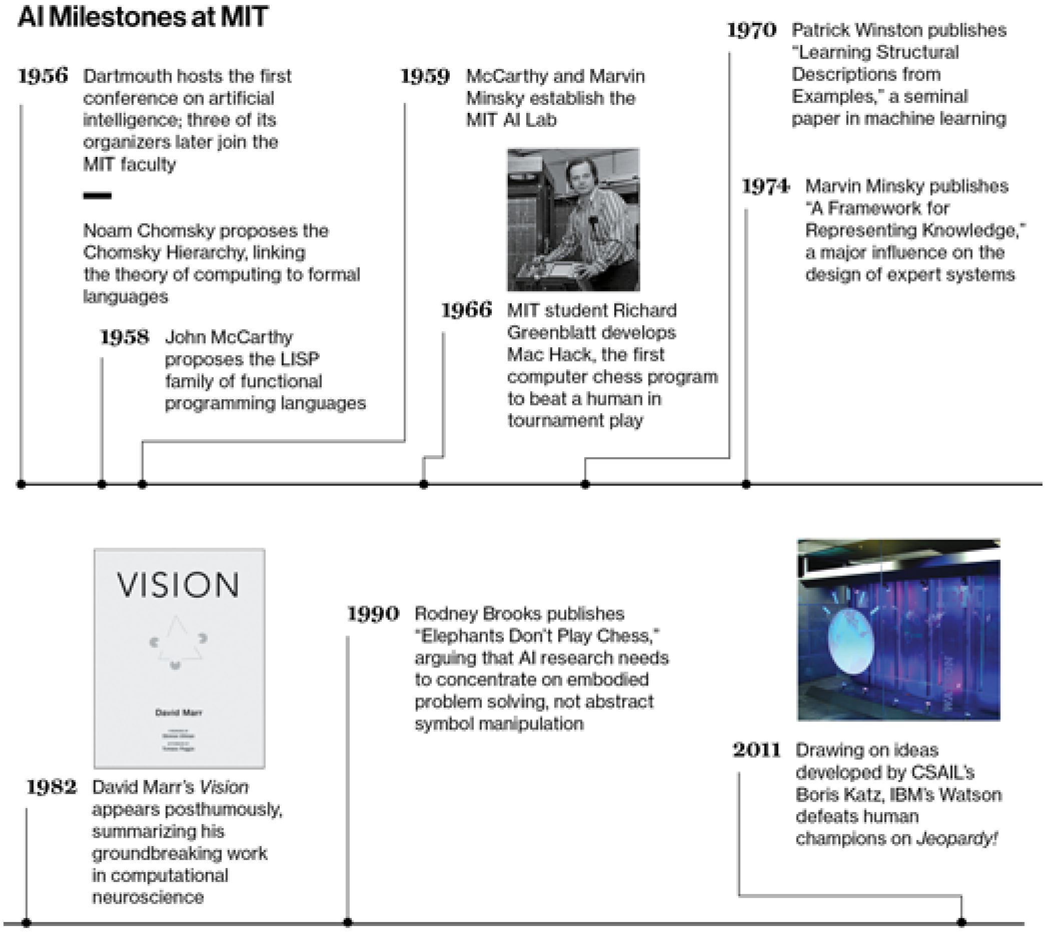

AI Milestones at MIT

1956 Dartmouth hosts the first conference on artificial intelligence; three of its organizers later join the MIT faculty

Noam Chomsky proposes the Chomsky Hierarchy, linking the theory of computing to formal languages

1958 John McCarthy proposes the LISP family of functional programming languages

1959 McCarthy and Marvin Minsky establish the MIT AI Lab

1966 MIT student Richard Greenblatt develops Mac Hack, the first computer chess program to beat a human in tournament play

1970 Patrick Winston publishes "Learning Structural Descriptions from Examples," a seminal paper in machine learning

1974 Marvin Minsky publishes "A Framework for Representing Knowledge," a major influence on the design of expert systems

1982 David Marr's *Vision* appears posthumously, summarizing his groundbreaking work in computational neuroscience

1990 Rodney Brooks publishes "Elephants Don't Play Chess," arguing that AI research needs to concentrate on embodied problem solving, not abstract symbol manipulation

2011 Drawing on ideas developed by CSAIL's Boris Katz, IBM's Watson defeats human champions on *Jeopardy!*

Patrick Winston, a professor in the Department of Electrical Engineering and Computer Science and CBMM's research coördinator, adds that technologies for investigating the problem have improved meaningfully in recent years. For one thing, Winston says, "computing is free: whatever type of computation needs to be done, it can be done." For another, "fMRI is now routine," he says, referring to functional magnetic resonance imaging, which can be used to study brain activity. He also points to technologies like transcranial magnetic stimulation, which can disrupt activity in targeted brain regions during cognitive tests, and optogenetics, a technique that uses light to selectively activate or silence genetically modified neurons in the brain. Optogenetics was pioneered by Ed Boyden '99, MEng '99, a Media Lab professor who is a principal investigator at the McGovern Institute and the new center.

Research at the center is organized into several major themes, or "thrusts": visual intelligence, which involves the integration of vision, language, and motor skills; circuits for intelligence, which will span research in neurobiology and electrical engineering; the development of intelligence; and social intelligence. Poggio, who is one of the primary investigators on visual intelligence, will also lead the development of a theoretical platform intended to tie together the work in the other areas.

Within each thrust, CBMM researchers are working to define a set of benchmark questions that they can use to assess their progress. Poggio offers one example, which relates to his own previous work on the visual system. Presented with an image of people interacting, an intelligent computer system should be able to provide plausible answers to five questions, ordered from easiest to hardest: What is in the image? Who is in it? What are the people doing? Who is doing what to whom? And what happens next?

Invariants

A theoretical framework for exploring all the questions surrounding human intelligence is a tall order. But Poggio's investigations into how the brain answers the first question on his list provide a sketch of what such a framework might look like.

Object recognition—developing computer systems that can answer the question "What is in the image?"—is a thriving area of artificial-intelligence research. Typically, object-recognition systems use some species of machine learning. Human beings label sample images, indicating which objects appear where, and the system tries to identify some common features that all images of the object share. "That's very different from human learning, or animal learning," he says. "When a child learns to recognize a bear or a lion, it's not that you have to show him pictures of a lion and a lion and a lion a million times. It's more like two or three times."

Poggio believes that unlike machine learning systems, the brain must represent objects in a way that is "invariant": the representation is the same no matter how big the object appears, where it is in the visual field, or whether it's rotated. And he also believes he has a plausible theory about what that representation might consist of.

Poggio's theory requires that the brain, or a computer system trying to simulate the brain, store one template of a few objects undergoing each type of variation—size, location, and rotation in the plane. The brain might, for instance, store a few dozen images of a human face tracing out a 360° rotation.

An unfamiliar object would then be represented as a collection of "dot products"—a standard computation in linear algebra—between its image and the templates. That collection would remain the same regardless of the object's size, location, or orientation.

One appeal of the theory is that the dot product reduces the comparison of two complex data sets, like visual images, to a single number. Collections of dot products, even for multiple templates, wouldn't take up much space in memory. Another appeal, Poggio says, is that "dot products are one of the easiest, maybe the easiest, computation for neurons to do."

In experiments, Poggio's system may not outperform machine learning systems. But it requires far fewer training examples, suggesting that it better replicates what the brain does. And for most computational tasks, the brain's approach usually turns out to be better.

Poggio believes that collections of dot products could anchor more abstract concepts, too. Templates that included different-shaped clusters of objects—arranged like dots on the face of a die, or in a line, or in a circle—could undergird the notion of number; a template of parallel lines viewed from different perspectives could undergird the notions of parallelism or perspective. "There may be more interesting things to explore," he says.

Fuzzy thinking

Like Poggio, Josh Tenenbaum is a professor in the Department of Brain and Cognitive Sciences (BCS) and a principal investigator in CSAIL. Although he leads CBMM's development thrust, which concentrates on the intuitive grasp of physics that even young children demonstrate, he has also done research that could contribute to theoretical work Poggio is leading.

The earliest AI research, Tenenbaum explains, focused on building a mathematical language that could encode assertions like "Birds can fly" and "Pigeons are birds." If the language was rigorous enough, researchers thought, computer algorithms would be able to comb through assertions written in it and calculate all the logically valid inferences.

But making sense of linguistic assertions turned out to require much, much more background information than anticipated. Not all birds, for instance, can fly. And among birds that can't fly, there's a distinction between a robin in a cage and a robin with a broken wing, and another distinction between any kind of robin and a penguin. Hand-coding enough of these commonsensical exceptions to allow even the most rudimentary types of inference proved prohibitively time-consuming.

With machine learning, by contrast, a computer is fed lots of examples of something and left to infer, on its own, what those examples have in common. (Given a million images of a lion, a machine learning algorithm can quantify its own guesses: 77 percent of images with these visual characteristics are images of lions.) But while this approach can work fairly well with clearly defined problems—say, identifying images of birds—it has trouble with more abstract concepts such as flight, a capacity shared by birds, helicopters, kites, and superheroes. And even flight is a concrete concept compared with, say, grammar, or motherhood.

Tenenbaum and his students have developed a new type of tool called a probabilistic programming language, which fuses what's best about AI old and new. Like the early AI languages, it includes rules of inference. But those rules are probabilistic. Told that the cassowary is a bird, a program written in Tenenbaum's language might conclude that cassowaries can probably fly. But if the program was then told that cassowaries can weigh almost 200 pounds, it might revise its probabilities downward.

"In the two earlier eras of AI, the biggest difference was symbols versus statistics," Tenenbaum says. "One of the things we've figured out on the math side is how to combine these, how to do statistical inference and probabilistic reasoning [with] these symbolic languages."

Reading people

The second of Poggio's five benchmark questions—Who is in the image?—has long been associated with the work of BCS professor Nancy Kanwisher, who is best known for using functional MRI to identify and analyze a region of the brain devoted to face perception.

Kanwisher leads the CBMM’s social-intelligence thrust, which she sees as the natural extension of her earlier work. “When you look at a face, you’re interested in more than just the basic demographic stuff, like what particular person that is, are they male or female, how old are they,” she says. “You can tell not just if a person’s happy or sad, but if they’re assertive or tentative, if they’re exuberant or passive—there’s a rich space of things that we can see in a face from a very brief glimpse.”

Similarly, Kanwisher says, humans can infer a great deal about people’s moods, intentions, and relationships with others from body language—which has the advantage of being amenable to computational modeling. She also points to the work of the late Nalini Ambady, the Stanford University social psychologist who developed the theory of “thin-slice judgments.”

“She videotaped TAs of Harvard courses teaching in front of their classes at the beginning of the semester,” Kanwisher says. “Then she showed very short segments of these videos to subjects in psychology experiments and said, ‘Rate the effectiveness of this teacher.’ All they have is a few seconds of a person in front of a room talking to a class—you can’t even hear what they’re saying. And she found that those ratings were highly correlated with the ratings of that person’s actual students.”

The first project of the CBMM’s social-intelligence thrust, Kanwisher says, will be to design a set of experimental tasks that allow researchers to quantify human social perception. Once the researchers establish a baseline, they can study such things as how performance on the tasks develops through childhood, or how autistic children’s performance differs from that of other children. They could also identify the brain regions involved in social perception by using fMRI to measure neural activity or transcranial magnetic stimulation to disrupt performance. And after collecting all of this data, they will try to computationally model what, exactly, the brain is doing.

Follow the story

The later questions on Poggio’s list—“Who is doing what to whom?” and “What happens next?”—fascinate Patrick Winston. He believes that the defining feature of human intelligence is the ability to tell and understand stories. That ability even plays a role in image labeling. As Winston likes to point out, a human subject will identify an image of a man holding a glass to his lips as that of a man drinking. If the man is holding the glass a few inches farther forward, he’s toasting. But a human will also identify an image of a cat turning its head up to catch a few drops of water from a faucet as an instance of drinking. “You have to be thinking about what you see there as a story,” Winston says. “They get the same label because it’s the same story, not because it looks the same.”

That’s one reason for dedicating a research thrust to the integration of vision, language, and social and motor skills. To illustrate another reason, Winston points to an experiment conducted by developmental psychologist Elizabeth Spelke, a former MIT faculty member who is now at Harvard and is one of the primary investigators in the development thrust. Spelke was intrigued by experiments in which researchers had placed rats on a rotating platform in the center of a room. Food was visibly placed in one corner of the room but then masked from view. Identical masks were placed in the other three corners, and the platform was rotated. Spelke decided to extend that study to human children and adults, hiding a toy or a ring of keys instead of food.

With all animals, children, and adults, once the rotating stopped, the test subject would head with equal probability to either the corner with the masked object or the one diagonally across from it, which had the same relationship to the subject. Both groups of researchers also varied the experiment, painting a different color on one of the walls adjacent to the corner where the object was placed. Animals and small children still selected either the correct corner or the one opposite it with equal probability, but adults could now reliably retrieve the object.

Here’s where things get interesting. If the adults were asked to listen to a passage of text and recite it back before heading to the object, they reverted to confusing the diametrically opposite corners. Hearing and reciting the text “consumes the human language processor, and that reduces them to the level of a rat,” Winston says. “Afterward they’ll say, ‘Yeah, I could see the blue wall, but I couldn’t quite use it.’”

Answering the highest-level questions on the CBMM researchers’ lists of benchmarks will probably take much longer than the five years of the initial NSF grant. But, Poggio says, “it’s time to try again. It’s been 50 years. We don’t know whether it will work this time. But if we don’t try, we won’t know.”

MIT News
ON CAMPUS AND AROUND THE WORLD

Illuminating neuron activity in 3-D

New technique allows scientists to monitor the entire nervous system of a small worm.

Anne Trafton | MIT News Office

May 18, 2014

Researchers at MIT and the University of Vienna have created an imaging system that reveals neural activity throughout the brains of living animals. This technique, the first that can generate 3-D movies of entire brains at the millisecond timescale, could help scientists discover how neuronal networks process sensory information and generate behavior.

The team used the new system to simultaneously image the activity of every neuron in the worm Caenorhabditis elegans, as well as the entire brain of a zebrafish larva, offering a more complete picture of nervous system activity than has been previously possible.

“Looking at the activity of just one neuron in the brain doesn’t tell you how that information is being computed; for that, you need to know what upstream neurons are doing. And to understand what the activity of a given neuron means, you have to be able to see what downstream neurons are doing,” says Ed Boyden, an associate professor of biological engineering and brain and cognitive sciences at MIT and one of the leaders of the research team. “In short, if you want to understand how information is being integrated from sensation all the way to action, you have to see the entire brain.”

The new approach, described May 18 in *Nature Methods*, could also help neuroscientists learn more about the biological basis of brain disorders. “We don’t really know, for any brain disorder, the exact set of cells involved,” Boyden says. “The ability to survey activity throughout a nervous system may help pinpoint the cells or networks that are involved with a brain disorder, leading to new ideas for therapies.”

Boyden’s team developed the brain-mapping method with researchers in the lab of Alipasha Vaziri of the University of Vienna and the Research Institute of Molecular Pathology in Vienna. The paper’s lead authors are Young-Gyu Yoon, a graduate student at MIT, and Robert Prevedel, a postdoc at the University of Vienna.

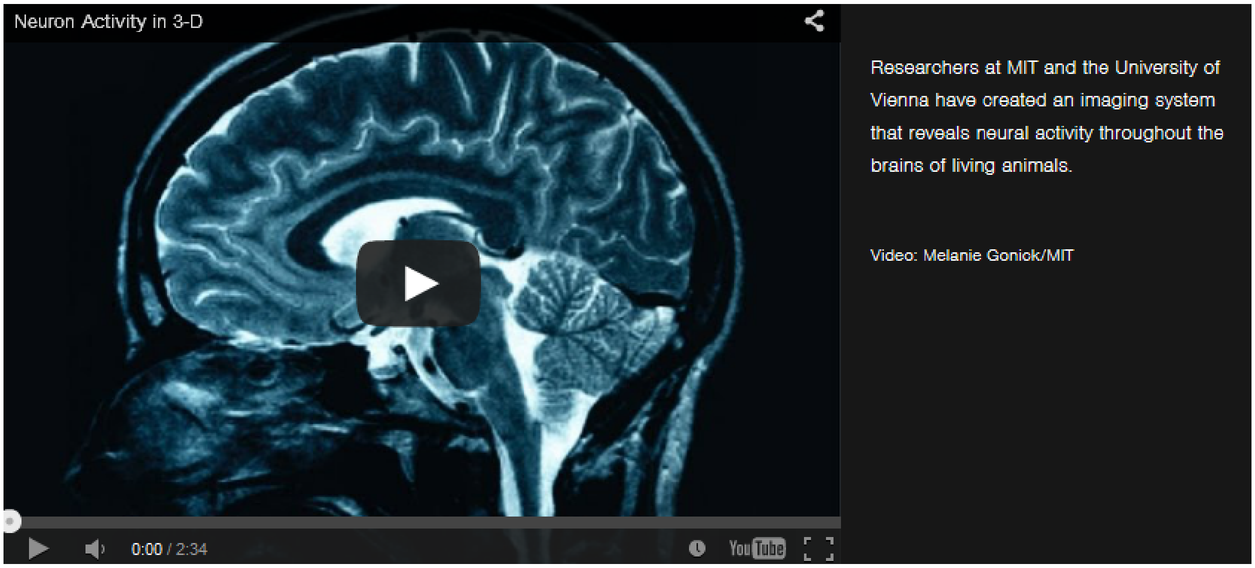

Researchers at MIT and the University of Vienna have created an imaging system that reveals neural activity throughout the brains of living animals.

Video: Melanie Gonick/MIT

High-speed 3-D imaging

Neurons encode information — sensory data, motor plans, emotional states, and thoughts — using electrical impulses called action potentials, which provoke calcium ions to stream into each cell as it fires. By engineering fluorescent proteins to glow when they bind calcium, scientists can visualize this electrical firing of neurons. However, until now there has been no way to image this neural activity over a large volume, in three dimensions, and at high speed.

Scanning the brain with a laser beam can produce 3-D images of neural activity, but it takes a long time to capture an image because each point must be scanned individually. The MIT team wanted to achieve similar 3-D imaging but accelerate the process so they could see neuronal firing, which takes only milliseconds, as it occurs.

The new method is based on a widely used technology known as light-field imaging, which creates 3-D images by measuring the angles of incoming rays of light. Ramesh Raskar, an associate professor of media arts and sciences at MIT and an author of this paper, has worked extensively on developing this type of 3-D imaging. Microscopes that perform light-field imaging have been developed previously by multiple groups. In the new paper, the MIT and Austrian researchers optimized the light-field microscope, and applied it, for the first time, to imaging neural activity.

With this kind of microscope, the light emitted by the sample being imaged is sent through an array of lenses that refracts the light in different directions. Each point of the sample generates about 400 different points of light, which can then be recombined using a computer algorithm to recreate the 3-D structure.

“If you have one light-emitting molecule in your sample, rather than just refocusing it into a single point on the camera the way regular microscopes do, these tiny lenses will project its light onto many points. From that, you can infer the three-dimensional position of where the molecule was,” says Boyden, who is a member of MIT’s Media Lab and McGovern Institute for Brain Research.

Prevedel built the microscope, and Yoon devised the computational strategies that reconstruct the 3-D images.

Aravinthan Samuel, a professor of physics at Harvard University, says this approach seems to be an “extremely promising” way to speed up 3-D imaging of living, moving animals, and to correlate their neuronal activity with their behavior. “What’s very impressive about it is that it is such an elegantly simple implementation,” says Samuel, who was not part of the research team. “I could imagine many labs adopting this.”

Neurons in action

The researchers used this technique to image neural activity in the worm C. elegans, the only organism for which the entire neural wiring diagram is known. This 1-millimeter worm has 302 neurons, each of which the researchers imaged as the worm performed natural behaviors, such as crawling. They also observed the neuronal response to sensory stimuli, such as smells.

The downside to light field microscopy, Boyden says, is that the resolution is not as good as that of techniques that slowly scan a sample. The current resolution is high enough to see activity of individual neurons, but the researchers are now working on improving it so the microscope could also be used to image parts of neurons, such as the long dendrites that branch out from neurons’ main bodies. They also hope to speed up the computing process, which currently takes a few minutes to analyze one second of imaging data.

The researchers also plan to combine this technique with optogenetics, which enables neuronal firing to be controlled by shining light on cells engineered to express light-sensitive proteins. By stimulating a neuron with light and observing the results elsewhere in the brain, scientists could determine which neurons are participating in particular tasks.

Other co-authors at MIT include Nikita Pak, a PhD student in mechanical engineering, and Gordon Wetzstein, a research scientist at the Media Lab. The work at MIT was funded by the Allen Institute for Brain Science; the National Institutes of Health; the MIT Synthetic Intelligence Project; the IET Harvey Prize; the National Science Foundation (NSF); the New York Stem Cell Foundation-Robertson Award; Google; the NSF Center for Brains, Minds, and Machines at MIT; and Jeremy and Joyce Wertheimer.