FastLane **PI/CO-PI Management -**
Proposal Functions | HOME

NSF Home | **News** | **Site Map** | **GPG** | **AAG** | **Contact Us** | **FastLane Help**
Change Password | **Logout**

Proposal Status | MAIN

Organization: Massachusetts Institute of Technology

Proposal Detail:

Proposal Information

Proposal Number:	1231216
Proposal Title:	A Center for Brains, Minds and Machines: the Science and the Technology of Intelligence
Received by NSF:	02/02/12
Principal Investigator:	Tomaso Poggio
Co-PI(s):	Matthew Wilson Lakshminarayan Mahadevan Haym Hirsh

This Proposal has been Electronically Signed by the Authorized Organizational Representative (AOR).

NSF Program Information

NSF Division:	Office of Integrative Activities
NSF Program:	Science and Technology Centers (Integrative Partnerships Program)
Program Officer:	Dragana Brzakovic
PO Telephone:	(703) 292-8040
PO Email:	dbrzakov@nsf.gov

Proposal Status

Status As of Today Dated: **02/06/12**

This proposal is currently being considered by the above Program. Consideration of proposals usually requires up to six months.

◀ Back to Status Search Results

Download Adobe Acrobat Reader for viewing PDF files

Project Personnel

Center Management:

Director: Tomaso Poggio (MIT)
Assistant Directors: Haym Hirsh (Rutgers); L. Mahadevan (Harvard) and Matthew Wilson (MIT)
Coordinator for Diversity: Mandana Sassanfar (MIT)
Coordinator for Education: Haym Hirsh (Rutgers)
Coordinator for Knowledge Transfer: L. Mahadevan (Harvard)
Coordinator for Research: Patrick H. Winston (MIT)

Key Personnel:

Boyden, Ed, MediaLab, BCS Dept., BioEng Dept., MIT
Buelthoff, Heinrich, Max Planck Institute, Tuebingen
Desimone, Robert, BCS Dept., MIT
Freiwald, Winrich, Rockfeller University
Goodman, Noah, Psych Dept., Stanford University
Hirsh, Haym, CS Dept, Rutgers University
Kaelbling, Leslie, EECS Dept., MIT
Kanwisher, Nancy, BCS Dept., MIT
Katz, Boris, CSAIL, MIT
Koch, Christof, Caltech and Allen Institute
Kreiman, Gabriel, Children's Hospital, Harvard University
Logothetis, Nikos, Max Planck Institute, Tuebingen
Mahadevan, L., Harvard University
Metta, Giorgio, Italian Institute of Technology (IIT)
Nakayama, Kenneth, Harvard University
Poggio, Tomaso, BCS Dept., CSAIL, MIT
Richards, Whitman, BCS Dept., CSAIL, Media, MIT
Rosasco, Lorenzo, Italian Institute of Technology (IIT)
Sassanfar, Mandana, Dept. of Biology, MIT
Saxe, Rebecca, BCS Dept., MIT
Schulz, Laura, BCS Dept., MIT
Shashua, Amnon, Hebrew University of Jerusalem; MobileEye
Smale, Steve, City U, Hong Kong
Sompolinsky, Haym, Harvard University and Hebrew University
Spelke, Liz, Psych Dept., Harvard University
Tenenbaum, Josh, BCS Dept., CSAIL, MIT
Ullman, Shimon, Weizmann Institute and BCS Dept., MIT
Valiant, Leslie, Harvard University
Verri, Alessandro, U. Genoa
Wilson, Matthew, Picower Institute, BCS Dept., MIT
Winston, Patrick, EECS Dept., CSAIL, MIT
Yuille, Alan, UCLA

External Advisory Committee (EAC)

Joel Oppenheim, Senior Associate Dean for Biomedical Sciences at NYU, and Director of the Sackler Institute of Graduate Biomedical Science

Partner Institutions and Organizations (organized by category)

Academic Institutions:
Caltech
Harvard University
MIT
Rockefeller
Rutgers University
Stanford University
UCLA

Minority-serving Partnerships:
Howard University
Hunter College (Prof. Susan Epstein, CS)
Universidad Central del Caribe (UCC), Puerto Rico (Prof. Maria Bykhovskaia, Chair, Neuroscience Department)
Wellesley College (Prof. Ellen Hildreth, Chair, CS)

International Partnerships:
City University, Hong Kong (Prof. Steve Smale)
Hebrew University (Profs. Haym Sompolinsky, Amnon Shashua)
Italian Institute of Technology (IIT) (Profs Giorgio Metta, Lorenzo Rosasco)
Max Planck Institute, Tuebingen (Profs Heinrich Buelthoff, Nikos Logothetis)
National Center for Biological Sciences, Bangalore, India (Professor K VijayRaghavan)
University of Genova (Prof Alessandro Verri)
Weizman Institute (Prof Shimon Ullman)

Industrial and technological partnerships (tentative)
Google (Dr. Peter Norvig, Director of Research)
Microsoft
IBM (Dr. David Ferrucci)
Orcam (Dr. Ammnon Shashua)
Mobileye (Dr. Ammnon Shashua)
DeepMind (Demis Hassabis)

2. Project Summary

Imagine a world where intelligence and how it emerges from brain activity is truly understood. Instead of narrow systems, such as Watson and Siri, imagine really smart systems that would change the world. Instead of knowing what works in K-12, researchers and educators would know why it works. The proposed Center for Brains, Minds, and Machines (CBMM) could revolutionize K-12, and also 0-K, and 12-life. This effort could revolutionize the education of people with special needs and unusual learning styles, and provide compassionate care for the aged and challenged with robots that anticipate needs. Systems that recognize how culture influences thinking could help avoid social conflict. The work of scientists and engineers could be amplified to help solve the world's most pressing technical problems. Mental health could be understood on a deeper level to find better ways to intervene. These accomplishments will take decades. But the time has come to actively and energetically pursue the kind of research that begins to achieve these ambitious goals in both scientific foundations and practical applications.

How projects become more than the sum of the parts Understanding intelligence encompasses four main areas of inquiry and real progress requires unified work in all four. First, intelligence goes beyond narrow expertise, bridging several domains including vision, planning, and language. Second, understanding intelligence requires an appreciation for how it develops and the twin roles of learning and innate structure. Third, understanding the physical machinery of intelligence spans multiple levels of analysis from neural circuits to large-scale brain architecture. Finally, intelligence emerges from the interactions among individuals – the product of social interactions. Therefore an integrated research agenda engages four major research thrusts with interlocking teams and working groups as described in this proposal.

Perhaps the most distinctive aspect of the Center's vision is a common theoretical, mathematical, and computational platform underlying the four thrusts, *"Enabling Theory."* The Center's success depends on the use and development of revolutionary new languages and tools for rigorous theory building based chiefly on machine learning, probabilistic modeling, and inference, that for the first time can span and unite multiple disciplines, multiple levels of analysis, as well as science and engineering goals.

Why is this the right time? The prospect of understanding intelligence is tantalizing and timeless. Yet, now is a uniquely appropriate time to pursue this vision. The full problem cannot be solved in a decade, but this integrated effort is expected to make fundamental, unprecedented progress with great value to science, technology, and society for decades to come.

Integrated efforts similar to those described here have been attempted at least twice before: first in the early days of AI, cognitive science, and neuroscience in the 1960's; and again in the mid 1980's with the advent of neural networks. Since then, transformative developments have made it possible to pursue this effort in ways that were simply impossible before. Integrating intelligence across vision, language, planning, and other domains is possible now because of advances in understanding the cognitive, neural, and computational bases of each domain. A developmental approach to intelligence is now possible because of revolutionary work in infant cognition based on looking time, reaching, and increasingly large-scale video data collection, and physiological and neuroimaging methods. Neuroimaging plus genetic techniques for dissecting the neural circuits underlying learning and decision now allow the study of intelligence and the brain at both large and small scales. Recent discoveries about brain systems and innate behavioral capacities for social cognition, and the development of robotic and AI systems that for the first time interact intelligently with humans, make it clear that we must be studying intelligence in a social context.

The education and human resources objectives of the STC are to: a) Produce a generation of leaders that is equally knowledgeable in computer science and neurocognitive science, and b) Attract even more underrepresented groups to interdisciplinary research and education. These objectives will be met by:

- Developing an integrated system of graduate courses to cross educate computer engineers and neuroscientists

- Establishing a graduate teaching consortium via the Internet that offers a broad selection of courses in the science and the engineering of intelligence,
- Involving minorities in research
- Nurturing student leadership by forming a student council
- Developing internships with international collaborators

A novel focus of this Education effort -- based at Rutgers -- is to pursue the use of technology for developing Internet-based courses offered to students worldwide with infrastructures such as the new MITx platform. The online courses will be designed collaboratively with instructors at partnering institutions, particularly those at minority serving colleges, to provide scalable online offerings accessible to suitably motivated individuals. Unique to the Education effort at Rutgers are online educational experiments designed to optimize and personalize the Internet-based curriculum by using modern data analysis tools and sound experimental design.

The diversity objective of the Center is to attract more women and minorities to computational and cognitive neuroscience, and artificial intelligence. The diversity effort -- based at MIT -- involves partnering with several minority-serving institutions to increase exposure to these fields through a seminar series, workshops, collaborative research, novel curriculum, and summer research opportunities for undergraduates. All faculty will contribute to this effort.

The knowledge transfer objectives are to: a) Widely disseminate research findings with a combination of Internet based and traditional methods; b) Reach out to the public and to establish partnerships with minority-serving institutions; and c) Develop collaborative relationships with commercial industries and international institutions devoted to the science and engineering of intelligent machines. The Center's knowledge transfer effort -- based at Harvard -- focuses on disseminating the teaching materials from the consortium courses through MITx, a new online teaching and certification platform (that replaces OpenCourseWare), and developing online videos for K-12 students to attract a new generation of students to the science and engineering of intelligence. International symposia will focus on the Center's themes.

The Center's interdisciplinary team includes science and engineering researchers at MIT, Harvard, and Rutgers University with expertise in: Biology, Computer Science, Cognitive Science, Neurobiology, Physics, and Statistics. Harvard will lead the research thrust focused on basic neurobiology, whereas MIT will lead the remaining three thrusts focused on systems and computation. All Institutions are involved in the theory platform, which includes the best computational neuroscientists in Cambridge. Rutgers is coordinating the Education effort and conducting research on its online implementations with potential implications for a real engineering of Education on Intelligence. The three institutions have formed a strong partnership focused on achieving the Center's mission. A faculty and student exchange program will facilitate collaboration among the Institutions.

The intellectual merit of the Center is its ability to elucidate the mechanisms and architecture of the most intelligent system known: the human brain. The Center aims to understand, at the computational level, how human intelligence emerges from the integration of nature and nurture—innate capacities and learning—and how we may reproduce this aspect of intelligent behavior in machines. Success in this project will ultimately enable us to understand ourselves better, to produce more intelligent machines, and perhaps even to make ourselves smarter.

The broader impact of this program will be 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 more 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.

4. Project Description

4.a. Rational for Center Approach

The two major objectives of the Center for Brains, Minds, and Machines (CBMM) are 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 integrates cognitive science, neuroscience, computer science, and artificial intelligence. Human intelligence is marked by its breadth, including the ability to understand people and natural phenomena by using vision and language.

Why? There is perhaps no greater scientific and engineering mission than understanding intelligence, because any progress toward improving intelligence in brains and machines contributes to other great challenges in science and society, both theoretical and practical. Research on intelligence will help us understand the human mind and brain, build more intelligent machines, and improve the mechanisms for collective decisions. These advances will be critical to future prosperity, education, health, and the security of our society.

Why now? The convergence and recent progress in technology, mathematics, and neuroscience has created a new opportunity for synergy across fields [3, 24, 29, 45, 56, 60, 87, 99]. The past two decades have brought significant developments in computational power, the mathematics of learning, and in our ability to interrogate neuronal circuits at unprecedented resolution. Practical successes [11, 35, 59, 79] illustrate the promise of what can be achieved as well as the need for more fundamental research.

Why a Center? We are uniquely poised at this juncture to make dramatic progress because of significant advances in several subfields of neuroscience and computer science, and the appearance of a core set of experimental tools and theoretical frameworks that span much of the area. To get to the next level we need an unprecedented interdisciplinary approach because the advances we anticipate require mastery of a large range of knowledge and scholarly methods. The STC would be designed to be a catalyst for a 'phase transition' in this domain. Recent successes in AI show that the key to understanding human intelligence and being able to replicate it in machines that can pass a Turing test is the deep problem of *integrating* multiple facets of human intelligence. Therefore, we need not only neuroscientists to work with computer scientists, but also vision researchers to work with language and social cognition researchers. No single discipline can solve this problem: intelligence is a truly interdisciplinary problem and the structure of our Center—with its four thrusts and enabling theory platform—is designed to provide this integration.

Why do we believe we can be successful? For the first time in history, our digital computers are approaching the level of human brains in terms of raw processing power and memory. Over the past decade we have made great advances in neuroscience; the neuroscience of the cortex, for instance, has led to vision systems as good or better than engineered ones in restricted tasks. More important, for the first time we have a common theoretical language across different disciplines: the language of learning theory and statistics, the mathematics of prediction, of probabilistic inference, and of modeling. We also have new tools (optogenetics, TMS) for switching on and off circuits of a few hundred thousand neurons and looking at the impact on intelligent behavior. Technologies developed over the past two decades allow us for the first time to record brain activity in humans (fMRI, MEG, intracranial arrays) while they perform intelligent tasks.

4.b. Research Plans and Objectives

Potential legacy and national and global impact of the proposed CBMM

In the 2007 report "Rising Above the Gathering Storm," the National Academies urged that "the United States must compete by optimizing its knowledge-based resources, particularly in science and technology..." To address these challenges and opportunities, scientists and engineers at MIT, Harvard, Rutgers, Rockefeller, Stanford, Caltech, and UCLA propose to form a partnership STC to conduct highly interdisciplinary research, train the next generation of leaders in research, engineering, and education, and widely disseminate this knowledge to impact society. Preliminary work that led to the present proposal started in 2008 with the beginning of an ***Intelligence Initiative*** ***(I^2)*** at MIT (http://isquared.mit.edu/). A central goal of I^2, which was strongly supported by the MIT administration, was to enable more interdisciplinary and integrative approaches than do conventional institutional structures. Thus, the "I^2" label also stands for integrative intelligence research. We organized a "Brains, Minds and Machines" symposium (http://mit150.mit.edu/symposia/brains-minds-machines) in May 2011 where the enthusiasm of the speakers and audience was palpable. We also started a small set of seed projects around the theme of intelligence (see http://isquared.mit.edu/research). This initial work has already helped forge new research collaborations and has illustrated broad interest in the scientific vision of the Center.

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 operational definition of intelligence provided by the Turing test best captures the ultimate goals of CBMM research: to understand the brain and to replicate the human mind in machines. The vision is based on an *integrative approach* combining experimental techniques in neuroscience and cognitive science with computational modeling. This integrative approach will also combine the study of the different aspects of perception, action, and cognition. We believe that real progress will come only from taking seriously four aspects of intelligence: the *integrative aspect* of intelligence, its *development*, the *wetware*, and *social* intelligence. We have therefore constructed a research agenda organized into **four major research thrusts**

- *3-6-12 Reverse Engineering the Infant Mind*
- *Neuronal circuits underlying intelligence*
- *Integrating intelligence: vision, language, and social interactions*
- *Social intelligence.*

together with the development of a unifying mathematical framework (*Enabling Theory).* These projects are interconnected in many ways and will be much stronger when tackled together, with interlocking teams and working groups as described in this proposal.

4.b.1 Thrust 1: *3-6-12 Reverse Engineering the Infant Mind* (Tenenbaum, Schulz, Goodman, Mahadevan, *Metta, Sandini, Saxe, Spelke, Ullman*)[1]

4.b.1.1 Vision and goals

Understanding the development of intelligence in a human infant is a key project of CBMM. This project engages the fundamental tradeoff between nature and nurture, or priors and data, and ultimately the origin of priors—how constraints are selected by evolution, encoded in genes, and instantiated in genetically wired brain circuits.

This project also represents a novel developmental approach to building human-like artificial intelligence systems and comprehensive computational models of human cognitive architecture. Instead of branching out from a single area of adult human intelligence, we start with an integrated cognitive model of multiple core capacities in the young child's mind, along with a set of developmental or learning mechanisms for scaling up to an adult model [86]. We build on the last two decades of research in cognitive development,

---

[1] Secondary members of the projects are in italics

specifically studies by Spelke [81], Carey [16], and others, suggesting that infants even as young as 3 months old have powerful knowledge of the physical and social worlds—an intuitive physics of objects and their dynamics, and an intuitive psychology of intentional agents, social groups and roles, and moral behavior—along with the beginnings of abstract symbolic thought, such as the ability to quantify sets of objects, that form the foundations for emerging language abilities. Research has also mapped out how these core systems develop over the first two years of life: what changes between 3 and 6 months, 6 and 9 months, 9 and 12 months, etc., and how developments in intuitive physics, psychology, and language build on each other [41, 42, 74].

The goal of the 3-6-12 project is to characterize in more precise computational terms the content of these systems of knowledge, how they support infants' inferences and actions in the world, and how they arise through mechanisms of learning and development. We will test our computational models quantitatively against infant and child behavior [91] and in state-of-the-art humanoid robotics technology [57, 101].

4.b.1.2 Approach

1.2.1 Core empirical phenomena. How infants (or adults) parse the world into objects and animate goal-directed agents can be studied with very simple displays of point-like objects with simple shapes (triangles, circles) moving rigidly in simple, two-dimensional spatial environments [4, 44, 94]. These stimuli can be fully characterized as low-dimensional time series yet they evoke strong percepts of physical force dynamics (bashing, piercing, breaking, pulling, pushing), future tendencies (precariousness, stability), and latent physical properties (mass, friction, elasticity). They also evoke rich psychological percepts: intentional states of single agents (beliefs, desires, plans, preferences) and social relations between agents (helping, hindering, chasing, fleeing). Infants perceive some version of all these abstract concepts; adults (and perhaps infants too, it is not known) can additionally perceive personalities, social alliances, emotions, achievements and frustrations, and embarrassments.

Our empirical work on this thrust consists of two phases. The first phase (year 1) maps out timelines (based largely on prior studies) for the development of the capacities described above over the first two years of life, focusing on physical and psychological reasoning that develop in significant ways over this period [100]. These timelines drive model development described below. The second phase (year 2 and following) focuses on designing new behavioral experiments with infants to test distinctive model predictions. Our experiments use standard looking-time and reaching methods. We also anticipate using novel methods with the potential to generate richer quantitative data on behavior and neural substrates, including computer-based eye-gaze tracking and NIRS.

1.2.2 Models. We model infants' intuitive psychology, intuitive physics, and language as forms of causal inference: infants build (implicit) probabilistic models of the causal structure in the world that gives rise to the data observed, and their behavior can be explained in terms of approximate Bayesian inference to the causes that best explain the data. Key questions are: what kinds of causal knowledge, what kinds of representations, and what kinds of approximate inference algorithms are instantiated in infants' minds? What is innate, and what learning algorithms build new models? How can these formalisms be tested in infant behavior?

We focus on a class of probabilistic models known as *probabilistic programs* [40]. Probabilistic programs are a generalization of *probabilistic graphical models*, a paradigm that has transformed AI [62], machine learning, and cognitive modeling over the past two decades with a powerful toolkit for describing and reasoning about complex environments with pervasive uncertainty. Probabilistic programs (see **Enabling Theory**) combine key features of graphical models with the expressiveness needed to handle the core abstract concepts that emerge in infants' intuitive physics and intuitive psychology (e.g., analogs of force, mass, utility, or belief), the capacity for recursive and reflective reasoning (e.g., reasoning about an infant's beliefs about how other agents form beliefs, the central focus of "theory of mind"), and the ability to cast learning the structure of intuitive theories as a well-posed probabilistic inference. For example, probabilistic programs capture intuitive psychology by representing intentional agents as approximately rational planners, who stochastically choose actions that tend to maximize their expected utilities given environmental constraints [40]. Goals are represented as sources of positive utility; actions such as moving in the environment have costs. Agents' beliefs (true or false) are represented in the expectations component of the expected utility computation [4]. Social goals are represented by recursive utility functions: A is trying to help B if A's utility depends positively on B's [94]. Judgments about agents' goals are cast as Bayesian inferences about which goals are most likely to have given rise to observed behavior under this probabilistic planning model.

Probabilistic programs allow us to model the development of infants' intuitive theories—how they learn from experience to more faithfully and richly model the world's causal structure—by combining ideas from hierarchical Bayesian learning, program induction, inductive logic programming, and stochastic search [86, 95]. Probabilistic programs also have a natural sampling-based semantics, suggesting inference mechanisms that are computationally efficient as well as cognitively and biologically plausible: instead of computing high-precision probabilities, an inference agent imagines a small sample of possible interpretations constructed stochastically to have high posterior probability given observed data. We have shown how both adult judgments and standard looking-time (violation-of-expectancy) measures of infants' knowledge can be predicted quantitatively by probabilistic models with relatively few posterior samples [91, 102]. We plan to explore combinations of sampling-based inference with other approaches based on stochastic search, bottom-up data-driven heuristics, and constrained optimization.

We already have achieved exciting preliminary results on all aspects of this integrated approach: drafts of partial developmental timelines, models of several core capacities of infant cognition, successful experimental tests of those models (specifically, how infants track moving objects [91], infer hidden properties of objects [42], and judge social goals of agents [94]), and proof-of-principle learning models for several very simple theories [86, 95]. These results suggest the viability of our approach and its potential for high impact; two recent papers on this work appeared in *Science* and one in *PNAS*, and they were covered prominently in the media (e.g., *Boston Globe* front page, and top stories of 2011 in *Discover* magazine).

1.2.3 Applications to robotics We plan to pursue several lines of application to building machine systems with more human-like intelligence in domains representing extensive technological interest and need. Specifically, we plan to explore intuitive theories of physics and psychology of mind by putting them into autonomous robots that interact with humans in close quarters, where these models should be invaluable for efficient and rewarding joint human-robot activities, and in prototype systems for autonomous urban driving, to reason about intentions of other drivers and pedestrians in more human-like ways. Robotics platforms also allow us to test the engineering adequacy our models of infant cognition in physically embodied forms, specifically on the iCub open-source test-bed provided by our partner IIT, which is equipped with biologically-inspired control and hardware, and is similar in size to a small child [36, 73, 88].

4.b.1.3 Synergies with other center thrusts

The project of reverse engineering infant cognition dovetails closely with other center thrusts. This project is a principal motivation for developing computational methods of hierarchical Bayesian learning and probabilistic programs (**Enabling Theory)**. The problems of scene understanding to be studied here closely parallel those studied in integrative adult cognition in the **Thrust 3**, but with simpler stimuli that are more readily studied in infants. Our study of infants' intuitive physics in the context of tracking and predicting the trajectories of objects in motion can ultimately be grounded in some of the best-understood neural circuits for statistical inference (**Thrust 2**) for motion analysis in the dorsal stream. Our sampling-based inference algorithms also have known potential neural mechanisms. Our emphasis on understanding how infants perceive goals and other mental states of intentional agents from observing their behavior parallels the nonverbal social perception side of **Thrust 4**.

4.b.1.4 Participants and their roles

Schulz, Spelke, and Saxe will do the empirical work (1.2.1). Spelke and Saxe will focus on intuitive physics and intuitive psychology in infants younger than 1 year. Schulz will focus on physical and psychological reasoning and their integration in infants between 1 and 2 years. Tenenbaum, Goodman, Mahadevan, and Ullman will build computational models (1.2.2). Goodman will develop probabilistic programming models of infants' physical and psychological reasoning. Mahadevan will study the models' theoretical properties and heuristics for inference. Ullman will integrate these models with visual cues and visual routines suitable for analyzing the stimuli infants observe. Tenenbaum will model learning and developmental transitions between stages of knowledge, and will coordinate all computational efforts. Metta and Sandini will implement these models in the ICub robotic testbed (1.2.3).

4.b.2 Thrust 2: ***Neuronal circuits underlying intelligence*** (Kreiman, Desimone, Boyden, Wilson, Sompolinsky, Freiwald, Koch, *Kanwisher, Valiant, Logothetis*)

4.b.2.1 Vision and goals

Abstract thinking and complex problem solving constitute paradigmatic examples of computation emerging from interconnected neuronal circuits. The biological hardware represents the output of millions of years of evolution leading to neuronal circuits that provide fast, efficient, and fault-tolerant solutions [19, 23, 75]. 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). Our team of theoreticians and neurophysiologists will focus on systematic, novel and integrative approaches to deciphering the neuronal circuits underlying intelligence. Understanding neuronal circuits that implement solutions to complex challenges is an essential part of scientific reductionism, leading to insights useful for developing intelligent machines.

4.b.2.2 Approach

We consider three paradigmatic examples of complex problems representing different aspects of intelligence to be investigated at the neurophysiological level:

2.2.1 Invariance in computations underlying recognition of objects and people. Recognizing objects and people poses important challenges for machines such as invariance to external transformations, robustness to internal sources of noise, rapid processing, fault-tolerance, and learning from few examples. Visual recognition requires the ability to detect and identify specific objects or people. The complexity of the problem arises from the infinite number of projections that the same object can cast on the retina [20, 27, 37, 38, 53, 71, 72, 76, 83, 85, 90, 92]. We will investigate the neuronal mechanisms underlying the theoretical framework described in **Enabling Theory**. We aim to describe how neuronal circuits provide invariance to external variations (including viewpoint, illumination, clutter, and occlusion) yet still learn from few examples. The theory and computational models will be empirically evaluated by neurophysiological experiments in macaque monkeys and humans.

2.2.2 Understanding interactions among people, objects and scenes. **Thrusts 3 and 4** emphasize that intelligent processing goes beyond detecting and identifying objects and people. As described in **Thrust 3** and illustrated in the example shown in **Figure 2.1,** parsing a scene requires the ability to understand how objects relate to each other, how to recognize actions, who implements those actions, and the context in which they take place. We will develop a neural model of the theory in **Thrust 3**, describing computational mechanisms for learning and recognizing associations, actions, and interactions. Simultaneously we will investigate how neuronal circuits in monkeys and humans represent information about people and objects. Building up from the efforts in **Thrust 3**, we will begin with images depicting interactions among two people, two objects, or people and objects, gradually increasing the image and task complexity to studying natural scenes [61] and tasks that rely on scene understanding.

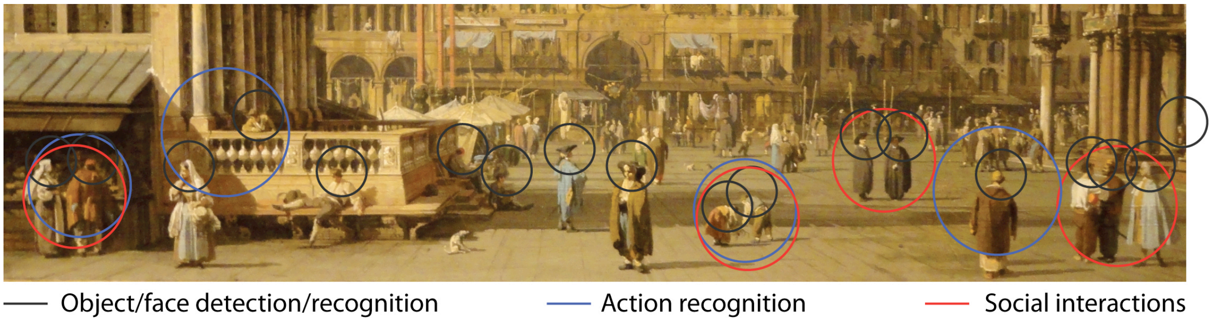

Figure 2.1. (Detail from Canaletto's The Clock Tower in Piazza San Marco) Example of complex and interacting tasks to required to parse an image including object recognition (black), action recognition (blue) and understanding goals and social interactions (red).

2.2.3 Interactions among brain areas. A critical aspect of intelligence, which proves difficult to implement in machines, is the capacity to integrate processes—interpreting sensory data, planning, decision-making, and motor output. Integration requires dynamic interactions across multiple brain areas [1, 12, 13, 22, 30, 46, 58, 84]. It is likely that *intelligence* and the ability to solve the complex problems illustrated above are not properties of individual brain areas [25, 28, 32, 47, 67, 68, 103]. Current theories suggest that these computations emerge from interactions among brain areas [48], a problem studied empirically in **Thrust 4**. It has been difficult to study interaction and integration between brain areas at the neuronal level and at millisecond resolution. We will develop the neuro-technology to build "functional microscopes" and tools to interrogate dynamical interactions in the brain: optogenetics, electrophysiology, transcranial magnetic stimulation, and functional imaging [10, 17, 18, 43, 89]. imaging. These tools will be applied to rodents, monkeys, and humans -- as dictated by the question and the technique. We plan to develop electrode arrays consisting of thousands of microwires to simultaneously monitor the activity of large neuronal ensembles in multiple brain areas [106]. This will enable us to extend the type of synergistic interactions possible between theory and experiments, typically limited by recordings from one or a few neurons. We will directly correlate the computational and theoretical efforts about neuronal network architecture and function with neuronal ensembles recordings while evaluating the type of tasks in 2.2.1 and 2.2.2 to provide direct mechanistic insights about how various abilities (e.g., vision, decision-making, motor output) are integrated.

4.b.2.3 Synergies with the other thrusts

Our theoretical and computational efforts are directly related to the main principles established by **Enabling Theory***.* *K*nowledge derived from the probabilistic inference models, learning theory, and ventral stream theory will be applied to synthesizing and understanding how neuronal circuits implement the computations required to support recognition tasks in 2.2.1 and 2.2.2. We anticipate a strong interaction with the efforts in **Thrust 1** in the characterization of learning processes (theoretically and neurophysiologically) because learning constitutes a central aspect of the biophysically plausible theories and computational models developed here. **Thrust 2** will directly benefit from and provide biological inspiration for **Thrust 3***.* We will provide insights about the underlying principles by which neurons integrate information in the context of recognizing parts, objects, and their interactions. **Thrust 2** will interact with **Thrust 4** given the importance of social interactions in 2.2.2. Across the four thrusts, we follow a specific example describing people and objects and their interactions. We will provide biological insights into the circuits and principles that govern these interactions in 2.2.2.

4.b.2.4 Participants and their roles

Valiant: theoretical constraints on computational primitives and related connectivity constraints. Sompolinsky: theoretical foundations of learning mechanisms at the level of individual neurons. Poggio: theoretical framework relating the computational goal of invariant perception to the architecture of cortex and cell tuning in the ventral stream. Kreiman: human neurophysiology and computational vision. Koch:

biophysics of computations in cortex. Boyden: neuro-technology, electrode arrays, optogenetics. Wilson: neuro-technology, rodent neurophysiology, interactions across brain areas. Freiwald: macaque neurophysiology, recognition of people and their actions. Logothetis: macaque neurophysiology and functional imaging, electrical stimulation, interactions among brain areas. Desimone: neurophysiology of scene understanding in macaque and interactions among brain areas. Kanwisher: functional imaging, transcranial magnetic stimulation, and interactions in the human brain. The collaborative research pursued here builds upon the expertise in each group but depart from the individual efforts in each laboratory – our goals involve multiple collaborations.

4.b.3 Thrust 3: Integrating intelligence: vision, language, and social interactions (Ullman, Winston, Yuille, Katz, Kaelbling, Poggio, Valiant, *Saxe, Tenenbaum, Buelthoff)*

4.b.3.1 Vision and Goals

The goal of this thrust is to combine vision with aspects of language and social cognition to obtain complex knowledge about the surrounding environment. Over the last decade, computational models have made significant progress in the task of recognizing hundreds of natural object categories under realistic viewing conditions [34, 65, 66, 76]. However, to obtain full understanding of visual scenes, computational models should be able to extract from the scene any meaningful information that a human observer can extract, about actions, agents, goals, object properties, scenes and object configurations, social interactions, and more. We refer to this 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 the image in a human-like manner.

The object domain goes ‘below’ and ‘above’ single objects i.e., the recognition of meaningful objects parts (door knob, zipper) and configurations (a table set for dinner). Agent domain queries include actions, goals, and interactions (hugging, quarreling). Understanding queries and formulating appropriate answers requires interactions between vision and natural language. Interpreting goals, and interactions requires connections between vision and social cognition. Answering queries also requires task-dependent processing, i.e., different visual processes achieve different goals. These problems can are divided into sub-tasks described below.

4.b.3.2 Approach

3.2.1 Objects, parts, and configurations. The goal is to exceed detecting and recognizing objects by going ‘below’ the object level to analyze meaningful parts and by going ‘above’ to analyze meaningful configurations of multiple objects. Object parts are often defined by their use rather than appearance, and detection may rely on context rather than appearance, raising major challenges to learning and interpretation. Our approach to learning semantic parts (e.g., door knob, cup-handle) combines part-learning with analysis of agent-object interactions. For example, the location of an agent’s hand when interacting with a door guides the learning of the typical location, shape, and use of a doorknob. Understanding meaningful object configurations raises similar challenges regarding the understanding of meaning that goes beyond object appearance. For example, a table set for dinner and a table after dinner includes similar objects but in different configurations. Meaningful object configurations are variable and recognition requires understanding goals and use. We approach this problem by first recognizing participating objects and their spatial relations, and applying the extended interpretation processes discussed below. This part will be coordinated with **Thrust 2** (part 2.2.1) concerning neuronal aspects of recognizing objects and people.

3.2.2 Understanding goals and interactions*.* The aim here is to recognize actions performed by agents, their goals, and interactions between agents. Understanding actions and goals requires the recognition and analysis of agents and their body parts, object parts and shapes, and agent-object interactions [39, 70]. For example, human-level recognition of drinking requires the localization and analysis of the agent’s hands, face, mouth, the object, whether it is opened from above, type of grasp, and relative location of the object. Our approach is to learn and execute extended routines, which might start by detecting body parts, localize the hands and held object, its shape and relationship to the body, face, and mouth. Learning the actions will be based on dynamic images, and recognition will be obtained from both static and dynamic input. The goal for recognizing the agents’ interactions is to recognize actions such as hugging, quarreling, or dog-walking. Little is known about the processes and the information used to perceive agents’ interactions. Psychophysical studies will help identify cues used by humans combined with computational models of the processes that use them to make inferences about agents' interactions, by combining visual and non-visual information. This part will be coordinated with the study of neuronal mechanisms of agents and agent-object interactions in ***Thrust 2*** (part 2.2.2) and **Thrust 4** (4.2.2).

3.2.3 Turing vision test. The goal is to answer a large, flexible set of queries about images in the domains discussed above, e.g., where is the door-knob, pocket, or side-view mirror, which object is closest to the window, what is the boy holding, who is touching the chair, what is the woman looking at. To answer appropriately, the query must be understood, a suitable visual process must be applied, and a correct answer must be composed [49].

3.2.4 Extended interpretation and goal-directed processing. To achieve these goals we must develop novel methods for extracting meaningful information from images. Semantic image interpretation often requires an extended process, directed to specific objects and relations, e.g., what is person X looking at or touching, is object Y stable. The model of the interpretation will be composed of multiple steps applied in a goal-dependent manner, i.e., embedded inference. This method combines probabilistic inference with policy learning to generate a sequence of operations applied to the image. The first stage constructs the most probable interpretation of the scene, and the second generates and applies an interpretation in a task dependent manner [93]; different processes can be synthesized in response to different queries. Current probabilistic inference methods will be extended to produce rich hierarchical representations [6] for the first components, and to deal with robustness and invariance. The second stage, generating extended goal directed processing, requires policy learning related to Markov Decision Processes and reinforcement learning.

4.b.3.3 Integration with other areas

Close interactions with **Thrust 4** involve understanding actions, goals, and agents' interactions. Close interactions with **Thrust 1** involve incorporating useful structures and biases derived from human developmental cognition. Hierarchical object recognition connects with **Thrust 2** for modeling cortical mechanisms of hierarchical representations and object recognition and connecting the computational constraints characterized in **Thrust 3** with neuronal circuits. Computational theories will be used to design stimuli for testing neuronal responses to objects configurations, agent-object interactions, and agents interactions (parts 2.2.2 and 2.2.3), and for predicting, testing, and analyzing interactions among brain areas (using developments in 2.2.3) in tasks that require integrating vision and social cognition (**Thrust 4**). Visual aspects of social interactions connect with **Thrust 4**, parts A and B. The focus in **Thrust 3** will be on what vision can deliver and how, and **Thrust 4** will focus on how representations of social knowledge use visual information to make social inferences. All projects will engage **Enabling Theory** concerning invariant recognition and probabilistic modeling and inference.

4.b.3.4 Participants and their roles

Ullman: objects and parts, action recognition, agents' interactions. Poggio: learning, hierarchical object representation, invariances, and relation to cortical models. Yuille: objects, parts and configurations, action recognition. Kaelbling: extended sequential processing by eye movements and internal processing, planning and vision. Winston: combining language and vision, using visual models to understand natural language queries and generate an appropriate response. Katz: understanding natural language queries, generating natural language answers.

4.b.4 Research Thrust 4: *Social intelligence* (Kanwisher, Saxe, Nakayama, Spelke)

4.b.4.1 Vision and Goals

Social cognition is at the core of human intelligence. It is through social interactions that we learn, and it is social interactions that drove much of the evolution of the human brain. Indeed, the neural machinery of social cognition comprises a substantial proportion of the brain. Finally, the greatest feats of the human intellect are often the product not of individual brains, but rather of groups of people working together in social groups. Thus, intelligence simply cannot be understood without understanding social cognition. Yet we have no taxonomy or theory of social intelligence, and little understanding of brain mechanisms, their development, or their underlying computations. Here we bring developmental, computational, and cognitive neuroscience to bear on a newly tractable component of social intelligence: nonverbal social perception (NVSP).

NVSP is the ability to discern multidimensional social information from a few seconds of a silent video. Human subjects can tell whether a person is happy or sad, confident or tentative, thoughtful or bored, attentive or disengaged, and expectant or resigned. They can infer what the person is doing and why they are doing it. If the video contains two people they can discern their relationship (boss/employee, parent/child, brother/sister), and whether they are interacting, or are engaged in unrelated activities, even if they don't address one another.

This project aims to characterize NVSP and the stimulus information that makes it possible, as well as the computational, developmental, and neural basis of this remarkable ability. Each question is addressed through a set of component projects, with different methods, leading synergistically to a theory of the mental, neural, and computational processes that enable us to represent the world in vivid social Technicolor.

4.b.4.2 Approach

4.2.1 The Psychophysics of Social Perception**.** This project uses psychophysical methods to characterize adult NVSP. The pioneering work of Nalini Ambady provides a rich natural history of NVSP, detailing myriad social judgments made from brief silent movie clips [105]. We will build on this foundation to provide a precise, quantitative, and analytical characterization of each perceptual ability. Central to this project is the design of stimuli and tasks that tap key components of social perception, which will be used in subsequent computational analyses (4.2.2), development studies (4.2.3), and brain-based studies (4.2.4).

Stimuli will be short video clips of real people either alone or interacting. Judgments will include the: presence or intentions of an agent, presence and nature of an interaction, nature of the relationship (e.g., status/age/power), and group membership structure of three or more agents. To analyze the perceptual cues supporting each social inference, subjects will be tested on altered versions of the original video in which a principled subset of the stimulus information has been isolated or filtered (e.g., still snapshots, videos played backwards, scrambled, or reduced to "point-light" displays). We will also consult with character animators for their insights about the most powerful cues in nonverbal social perception, which will be manipulated in character animation movies that manipulate these cues.

4.2.2 Computational Theory of Social Perception**.** How can social perception be understood in computational terms – what are the representations and algorithms that enable such rich information to be extracted so quickly from short image sequences? We hypothesize that core machinery is shared with infants' intuitive psychology, and characterize adults' social perception using the same basic approach of Bayesian inference over probabilistic models of intentional agents as boundedly rational planners (see **Thrust 1**). In preliminary studies with highly simplified stimuli, these models have already successfully predicted adults' judgments of goals, beliefs, preferences, and social orientation [4, 94]. This project extends these models to work with more complex, naturalistic, perceptual input, drawing in part on computational models for integrated visual scene understanding (**Thrust 3**), and to capture richer aspects of social perception not yet present in infancy. We will also test the predictive power of probabilistic planning-based models against simpler alternative models based solely on bottom-up cues, and study how these two approaches may be integrated to explain how social perception can be both so rich and so quick.

4.2.3 The Development of Social Perception. How does nonverbal social perception develop in childhood? A new study by Saxe and Kanwisher shows protracted development of the ability to perceive whether an individual is engaged in a social interaction, continuing from age 5 to 10 years [5]. In contrast, basic face discrimination abilities do not show domain-specific development after age 4 [55, 104]. The psychophysical tasks developed in Part 4.2.1 will be adapted for children from the youngest age possible through the early teens. The development of distinct NVSP abilities along different trajectories would complement other methods (e.g., fMRI) of discovering the functional components of nonverbal social perception.

4.2.4a Standard fMRI Studies of Social Perception in Adults**.** The superior temporal sulcus (STS) plays a special role in nonverbal social perception. STS subregions have been implicated in the perception of dynamic (not static) faces, voices, biological motion, intentional action, social interactions, and theory of mind. Little is known about the precise relationship between these functions. Saxe and Kanwisher will collaborate on imaging experiments focused on the functional organization of the STS for social perception and social cognition, using the stimuli and tasks devised in part 4.2.1. Some studies will apply standard fMRI methods to test aspects of NVSP thought to be distinct based on data from parts 4.2.1 - 4.2.3; others will apply newer data-driven methods (e.g., clustering) to discover structure in the brain's response to a large sets of movie clips.

4.2.4b Other brain-based methods. To watch early development of brain regions engaged in NVSP, Saxe and Kanwisher will scan 6-month-old human infants viewing social movies in an improved version of our specialized infant coil/recliner designed for high signal to noise ratio. Social perceptual inferences are largely automatic so short scans (of around ten minutes) may suffice. TMS studies in adults will test the causal role of each STS subregion in NVSP and related abilities. Diffusion-weighted imaging will identify connections between distinct STS subregions, and with the rest of the brain. Social perceptual abilities in macaques rival humans [78], affording the possibility of fMRI studies on NVSP in macaques.

4.b.4.3 Integration with other areas

Connections to **Trust 1** include the computational analyses of nonverbal social perception in infants and adults, and how these competences develop. **Thrust 2** work with social stimuli (faces) connects with the human fMRI and macaque models of nonverbal social perception. High-level perceptual representations examined in **Thrust 3**, and models of the social and physical world, are closely connected to this thrust.

4.b.4.4 Participants and their roles

Kanwisher will be involved in all projects and their tight integration. Nakayama will lead the psychophysics project (4.2.1) jointly with Kanwisher. Tenenbaum will lead the computational modeling project (4.2.2). Saxe will work closely with Kanwisher on the developmental project (4.2.3) and the brain-imaging project (4.2.3). Spelke will collaborate on the developmental project (4.2.3).

4.b.5 A common platform: Enabling Theory (Poggio, Goodman, Tenenbaum, Yuille, Sompolinsky, Verri, Rosasco, Smale, Shashua, Valiant, Ullman, Mahadevan, Koch, Hirsh)

Understanding intelligence and the brain requires theories at different levels (see afterword and foreword in [54]), ranging from the biophysics of single neurons, to algorithms and circuits, to overall computations and behavior, and to a theory of learning. In the past few decades, advances have been made in multiple areas from multiple perspectives including: PAC learning, connectionism, statistical learning theory, machine learning, probabilistic inference, and the biophysics of computation. Several of the main contributors to these advances are members of our team.

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.

A core *mathematics of intelligence* comprising *learning*, *inference,* and *neural computation has emerged from these advances*. The goal of the **Enabling Theory** platform is to extend and formalize this theoretical framework and to use it as a unifying language across the Center's thrusts to guide, analyze and interpret experiments. We list here some of planned work in each of the three main directions of theoretical research:

Learning theory is the modern synthesis (due to Vapnik, Valiant, and Smale among others) of diverse fields in modern mathematics such as high dimensional probability and empirical process theory, computational harmonic analysis, computational geometry and topology, optimization theory and convex analysis [2, 9, 21, 26, 63, 77, 80, 82, 96–99]. Hierarchical “deep” architectures for learning represents a promising area for theoretical work leading to a new learning theory inspired by the basic organization of the cortex. We plan to organize and extend recent work to a comprehensive kernel-based theory ranging from supervised learning to a broader range of tasks including semi-supervised and unsupervised learning, multiclass categorization, clustering, and taxonomy learning. This extension will include statistical approaches in a time dependent manner with varying parameters to understand robust learning and adaptation to a varying environment.

Probabilistic modeling and inference are central tools for acting intelligently in a complex world with pervasive uncertainty. Probabilistic graphical models are our starting point, casting perception, reasoning, learning, prediction, and planning in a unified framework as Bayesian inferences about unobserved variables (latent causes or future outcomes) conditioned on observed data (effects). Hierarchical and nonparametric Bayesian methods extend this approach to explain how priors may be learned from experience over longer time scales and broader domains of tasks, and how learners can grow models over a lifetime that balance simplicity and data fit, introducing new degrees of freedom only as needed based on experience[86]. Probabilistic grammars allow uncertain reasoning for structured interpretations of scenes and sentences[51]. Probabilistic programs generalize all these methods, marrying Bayesian probability with universal computation[40]. Instead of modeling joint distributions over a set of random variables, as in traditional graphical models, probabilistic programs model distributions over the execution histories of programs, including programs that analyze, transform and write other programs. This provides the expressiveness to capture core abstract concepts (e.g., analogs of force, mass; utility, belief) and the capacity for recursive and reflective reasoning (e.g., beliefs about beliefs) needed for intuitive physics and psychology in infants, and integrative scene understanding, language and social perception in adults. We plan to extend and unify these methods under the Church probabilistic programming language [40], to develop new fast approximate inference schemes, and to study the integration of inference and learning, to explain why probabilistic reasoning works so well only in domains where humans have appropriate experience.

Neural Computation comprises several complementary modeling approaches that have been developed to link intelligent behavior and the brain mechanisms underlying it [69]. Work is planned on neural circuits that may implement probabilistic inference (including representations of constraints and priors) [7, 8, 14]. We will also investigate a recent theory that attempts to explain and predict cortical architecture and properties of neurons in different visual areas [64]. The starting assumptions are: a) The main computational goal of the ventral stream is to learn-and-discount image transformations, and b) Hebbian-like plasticity occurs during the development of sensory cortex [15, 33, 50]. The theory shows how the hierarchical architecture of visual cortex as well as properties of neurons in different visual areas follow from these two assumptions as a consequence of symmetry properties of the geometry of affine transformations -- very much in the spirit of modern physics.

Participants and their roles

Several members of the team play key roles in the theory platform. Goodman, Tenenbaum and Yuille: Theories of probabilistic reasoning and inference for language, social intelligence, for vision. Sompolinsky: connecting Bayesian models to neurons. Verri, Rosasco, Smale, and Shashua: Techniques for supervised and unsupervised learning. Valiant: Connecting cortical connectivity to computational primitives needed for the various aspects of intelligence, i.e., can the same basic machinery support computations ranging from vision to language and motor control? Ullman, Mahadevan, Koch and Poggio: Theory of the ventral stream -- feed forward and feedback -- including attention and visual routines.

4.c Education and Human Resource Plan

The goal of CBMM's education program is to provide the next generation of researchers with unprecedented interdisciplinary and integrated training in brain, cognitive, and computational sciences. *At the graduate level* we will: 1) Develop coursework that gives students from diverse backgrounds a shared interdisciplinary underpinning for studying intelligence; 2) Design a teaching consortium that offers courses across participating institutions that provide integrated knowledge and that leverage curricular offerings that are more expansive than what is available currently; and 3) Provide hands-on experience in interdisciplinary research by participation in center research and in internships across institutions and industrial partners. *At the undergraduate level* we will: 1) Develop interdisciplinary curricula that prepares students for graduate study with a broad background; and 2) provide advanced undergraduates with experience in interdisciplinary research on intelligence at CBMM institutions. At both levels we plan to lead by example, producing model curricula and learning materials that can be adopted by other institutions, connecting with Rutgers efforts in the NSF-funded IGERT program in Perceptual Science, the I3 award on "Institutionalizing the IGERT Innovations at Rutgers," and the recent REU Site on Perceptual Science and Technology. We will aggressively pursue our educational activities to exploit the rapidly evolving educational technologies while exploring new modes of instruction, pedagogy, and assessment.

Graduate Education

Course Development. Students new to CBMM will be offered three new courses to provide a shared interdisciplinary foundation for studying intelligence:

- What Is Intelligence? This broad introduction to the science of intelligence, to be taken in a student's first semester, provides a comprehensive overview several areas of study. Lectures will be given by center faculty, providing an interdisciplinary view of each area spanning from computational models to the neural circuits. An initial version of this course was offered at MIT this past fall, by Shimon Ullman and Tomaso Poggio, which was taught through lectures by many of the investigators of this project. This course provides the "big picture" of intelligence from a research perspective to frame students' graduate research experience, and to help them meet and select the faculty with whom they will undertake their research.
- Computational Cognition: This first year course provides in-depth experience in the computational and probabilistic modeling approach to cognitive science, in which the elements of intelligence are understood as inference of and in complex probabilistic models. Examples will be drawn from concept learning, causal reasoning, social cognition, and language understanding. Formal modeling ideas and techniques are discussed in concert with relevant empirical phenomena. Initial versions of this course, grounded throughout with software in the Church programming language for probabilistic modeling, have been offered by Josh Tenenbaum at MIT and Noah Goodman at Stanford; they will co-teach the course for CBMM to make it available more broadly as discussed further below.
- Computational Neuroscience: This first year course provides in-depth experience in the computational and probabilistic modeling approach to neuroscience, in which neural coding and dynamics are understood mathematically and computationally. Formal modeling ideas and techniques are discussed in tandem with relevant empirical phenomena. Versions of this class are already offered at MIT by Michael Fee, Matt Wilson, and others, and at Harvard by Haim Sompolinsky.

We also plan to adapt existing graduate courses on intelligence to heighten their interdisciplinary nature, focusing on how existing courses intersect with intelligence in neuronal, cognitive, and computational ways. We will make the learning materials from these courses available online, as further below, so that instructors at other institutions might adapt their own courses for greater interdisciplinary coverage. In the first year we plan to offer CBMM students updated versions of *What Is Intelligence, Computational Neuroscience*, and *Computational Cognition*, the last offered jointly by Goodman and Tenenbaum. We will add Statistical Learning Theory and Applications (Poggio), Visual Object Recognition (Kreiman), and Collective Intelligence (Hirsh) in year two. Each of these courses will be revised so they can be taught at one institution and offered remotely to students at other CBMM institutions, and later for an Internet-based course.

All CBMM graduate courses will be offered as part of a Graduate Teaching Consortium (GTC) that offers a selection of courses to advance integrative interdisciplinary training of graduate students at participating institutions. The goals of the GTC are to: 1) Offer students a broad selection of advanced graduate courses beyond those available at the local institutions; 2) Expose GTC students to world experts on these subjects; 3) Establish new models in cross-university education collaboration; and 4) Broaden the impact of CBMM's educational efforts.

Each GTC course will be offered by one institution but made accessible in "real time" to all participating institutions. Each course will have a lead instructor at that institution who is responsible for all lectures, problem sets, exams, and grading for all registered students at participating institutions. Our goal is to have these courses offered at one institution, and also offered to all members of the GTC for academic credit. Their design will include consideration of curricular sensibilities and course approval needs at CBMM's multiple institutions. All courses will be videotaped for subsequent use and subject to careful assessment, as discussed further below.

Model Curriculum. CBMM graduate students will be expected to follow a course of activities that trains them broadly, coupled with hands-on, diverse research experiences. Because each institution has different course offerings and degrees offered, we cannot dictate a single curriculum for all students. Instead, the initial year is a template educational program that can guide all member institutions – as well as other institutions that may subsequently adopt them – on the interdisciplinary training of graduate students. This template will be the basis for evaluating the impact of the STC on the curriculum at participating institutions. For example, regardless of their primary research students should have coursework in neuroscience; gain experience in probabilistic inference; learn about computational models in multiple areas of intelligence; and take a diverse set of courses involving cognition. Students will be expected to conduct two research projects with different advisors to gain first-hand experience in different research areas related to the student's thesis research. One of the projects should be connected to the final dissertation research; and the other in a second research area, either with a second faculty member or as a relevant industrial internship. Funds will be provided to offer students the opportunity to conduct their second research project at a second CBMM institution over a six-month stay.

Undergraduate Education

Model Curriculum. A well-designed undergraduate program should have sufficient scope to allow a student's research to flow in the most appropriate directions. Any list of topics we might consider for inclusion in such a program (e.g., linguistics, perception, developmental biology, neurophysiology, genetics, probability and statistics, computer programming) will undoubtedly be larger than can fit in a student's college stay and yet some relevant intellectual strands will be left out. Yet we can establish some expectations for a student targeting an eventual graduate program in intelligence studies. Because every institution is different, we cannot expect to design a specific course of study for all. Instead, in our second year (based on our experience in the first year at the graduate level) we will design a curriculum template that can frame our thinking about undergraduate study in the intelligence sciences, interacting especially with our partnered minority serving institutions as well as our Education Advisory Board (discussed below). This curriculum, for example, would ensure that students receive at least some treatment of neuroscience; some experience with probability and statistics; some computer programming coursework; plus a suitably diverse set of courses involving cognition.

Course and Module Development. While our model curriculum will focus on key intellectual threads, we also plan to offer a new undergraduate course in year 3, "What is Intelligence" modeled after our new graduate course on the topic, to bring them together for the students. The course will provide undergraduate students an overview of our understanding of intelligence, from neural substrates to computational models. Our goal is to remain vigilant in transitioning intellectual material from graduate coursework to the undergraduate level when the content makes it possible.

Many of the courses that we would like to see an undergraduate take are not specific to intelligence. A math professor may teach a course on probability without ever connecting it to the study of intelligence; similarly, a course on programming may never connect graph algorithms to questions studied in computational intelligence. We plan to develop modules and technical support designed to make it easy for the instructor of a course not specifically centered on intelligence to include intelligence-centered content. For example, an instructor of probability, or of game theory, or of introductory programming, would be able to use modules developed by the center that would give one or two lectures' of slides, course assignments, and exam questions, making it easier for an instructor to connect the course topic to the study of intelligence. These materials will be designed to exploit the affordances that the Internet and computing technology may provide.

Undergraduate Research and Student Leadership. Each of the three core institutions has well-established programs for undergraduate research. Particularly distinctive is MIT's UROP (Undergraduate Research Opportunities Program), offering a diverse range of research experiences well-integrated into the undergraduate experience. We will utilize the UROP resource for those students at MIT, and plan to use lessons learned from the UROP program in the design of UROP-like summer REU research projects at each institution. Our goal is to engage junior-level students in such activities so that in their senior year, after their summer research experience, they could develop their leadership skills by actively organizing seminar series and mentoring junior students.

Postdoctoral Researchers. Several postdoctoral researchers will be supported by the Center to conduct research aimed at understanding several aspects of intelligent behavior. Our Postdoctoral Researchers Mentoring Plan describes our flexible framework for supporting individual postdocs' professional and career development needs.

Technology and Education. Our goal is to pursue aggressively the use of technology in CBMM education. For the purposes of students within CBMM who may be at one institution and taking a course at a second, we plan to instrument classrooms so that students in different locations feel like fellow students. Video-wall displays in specially instrumented classrooms will be installed in year one for students to see and hear each other in classrooms that go beyond Skype-like videoconferences. These classrooms would also be used for research group meetings for CBMM.

We will design our courses so they can be adapted to Internet scales, such as seen in Norvig and Thrun's introductory AI course this past fall, which offered an educational experience to tens of thousands of students. In particular, we plan to develop -- collaboratively with instructors at our partner institutions, particularly those at minority serving colleges -- Internet-based courses, which could be offered to tens of thousands of students worldwide with infrastructures such as the new MITx platform. The goal is not to replace CBMM classroom teaching with technology, but to extend the reach of our educational efforts. The course development described above will follow a two-staged approach, with classroom lectures designed in separable small chunks so that they can be videotaped and subsequently augmented with gatekeeper questions and exercises, coupled with social technologies for student discussion and question answering, in service of our online offerings. These online courses will be designed collaboratively with instructors at our partner institutions, particularly those at minority serving colleges, to design online offerings that scale to all our partners and is accessible to suitably motivated individuals. Further discussion of Internet-based courses can be found in the next section.

We will use our curricular offerings to advance our understanding of the use computing technologies for teaching and learning in this domain. We will use modern data analysis tools on the ongoing data that we generate about our educational programs (discussed further below). We will develop learning materials, particularly module development, via a range of stakeholders, such as faculty, students, paid workers, and volunteers, so as to grow our understanding of ways to do so [52].

Ethics. All students and personnel will be required to participate in instruction on the Responsible Conduct of Research and IP Rights, as discussed further in the Ethics Plan.

Evaluation. CBMM's Evaluation Team, led by external evaluator Lizanne DeStefano, College of Education, UIUC, will administer surveys and other assessments, and conduct personal interviews and focus groups to collect data and track CBMM's learning outcomes over time. DeStefano plays a central role for assessment in an existing STC (EBICS), which will allow us to adopt lessons learned there to CBMM. The evaluation will utilize an educative, value-engaged approach (EVEN). The EVEN approach, developed with NSF-EHR support, defines high quality STEM educational programming as that which effectively incorporates cutting-edge scientific content, strong instructional pedagogy, and sensitivity to diversity and equity issues. This evaluation involves:

- Implementation: Are CBMM's education programs being implemented on schedule and as planned?
- Effectiveness: Are CBMM's education programs operating effectively? How might they be improved?
- Impact: What is the added value that CBMM uniquely provides?
- Institutionalization: How and to what extent are CBMM's activities and educational products becoming institutionalized at other participating institutions? What opportunities and barriers exist to broad institutionalization?

Evaluation will focus on four high priority areas: 1) Efficacy and impact of the Graduate Teaching Consortium; 2) Efficacy and impact of our technologically facilitated course offerings; 3) long-term impacts of the trainee experience; and 4) Dissemination and efficacy of CBMM's educational programs.

Graduate Teaching Consortium. Because the GTC seeks to develop a new program for inter-institutional study, we are particularly interested in accessing the student GTC experience. Evaluation data collected from student surveys of every GTC course, in addition to a yearly internal assessment of course content and relevance to the goals of CBMM, will be used to inform decisions about curriculum, course content and materials, teaching strategies, and communication media. The student survey results in combination with observations, document review, and faculty course surveys and interviews will be used to (1) inform the design of subsequent courses that will be offered, and (2) develop GTC best practices guidelines for faculty and teaching assistants.

Technologically Facilitated Courses. As CBMM researchers aggressively exploit the use of technology in CBMM education, the evaluation team will seek to develop accompanying evaluation tools and processes to support teaching and learning at the Internet scale. The evaluation will document and compare student experiences and outcomes in face-to-face and remote sites to identify effective on-line instructional strategies, measure increased access and diversity attributable to CBMM efforts, and document the sustainability and scalability of the model technology. Both the courses and the evaluation system will be products of this activity.

Longitudinal Study. Because one CBMM objective is to prepare students as future researchers and education leaders, another primary interest is the long-term impact of the programs, scientific training, and educational experiences. The Center will assess impact by conducting a cross-institutional Longitudinal Study of trainee experiences and outcomes (i.e., courses taken, research experiences, advisement, retention, graduation, awards, publications, and initial employment). This includes tracking REU students, undergraduates, graduate students, and postdoctoral associates throughout their involvement in the Center and beyond. We are interested in trainee career trajectories and the long-term retention of women and under represented minorities in STEM careers. We plan to use social media tools such as Linkedin to maintain professional connections with our trainees as they move on from the Center. Our approach will be modeled after the evaluation program at the EBICS STC, and, for example, we plan to adapt and adopt their evaluation templates and activities tracking software for CBMM.

Educational Programs. The Evaluation Team will also track development, use, dissemination, and efficacy of our educational materials by others. We will use short "satisfaction" surveys to evaluate participant experiences and to collect information on the target groups we reach with our programs.

Education Advisory Committee. CBMM will assemble an Education Advisory Committee to be a resource in our educational aspirations. It will be comprised of domain experts and educators as well as experts on education in such areas as online learning and at institutions that face severe budgetary constraints on education.

Degree	Total #	# Women	# African American	# Hispanic	# Native American	# with Disabilities	Average Time to Completion
MS	35	12	0	2	0	0	2 years
PhD	46	12	1	1	0	0	5.1 years

Country of origin: United States (39); Brazil (1); Canada (1), China (3), Cypress (1), Great Britain (1), France (1), Germany (2), India (2), Indonesia (1), Iraq (1), Israel (8), Italy (1), Mexico (1), Switzerland (1), Taiwan (1), Turkey (2)

4.d. Diversity Objectives

Women and Minority-serving Institutions

The goal of our diversity and outreach activities is to attract more women and minorities into the complex, interdisciplinary field of cognitive neuroscience and artificial intelligence.

MIT has had a long-term commitment to increasing women and minorities in science, engineering, and technology and has been very public about the challenges of recruiting and retaining minorities and women, especially faculty (45.3% of undergraduates, 32% of graduate students, and 12.5 % of faculty are women). 20% of MIT's undergraduates are underrepresented minorities, and there are more women in Brain and Cognitive Sciences than in Computer Science and Artificial intelligence; although only 6% of MIT's graduate students are underrepresented minorities. Rutgers also has long-term commitment to underrepresented populations in science. Douglass College is a women's residential college at Rutgers founded in 1911; it is the largest women's college in the US. Undergraduate resources include the Bunting Program for non-traditional students and the Douglass Project for Women in Math, Science, and Engineering. The Rutgers Newark campus is located in a highly urban setting, and has been rated the most diverse college in the US by both Forbes and US News (the latter for 14 straight years). Rutgers also has a very high proportion of first-generation scholars. These are student populations that we would have access to for our diversity efforts.

The following four partner institutions have agreed to participate in our diversity outreach efforts: Hunter College, Howard University, the University of Puerto Rico, and Wellesley College (via letters from Vita Rabinowitz, provost of Hunter college, Jose Lasalde, VP of research of the University of Puerto Rico, and Kevin Chamness director of Sponsored Research at Wellesley college). This diverse community includes women, under-represented minorities (particularly African Americans and Hispanics), first generation college students, students with disabilities, and students on financial aid. We believe that diversity can best be achieved by exposing undergraduate students to scientific disciplines early and with strong mentorship. Students and faculty at partnering institutions will be offered numerous opportunities for collaborative research, academic enrichment, and networking, as well as summer research internships for students, workshops, and access to novel and online curricula. We expect faculty from partner institutions to play an active role in the development and teaching new courses, and in training and mentoring women and minority students in research.

Hunter College, an urban public college located in central Manhattan, is within commuting distance to MIT, Harvard and Rutgers. It is the largest college in the City University of New York (CUNY) system, and is known for the diversity of its student body, educational opportunities for women and minorities, and non-traditional students. Most students work, many are parents, and more than one third are the first generation college students. This partnership is a bridge to the entire CUNY population. Faculty contacts include Professors Susan Epstein and Virginia Teller. Dr. Epstein conducts research in machine learning and cognitive modeling, and received Hunter's first award for excellence in research and undergraduate teaching. A past chair of the Cognitive Science Society, she is known for her work in interdisciplinary collaboration [31], and has active ongoing work with roboticists, linguists, and bioinformaticians. Dr. Teller, the Chair of the Department of Computer Science, specializes in natural language processing, and artificial intelligence, and is known for recruiting and retaining under-represented groups in the computer sciences. Both teach undergraduate and graduate courses, and are eager to work with CBMM to develop and implement novel curricula at Hunter, and to strengthen and develop research collaborations with faculty at MIT, Harvard, and Rutgers. Co-PI Dr. Haym Hirsh, our Education coordinator, is a long-time colleague of Dr. Susan Epstein; Dr. Mandana Sassanfar, our Diversity coordinator, has been a regular visitor to the Minority Access to Research (MARC) and Minority Biomedical Research Support (*MARC* / MBRS) programs for the past decade.

Howard University is a private historically black college located in the heart of Washington D.C. Professor Poggio, the director of the proposed CBMM, has recently met with Howard's provost and associate provost to discuss the scope of the center and potential collaboration in machine learning. Dr. Sassanfar has strong ties to the science faculty at Howard and is currently working with several faculty members at Howard to establish collaborations with MIT. Faculty collaborators for the proposed center include Professor Mohamed Chouikha, chair of the Department of Electrical and Computer Engineering. His research interests include Machine Learning, and intelligent control and he has mentored many undergraduate and graduate students. He will develop and implement new curricular material and hands-on learning. Prof. Mugizi Rwebangira, a graduate from Carnegie Melon, is another expert in machine learning at Howard and will be an active collaborator within CBMM.

The University of Puerto Rico is a Hispanic public research university with 11 campuses. The largest campus is in Rio Piedras, the metropolitan area of San Juan. UPR-RP is the top baccalaureate source for Hispanics who go on to earn PhDs in the Sciences, and ties as the fourth highest source of Hispanics PhDs (NSF 1999-2003). Many faculty members are able and willing to contribute to the development of programs that will increase diversity in Neuroscience. Prof Sandra Peña de Ortiz, and Prof. Irving Vega from the biology department at the Rio Piedras campus, Guillermo Bernal, director and Professor in the dept of Psychology, and Prof. Maria Bykhovskaia, Chair of the Department of Neuroscience at the Universidad Central del Caribe seek partnerships with research-intensive institutions to provide cutting-edge research opportunities for their students. Dr. Vega is director of the Neuro-ID program. Dr. Bykhovskaia has an on-going collaboration with Troy Littleton at MIT, and will work with Dr. Sassanfar to identify suitable students for the summer research internships and for workshops.

Wellesley College is a highly selective private college for women in Eastern Massachusetts, with closely connected Computer Science and Neuroscience programs. Nearly half its students are from underrepresented minorities. Including Wellesley as an outreach partner serves two goals: 1) Attract more women to a currently male dominated field, and 2) Help strengthen the undergraduate curriculum at our other partner institutions. Wellesley already has a strong educational relationship with MIT. Ellen Hildreth, Professor and Chair of the Computer Science department and member of the Neuroscience Advisory Committee, is a long-time colleague of Dr. Poggio. Her research on human visual processing integrates computational models and visual psychophysics. Other faculty collaborators include Bevil Conway, who studies the neural basis of color using physiology, behavior, and modeling, and Mike Wiest, who studies perception and attention using chronic multi-electrode recordings in behaving rats. Profs. Hildreth, Conway, and Wiest teach interdisciplinary courses in the Computer Science and Neuroscience programs and hold regular joint meetings of their research groups.

We will encourage diversity in three specific ways that focus on research and education: summer research opportunities for rising juniors and seniors, mini-sabbaticals for faculty from minority-serving institutions, and several on-going workshops and seminars.

Research opportunities for Undergraduates

Each institution will reserve 8 spots for 8-10 week undergraduate summer research internships in the laboratories of faculty involved in the center. (We will allocate $60,000 per year each for Harvard, MIT, and Rutgers to cover stipends, housing, travel and academic enrichment activities.) Priority will be given to students from our four partner institutions. All summer students will meet for a joint poster session at the end of the summer. The centers key faculty and at least one faculty member from each outreach partner institutions will attend the poster session. Because quality mentorship is a key component of the summer research experience, students will be carefully matched with research mentors. Students will have access to cutting edge research facilities, and many opportunities to interact with graduate students and faculty in academic and social settings.

MIT, Harvard, and Rutgers will host three faculty members from our outreach partner institutions or other minority-serving institutions for a 10-week **Summer Sabbatical** to help explore opportunities for future research collaboration. We have allocated $90,000 to fund three faculty members each summer. The funds will be used to complement the faculty's summer salary, provide housing and travel expenses, and support one graduate or undergraduate student from the visiting professor's laboratory.

5-day Workshops in neuroscience, cognition, and machine learning will be offered by MIT, Harvard, and Rutgers, who will take turns organizing and hosting them. The workshops will reserve slots for up to 5 students and 2 faculty members from each outreach partner institution. The workshops will take place early in January or late in May to avoid interference with the regular academic year and the summer research internships. The workshops will include faculty lectures, and hands-on activities in each of the three areas of the center led by graduate students and postdoctoral fellows, offering them the opportunity to teach. We have allocated $45,000 per year for lodging and travel.

The center will organize a **Lecture Series** that targets undergraduate audiences at minority-serving institutions. These lectures will be videotaped in front of a live audience and posted on the Internet for wide access. The series will contain 6 to 8 annual lectures for a cost of $18,000 per year (to cover travel expenses, videotaping, and a small honorarium). Lecturers will be chosen for pedagogical skill and knowledge.

The main strength of the diversity program will be to promote strong working relationships between the outreach partners on **novel curriculum development** to be made broadly available to all undergraduate institutions. Faculty at the four partner institutions will work together with faculty at MIT, Harvard, and Rutgers to develop a robust undergraduate curriculum applicable to institutions with limited resources. Each faculty member from our partner institutions will bring the skills, expertise, and ideas to help create a productive partnership and effective outcome.

Because **good communication** among the various partners will be essential, we plan to have an annual meeting with key faculty from each institution to review CBMM progress and to discuss new approaches and strategies to help the center reach its diversity goals. In addition, we will hold videoconferences to include a larger group of faculty in the dialogue and planning. We have requested $15,000 per year for travel expenses.

Finally, the center will take an active role in minority recruiting by having its members attend major national conference for minority students and scientists such as the Annual Biomedical Research Conference for Minority Students (ABRCMS), and the Society for Advancement of Chicanos and Native Americans in Science (SACNAS). $10,000 will be allocated annually for these conferences. Some of these funds will be used to send undergraduate and graduate students to the conferences.

Once the diversity programs are implemented we will reach out to additional institutions such as Spelman College and the University of Maryland Baltimore County whose students and faculty would greatly benefit from taking part in the center’s activities.

4.e. Knowledge Transfer Plans and Objectives

The main expected impact of the CBMM is to develop a broad theoretical and experimental framework for capturing aspects of intelligence at the interface between well-established domains of research such as language, vision, perception, planning, and movement control. This new integrated approach encourages breakthroughs in our understanding of computations in the brain. We expect this research to have an important impact on data-rich applications in engineering (computer vision, image processing) and natural sciences (linguistics, genomics).

Knowledge transfer is essential to furthering the objectives of the CBMM. We anticipate three types of transfer: a) Between the institutions in the Center—MIT, Rutgers, and Harvard, as well as our minority-serving partners—to ensure that research collaborations are close, and that Research, Education, and Diversity objectives are integrated; b) To industry, to ensure that new breakthroughs are rapidly translated into new products; and c) To society, to ensure that results are communicated to the widest possible audience of public leaders and citizens.

Knowledge Transfer Between Institutions

An essential aspect of this field is that it is interdisciplinary. Our Center will foster interdisciplinary collaborations across its participant Institutions with seminars, workshops, colloquia, and conferences.

As an additional means of knowledge transfer, we will support a postdoctoral exchange program. These will be researchers who are working on collaborative projects between the participating institutions and who will spend periods of up to 6 months at a collaborating institution. Funding will come from the research funds directed to each institution.

A faculty exchange program will be organized to support extended stays of 3-6 months at another institution. Applications will be solicited on a yearly basis for both the postdoctoral and faculty exchange programs. In order to build and sustain interactions among the institutions, we envision a program that would fund postdoctoral researchers who are jointly supervised by at least two faculty members, dividing their time between respective institutions.

Seed grant program. The new frontier of Intelligence research will require adventurous exploration, particularly the type of high-risk/high-reward projects that are not easily funded by traditional grant mechanisms. We also need mechanisms to encourage and enable scientists to jump fields and explore new ideas. Here, too, traditional grant mechanisms are unsympathetic to newcomers to any field who might have innovative new ideas but no proven track record. To bypass these funding difficulties, our Center will annually award money to short-term seed projects to foster creativity and cross-disciplinary endeavors.

Each year, we will solicit scientists—both members of the Center as well as the broader communities at MIT, Rutgers, and Harvard—for their best new ideas in Intelligence research. We will ask for short research proposals; 2 pages should be sufficient to describe the “big picture” goal of each seed project in broad strokes. The CBMM administration will award seed grants to investigators on the basis of the novelty and ambition of their proposals. Projects that are deemed fundable by traditional mechanisms will be discouraged. The duration of seed grants will be short, typically 1-2 years or enough to get the ball rolling. We will actively seek ideas with the greatest transformative potential. In fact, this seed grant mechanism may be a way to naturally evolve the leading direction of our Center, as what starts as a seed may grow and eventually bear fruit.

Symposia. Major meetings in the various disciplines that intersect within our proposed CBMM host symposia and workshops that we can use to attract the next generation of researchers in the field of Intelligence. We initially will target two meetings:
Annual meeting on Neural Information Processing Systems (NIPS)
Computational and Systems Neuroscience (COSYNE)

Weekly Seminars. The Center will host a weekly lunchtime seminar throughout the academic year and summers, allowing each trainee to present their research progress about twice a year to other trainees and faculty members. The weekly lunch seminars will provide all trainees with regular feedback from their physics and neuroscience colleagues, as well as valuable experience in delivering lectures and seminars.

Brains, Minds and Machines Seminar Series. The Center will sponsor a series of weekly talks on intelligence research, which will alternate between the MIT and the Harvard campus. This series will replace the present Intelligence Initiative series, which takes place each week or so in building 46 at MIT (see http://isquared.mit.edu/newsevents).

Annual Summer Retreat. Each year, an annual retreat for the CBMM will be hosted by rotation between MIT, Harvard, and Rutgers. Participating faculty, students, and postdocs from these three institutions plus Hunter College, Howard University, Wellesley College, and the University of Puerto Rico all would attend and many would present. Every effort would be made to ensure full participation at all levels by women and minorities.

International Programs. Three Center members have close ties to international programs that we will further develop if funding is awarded. Haim Sompolinsky is both a Professor at The Hebrew University in Jerusalem and a long-term Visiting Professor at Harvard. Consequently we are well suited to implement an international component to our postdoctoral training. Hebrew University has a particularly strong track record in training computational neuroscientists, who often end up in the United States, so this would be of great bi-directional benefit. The Safra Center at Hebrew University is a world-class institution for theoretical neuroscience. Haim Sompolinsky's presence there will facilitate a study abroad program in which postdocs will be able to spend time in Israel, taking an intensive course in Theoretical Neuroscience and undertaking a research project under the guidance of a faculty member at Hebrew University. Additionally, we plan joint meetings to disseminate advances with our international partners more effectively. Shimon Ullman is both a Professor at the Weizman Institute in Rehovot and an Adjunct Professor at MIT. We expect a few postdoctoral scholars each year to be interested in theoretical neuroscience. These young scientists would benefit from an immersive experience in theoretical neuroscience, which we will be able to offer through an international partnership with Hebrew University and with the Weizman Institute.

L. Mahadevan, who is a Professor at Harvard with close ties to National Center for Biological Sciences in Bangalore, will facilitate an exchange for postdoctoral scholars with India. Postdocs will be shared jointly between our proposed Center and the Center in Bangalore, spending two years in the US and two years in India. As with Israel, we plan joint meetings to rapidly spread the advances gained in our proposed Center. Similar arrangements will be made with Tübingen (Max Planck Institute) and Genova (IIT). There is already a IIT-MIT lab at MIT which supports with IIT funds long and brief visits by IIT scientists at CBCL.

Knowledge Transfer to Industry

It is not difficult to imagine the major economic and industrial impact of advances in our ability to understand, capture, and engineer intelligence. It is far more difficult at the outset of our ambitious program to pick winners—individual companies that are best poised to take advantage of the scientific advances our Center will produce. One might guess that today's major players—the research arms of Microsoft and Google, existing robot and computer vision system manufacturers, for example—would be primary beneficiaries, but it is more likely that entirely new companies will emerge to harness these discoveries.

The CBMM will offer summer internships to selected researchers from our industrial partners. We have identified this as the most effective way to efficiently transfer knowledge about the developments in the science and technology of intelligence to companies that are in position to develop and deploy these ideas as new products. Members of our scientific team already have strong contacts with leaders on the industrial side. We are confident that as our research progresses, this list, and hence the breadth of summer internship contacts, will grow.

- Google (Dr. Peter Norvig, Director of Research)

- Microsoft (Dr. Andrew Blake, Managing Director, Cambridge, UK)
- IBM (Dr. David Ferrucci, IBM’s T.J. Watson’s Research Center)
- Orcam (Dr. Amnon Shashua, CEO)
- Mobileye (Dr. Amnon Shashua, CTO)
- DeepMind (Demis Hassabis, CEO)

There will also be closer ties with several of the companies listed above as well as others in the future. In particular, IBM T.J. Watson’s Research Center has an existing internship program for graduate students. As discussed with Ferrucci, IBM is looking forward to interacting closely with the center to develop internship opportunities of mutual benefit to CBMM, IBM, and relevant students. Similarly, Microsoft Research (MSR) has an existing, robust internship program for both graduate and undergraduate students. There are multiple common interests between the proposed programme and the research interests of MSR, and also excellent collaborative connections between scientists in the center and in MSR. Blake anticipates that several internships will arise from these connections, and this will significantly enrich the proposed program of work. In general, through the extensive industry contacts of the CBMM faculty and generally at MIT and Harvard, we are and will continue to be in constant contact with companies that may benefit from interactions with the CBMM.

Knowledge Transfer via Cyberinfrastructure

This topic is covered in detail in the earlier section on Education, and funded through that mechanism. As noted there, for the past decade, MIT has been a leader in making advanced course materials available online. MIT began making materials from most courses freely available online via OpenCourseWare (OCW) in 2002. Remarkably, OCW now covers over 2,000 courses, and has logged 127 million user sessions by an estimated 90 million visitors. Just this month, MIT announced the latest advance, dubbed M.I.T.x, which, beginning in 2012 will provide online certification for successful completion of online coursework, without any admissions requirement. New components will include student-to-student and group discussion, online tutors, online laboratories, crowd-sourced grading of programs, and machine learning approaches to tailor material.

4.f Management Plan for Research, Education, Diversity, and Knowledge Transfer

The CBMM will be an inter-institution entity serving an academic community interested in integrated research on intelligence. Management challenges include coordinating activities carried out at different institutions. It is critical to have a good decision making process and efficient operational procedures. The Center's base of operations will be MIT, but with close organizational and infrastructural connections to Harvard, Rutgers, UCLA, Stanford, and Rockefeller. The organizational chart identifies the Director (Poggio) and Associate Directors from Harvard (Mahadevan) and MIT (Wilson). H. Hirsh will represent the other Institutions (Rutgers, UCLA, Stanford, Caltech, Rockefeller).

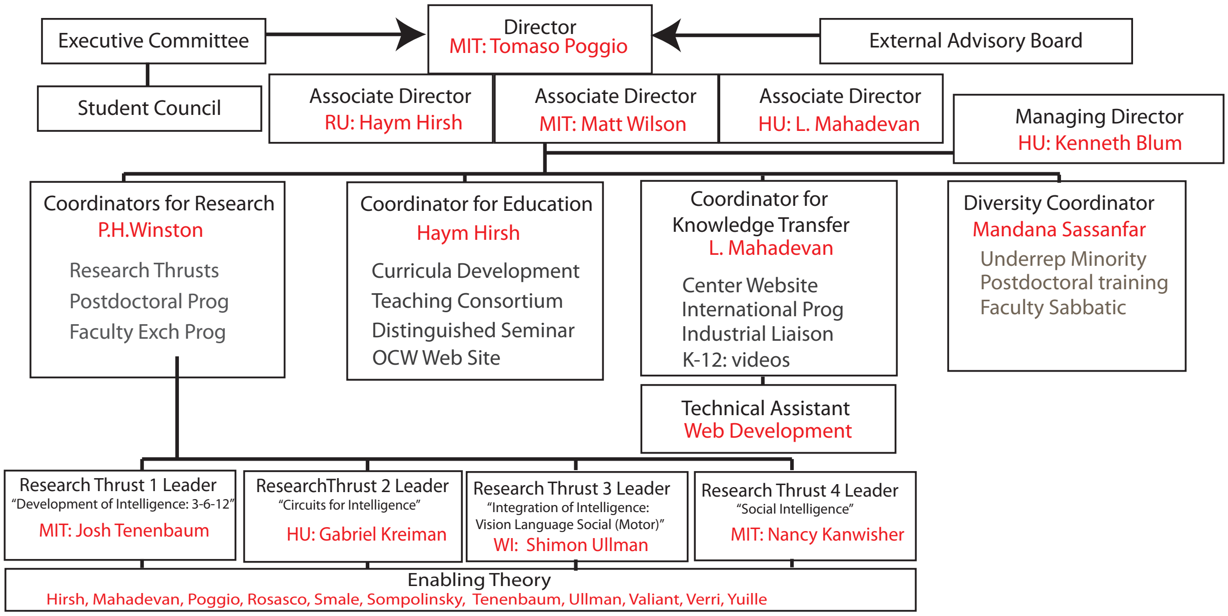

Hirsh, Mahadevan, and Sassanfar will coordinate Education, Research & Knowledge Transfer, and Diversity, respectively. This team of six comprises the top tier Executive Committee, which will meet monthly by Skype. In addition, the Executive Committee will meet with the leaders of the Research Thrusts yearly to review the past year's activities and plan for the next. The Center will be administered through a central office in the MIT CBCL Center@MIT, which will be renamed Center for Brains, Minds and Machines. Each person in a leadership role has strong management credentials. The Director (Poggio) is currently the Director of the Center for Biological and Computational Learning and a co-leader of the MIT Intelligence Initiative. Of the three Associate Directors, one (Hirsh) has previously been Director of the Division of Information and Intelligent Systems at the National Science Foundation, and one (Wilson) is the associate chairman of the BCS Department. The Research coordinator (Winston) has been Director of the MIT Artificial Intelligence Laboratory for 25 years.

An Executive Committee for Research (ECR) will be comprised of the Director, three Associate Directors, and the four Thrust Leaders. **An Executive Committee for Policy Issues (ECP)** will be comprised of the Director, three Associate Directors, and the Coordinators for Diversity, Education, and the Management Coordinator (K. Blum). The ECR will meet monthly by Skype. Its responsibilities include charting the research direction of the CBMM reviewing existing projects, and selecting new projects, both primary and seed. The Executive Committees will also meet yearly at the Summer Retreat to review the past year's activities and to plan for the next. The ECP will meet in person a second time each year. We anticipate a dynamic funding model, which is flexible from year to year to accommodate new ideas and to rotate researchers.

Responsibilities of Lead and Partner Institutions. Physical facilities necessary for the CBMM programs are based at each institution. A list of facilities and their capabilities at all partner institutions will be maintained on the CBMM website. These facilities will provide office and laboratory space for the students participating in exchange programs and for collaborative research. Initially, access to the facilities is provided to CBMM students and postdocs only. We foresee that, as research activities increase, these facilities may be opened to other researchers interested in investigating new ideas before investing in their own equipment. Each partner institution will actively participate in all CBMM activities: research, education, diversity, and knowledge transfer. However, each will have one primary responsibility. MIT: the coordination of all CBMM activities; Rutgers: Education, including Summer Schools and the Graduate Teaching Consortium; Harvard: Knowledge transfer; MIT: Diversity, including outreach programs coordination with our minority serving institutions, and administration of the URM research grant program. The research activities are organized into four thrusts; Thrusts1, 3 and 4 conducted at MIT, Thrust 2 at Harvard.

External Advisory Committee (EAC). An External Advisory Committee comprised of leaders from academia and industry will meet annually to ensure the scientific and technological relevance of the programs. Joel Oppenheim (Senior Associate Dean for Biomedical Sciences at NYU, and Director of the Sackler Institute of Graduate Biomedical Science) has agreed to serve on the EAC, subject to NSF approval. We look forward to working with NSF in selecting further EAC members, especially to ensure that diversity and outreach are adequately addressed. Yearly meetings of the EAC will rotate among the three primary institutions. The EAC will advise on all aspects of the CBMM and specific members will cover research, education, diversity, and knowledge transfer.

STC Research Coordination*.* To make the CBMM more than the sum of its parts, we encourage insights from colleagues in adjacent research areas, and put a premium on collaborations. Collaboration will be maximized by updating the membership of the research faculty teams from year to year. Our plan is to start with a relatively large number of faculty (> 20, counting US Institutions) in the Center to let the most promising cooperative research to emerge. We will then select at the end of the first year -- and in subsequent years -- the most promising and interesting projects and researchers, expecting a smaller number of faculty (< 20) at regime. Mechanisms to encourage collaboration include the large number and diverse interests of CBMM members. At the end of each year we will select the most promising and interesting projects and researchers through a careful review process involving external advisors. The ECR will be responsible for oversight of all research activities including monthly reviews of the individual projects to be conducted by the Thrusts leaders. A major review of each Thrust will also take place at each Summer Retreat, where participants will present the latest results, and outside reviewers will provide comments and criticisms. Criteria for evaluation will be the: a) Quality of the research, b) Progress, c) Coherence with the overarching goals off the Center, and d) Collaborative character of the project. The day after these presentations, the project leaders, post docs, and other researchers will meet for a "brain-storming" session to identify possible new inter-thrust initiatives. Based on the project review and "brain-storming" session, the ECR and the external advisors will adjust budgets, terminate projects, initiate projects, and renew individual projects and research teams. A portion of the research budget ($150K) will be set aside to fund 3 seed projects, which must be collaborations between investigators at more than one institutional partner. The seed grants will ensure the evolution of ideas within the CBMM. Projects that move research into areas of high risk and high potential will be emphasized. Seed projects will be funded for a maximum of two years, and then become regular projects or terminated. All of the Postdocs, across all thrusts, will convene monthly to "think out side the box."

Educational Programs. Rutgers and the Education Coordinator (Hirsh) will provide oversight of all educational activities; Hirsh will also have the responsibility of the overall coordination of Education, Knowledge Transfer, and Diversity efforts. The formation and coordination of a new graduate program to be adopted in Schools beyond Harvard and MIT will require close coordination, especially during start-up. The ECP will meet regularly, especially during the first two years, and coordinate with the graduate committees at each institution. The success of the Graduate Teaching Consortium (GTC) requires effective management and a plan for evaluation and renewal. Oversight of the GTC will be provided by the ECP with input from a the Education Coordinator, the Education Director, and one member from each of the participating universities (the GTC Oversight Committee). The GTC Oversight Committee will meet once each year at the Summer Retreat where courses will be reviewed and evaluated based on student participation, student evaluations, diversity, and coherence with the goals of the Center. Proceedings will be submitted to the Director and reviewed at a regular ECP meeting. We expect competition among the members of the GTC for their courses showcased by the CMBB. Courses are to be taught by recognized leaders and to be developed also for MIT-X. New courses will be solicited each year, both from existing GTC members and other universities that wish to participate. Funding will be included in the UIUC sub-contract and managed by the Education Director.

Knowledge Transfer. The Knowledge Transfer activities will be managed at Harvard by the Coordinator (Mahadevan) with help from the Managing Director (Blum). Knowledge transfer is a top priority because of the potentially wide-ranging implications of the proposed research. Both the ECR and ECP will provide oversight of various aspects of knowledge transfer, the former in connection with broad dissemination of the research findings, and the latter with the outreach activities described elsewhere. The Managing Director will administer funds.

Management of Diversity Programs. Primary contacts have been identified at the Wellesley, Puerto Rico, Hunter, and Howard campuses. These individuals and the Diversity Coordinator will form a Diversity Committee that meets once each year at the Summer Retreat. The CBMM Director and the ECP will review proceedings of the Diversity Committee meetings. All diversity-related programs will be coordinated from MIT by the Diversity Coordinator (Sassanfar). Funds for research activities at the minority serving institutions are in the MIT budget and will be managed by the Diversity Coordinator.

Ethical considerations. Intelligent machines are increasingly becoming part of everyday life. Ethical concerns need to be recognized and discussed openly and vigorously from the outset. These considerations will need to be addressed in a systematic manner, and for that purpose we will seek assistance from organizations dealing with neuroethics and with machine ethics.

Replacing the Center Director. If it becomes necessary to replace the Director, an open search will be undertaken across all three participating institutions. The search committee will be named by the MIT Dean of Science and be comprised of current members of the Center. While the search is being conducted, if the former Director is unable to perform the duties as Director, the MIT Associate Director (Wilson) will serve as Acting Director.

5. Facilities, Equipment and Other Resources

The CBMM offices at MIT will be within Prof. Tomaso Poggio's Center for Biological and Computational Learning including and additional space offered by the McGovern Institute. The existing CBCL will change his name to the new Center for Brains Minds and Machines. CBCL is now housed in a new multidisciplinary neuroscience research building which unites all of MIT's diverse brain and cognitive sciences research projects under one roof and presently includes the Brain and Cognitive Science Department, the McGovern Institute for Brain Research and the Picower Institute for Memory and Learning within Building (46) located opposite to the Computer Science and Artificial Intelligence Laboratory (CSAIL). Building 46 houses the laboratories of faculty in systems and computational neuroscience; imaging and cognitive neuroscience; genetic and cellular neuroscience; learning and memory neuroscience; and in the areas of perception, attention, and consciousness. The MIT faculty involved with the Center are housed in Building 46, CSAIL, and the Media Lab. The main Harvard campus is two T stops away. We will rely on the core administrative infrastructure of CBCL and especially of the McGovern Institute (also in Building 46), of which CBCL is also part, especially for financial, human resources and outreach part of our activities. The McGovern Headquarters staff, which consists of seven staff members reporting to the director, provides administrative support to all McGovern faculty, in addition to that provided by their own lab administrators. This central support includes grant management, event organization, donor outreach and communications, including press releases, media outreach and miscellaneous promotional materials (newsletter, website etc) to raise the visibility of McGovern faculty members and their research. Of course the BMM Center plans to complement the existing resources it with specific hires described in the budget justification.

For running computational models, CBCL has a private compute cluster consisting of 112 CPU cores augmented by 31 high-end NVIDIA GPUs, each of which can outperform a small CPU cluster using CBCL's CNS simulation framework. The cluster, 13 terabytes of RAID storage, and 30 high-end workstations are connected via a private gigabit network. Storage is backed up nightly to MIT's offsite TSM data silo. CBCL also has access to MIT's 250-node CSAIL (Computer Science and Artificial Intelligence) cluster and to its technical support team. CBCL employs a systems administration group that installs and monitors its workstations, servers, and cluster and maintains its network security, backups, mail and network connectivity, and web server.

The Center for Brain Science (CBS) will be the primary locus of the efforts at Harvard University, and will administer the subcontract for this grant, if awarded. CBS integrates the theoretical and computational neuroscience efforts at Harvard. It houses a core facility for magnetic resonance imaging, with a 3 Tesla Siemens Magnetom Trio A Tim human brain scanner and a small bore Bruker 4.7 Tesla animal scanner. This facility is used for cognitive neuroscience, technology development, and for undergraduate education. The Center scientists will have access to the Odyssey ten thousand-plus core Linux cluster and several high-end scientific computing systems, ranging from a 4096-processor IBM Bluegene/L supercomputer to a state-of-the-art 7680-core Nvidia Tesla GPGPU cluster to a Cloud compute cluster that provides flexible, on-demand compute capabilities to the SEAS community. CBS also houses an administrative structure to coordinate scientific activities and administer funds, with office space that can accommodate the additional support staff requested in this proposal.

Harvard also has a Research Experience for Undergraduates (REU) program that we plan to leverage for the proposed CBMM. The REU is a mechanism to provide summer research opportunities for a diverse group of undergraduates. Harvard provides a coordinated educational and research experience to inspire and encourage young scientists to continue on to graduate school. Professional development workshops, faculty seminars on research and ethics, student seminars to gain experience presenting original research, a final written report to practice scholarly scientific communication, and community activities all are integrated into this successful program.

7. References Cited

[1] Abeles, M. 1991. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press.

[2] Amit, D.J., Gutfreund, H. and Sompolinsky, H. 1985. Spin-glass models of neural networks. *Physical Review A*. APS.

[3] Anderson, J.R. 1983. *The Architecture of Cognition*. Harvard University Press.

[4] Baker, C.L., Saxe, R.R. and Tenenbaum, J.B. 2011. Bayesian Theory of Mind: Modeling Joint Belief-Desire Attribution. *Proceedings of the Thirty-Third Annual Conference of the Cognitive Science Society* (2011).

[5] Balas, B., Saxe, R. and Kanwisher, N. (in press). Thin-slice perception develops slowly. *Journal of Experimental Child Psychology*.

[6] Barto, A.G. and Mahadevan, S. 2003. Recent Advances in Hierarchical Reinforcement Learning. *Discrete Event Dynamic Systems*. 13, 4 (2003), 341-379.

[7] Beck, J.M., Latham, P.E. and Pouget, A. 2011. Marginalization in Neural Circuits with Divisive Normalization. *Journal of Neuroscience*. 31, 43 (2011), 15310-15319.

[8] Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E. and Pouget, A. 2008. Probabilistic population codes for Bayesian decision making. *Neuron*. 60, 6 (2008), 1142-1152.

[9] Bousquet, O., Boucheron, S. and Lugosi, G. 2004. Introduction to Statistical Learning Theory. *Advanced Lectures on Machine Learning Lecture Notes in Artificial Intelligence*. O. Bousquet, U. von Luxburg, and G. Ratsch, eds. Springer. 169-207.

[10] Boyden, E.S., Zhang, F., Bamberg, E., Nagel, G. and Deisseroth, K. 2005. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*. 8, 9 (2005), 1263-1268.

[11] Brin, S. and Page, L. 1998. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*. 30, 1-7 (1998), 107-117.

[12] Brunel, N. 2000. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*. 8, 3 (2000), 183-208.

[13] Bullier, J. 2001. Feedback connections and conscious vision. *Trends in Cognitive Sciences*. 5, 9 (2001), 369-370.

[14] Burak, Y., Rokni, U., Meister, M. and Sompolinsky, H. 2010. Bayesian model of dynamic image stabilization in the visual system. *Proceedings of the National Academy of Sciences of the United States of America*. 107, 45 (2010), 19525-19530.

[15] Cang, J., Rentería, R.C., Kaneko, M., Liu, X., Copenhagen, D.R. and Stryker, M.P. 2005. Development of precise maps in visual cortex requires patterned spontaneous activity in the retina. *Neuron*. 48, 5 (2005), 797-809.

[16] Carey, S. 2009. *The Origin of Concepts*. Oxford University Press.

[17] Chow, B.Y. and Boyden, E.S. 2011. Synthetic Physiology. *Science*. 332, 6037 (2011), 1508-1509.

[18] Chow, B.Y., Han, X., Dobry, A.S., Qian, X., Chuong, A.S., Li, M., Henninger, M.A., Belfort, G.M., Lin, Y., Monahan, P.E. and Boyden, E.S. 2010. High-performance genetically targetable optical neural silencing by light-driven proton pumps. *Nature*. 463, 7277 (2010), 98-102.

[19] Churchland, P.S. and Sejnowski, T.J. 1988. Perspectives on cognitive neuroscience. *Science*. 242, 4879 (1988), 741-745.

[20] Connor, C.E., Brincat, S.L. and Pasupathy, A. 2007. Transformation of shape information in the ventral pathway. *Current Opinion in Neurobiology*. 17, 2 (2007), 140-147.

[21] Cucker, F. and Smale, S. 2001. On the mathematical foundations of learning. *Bulletin of the American Mathematical Society*. 39, 01 (2001), 1-50.

[22] D Mumford 1992. On the computational architecture of the neocortex. *Biological Cybernetics*. 66, 3 (1992), 241-251.

[23] Dayan, P. and Abbott, L. 2001. *Theoretical Neuroscience*. MIT Press.

[24] Dehaene, S. 2009. *Reading in the brain*. Viking.

[25] Van Den Heuvel, M.P. and Sporns, O. 2011. Rich-Club Organization of the Human Connectome. *Journal of Neuroscience*. 31, 44 (2011), 15775-15786.

[26] Devroye, L., Györfi, L. and Lugosi, G. 1996. *A Probabilistic Theory of Pattern Recognition*. Springer.

[27] DiCarlo, J.J. and Cox, D.D. 2007. Untangling invariant object recognition. *Trends in Cognitive Sciences*. 11, 8 (2007), 333-341.

[28] Duncan, J. 2010. The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends in Cognitive Sciences*. 14, 4 (2010), 172-179.

[29] Elman, J.L. 2005. Connectionist models of cognitive development: where next? *Trends in Cognitive Sciences*. 9, 3 (2005), 111-117.

[30] Engel, A.K., Fries, P. and Singer, W. 2001. Dynamic predictions: oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience*. 2, 10 (2001), 704-716.

[31] Epstein, S.L. 2005. Making interdisciplinary collaboration work. *Interdisciplinary Collaboration An Emerging Cognitive Science*. S.J. Derry, C.D. Schunn, and M.A. Gernsbacher, eds. Lawrence Erlbaum. 245 -264.

[32] Fan, J., Gu, X., Guise, K.G., Liu, X., Fossella, J., Wang, H. and Posner, M.I. 2009. Testing the behavioral interaction and integration of attentional networks. *Brain and Cognition*. 70, 2 (2009), 209-220.

[33] Feldman, D.E. 2009. Synaptic Mechanisms for Plasticity in Neocortex. *Annual Review of Neuroscience*. 32, 1 (2009), 33-55.

[34] Felzenszwalb, P., McAllester, D. and Ramanan, D. 2008. A discriminatively trained, multiscale, deformable part model. *IEEE Conference on Computer Vision and Pattern Recognition (2008)*. 08, [9 (2008), 1-8.

[35] Ferrucci, D., Brown, E., Chu-carroll, J., Fan, J., Gondek, D., Kalyanpur, A.A., Lally, A., Murdock, J.W., Nyberg, E. and Prager, J. 2010. Building Watson : An Overview of the DeepQA Project. *AI Magazine*. 31, 3 (2010), 59-79.

[36] Fitzpatrick, P., Metta, G. and Natale, L. 2008. Towards long-lived robot genes. *Robotics and Autonomous Systems*. 56, 1 (2008), 29-45.

[37] Freiwald, W.A. and Tsao, D.Y. 2010. Functional Compartmentalization and Viewpoint Generalization Within the Macaque Face-Processing System. *Science*. 330, 6005 (2010), 845-851.

[38] Földiák, P. 1991. Learning Invariance from Transformation Sequences. *Neural Computation*. 3, 2 (1991), 194-200.

[39] Gergely, G., Bekkering, H. and Király, I. 2002. Rational imitation in preverbal infants. *Nature*. 415, 6873 (2002), 755.

[40] Goodman, N., Mansinghka, V., Roy, D. and K 2008. Church: a language for generative models. *Uncertainty in Artificial*. (2008).

[41] Gweon, H. and Schulz, L. 2011. (in press). 16-month-olds rationally infer causes of failed actions. *Science*. 332, 6037 (2011), 1524.

[42] Gweon, H., Tenenbaum, J.B. and Schulz, L.E. 2010. Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences of the United States of America*. 107, 20 (2010), 9066-9071.

[43] Han, X. and Boyden, E.S. 2007. Multiple-Color Optical Activation, Silencing, and Desynchronization of Neural Activity, with Single-Spike Temporal Resolution. *PLoS ONE*. 2, 3 (2007), 12.

[44] Heider, F. and Simmel, M. 1944. An Experimental Study of Apparent Behavior. *The American Journal of Psychology*. 57, 2 (1944), 243-259.

[45] Hinton, G.E. 2007. Learning multiple layers of representation. *Trends in Cognitive Sciences*. 11, 10 (2007), 428-434.

[46] Hochstein, S. and Ahissar, M. 2002. View from the Top : Hierarchies and Reverse Hierarchies Review. *Neuron*. 36, 3 (2002), 791-804.

[47] Jeneson, A. and Squire, L.R. 2011. Working memory, long-term memory, and medial temporal lobe function. *Learning and Memory*. 19, (2011), 15-25.

[48] Jung, R.E. and Haier, R.J. 2007. The Parieto-Frontal Integration Theory (P-FIT) of intelligence: converging neuroimaging evidence. *Behavioral and Brain Sciences*. 30, 2 (2007), 135-154; discussion 154-187.

[49] Katz, B., Borchardt, G. and Felshin, S. 2006. Natural Language Annotations for Question Answering. *Proceedings of the 19th International FLAIRS Conference* (2006).

[50] Katz, L.C. and Crowley, J.C. 2002. Development of cortical circuits: lessons from ocular dominance columns. *Nature Reviews Neuroscience*. 3, 1 (2002), 34-42.

[51] Kokkinos, I. and Yuille, A.L. 2011. (in press). Inference and Learning with Hierarchical Compositional Models. *International Journal of Computer Vision*. (2011).

[52] Kraut, R.E. and Resnick, P. 2011. *Building Successful Online Communities: Evidence Based Social Design*. The MIT Press.

[53] Logothetis, N.K. and Sheinberg, D.L. 1996. Visual object recognition. *Annual Review of Neuroscience*. 19, 1 (1996), 577-621.

[54] Marr, D. 2010. *Vision*. MIT Press.

[55] McKone, E., Crookes, K. and Kanwisher, N. 2009. The Cognitive and Neural Development of Face Recognition in Humans. *The Cognitive Neurosciences*. M. Gazzaninga, ed. 467-482.

[56] Meltzoff, A.N., Kuhl, P.K., Movellan, J. and Sejnowski, T.J. 2009. Foundations for a New Science of Learning. *Science*. 325, 5938 (2009), 284-288.

[57] Metta, G., Natale, L., Nori, F., Sandini, G., Vernon, D., Fadiga, L., Von Hofsten, C., Rosander, K., Lopes, M., Santos-Victor, J., Bernardino, A. and Montesano, L. 2010. The iCub humanoid robot: an open-systems platform for research in cognitive development. *Neural Networks*. 23, 8-9 (2010), 1125-1134.

[58] Miller, E.K. and Cohen, J.D. 2001. An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*. 24, 1 (2001), 167-202.

[59] MobilEye: *http://www.mobileye.com/*.

[60] Munakata, Y. and McClelland, J.L. 2003. Connectionist models of development. *Developmental Science*. 6, 4 (2003), 413-429.

[61] Oliva, A. and Torralba, A. 2006. Building the gist of scene: the role of global image features in recognition. *Progress in Brain Research*. 155, (2006), 23-36.

[62] Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.

[63] Poggio, T. and Smale, S. 2003. The Mathematics of Learning : Dealing with Data 1 Introduction. *Notices of the American Mathematical Society*. 50, 5 (2003), 537-544.

[64] Poggio, T.A. 2011. The Computational Magic of the Ventral Stream: Towards a Theory. *Nature Precedings*. (2011), 1-26.

[65] Ponce, J., Berg, T.L., Everingham, M., Forsyth, D.A., Herbet, M., Lazebnik, S., Marszalek, M., Schmid, C., Russell, B.C., Torralba, A., Williams, C.K.I., Zhang, J. and Zisserman, A. 2006. Dataset Issues in Object Recognition. *In Toward Category-Level Object Recognition. Springer-Verlag Lecture Notes in Computer Science*. J. Ponce, M. Hebert, C. Schmid, and A. Zisserman, eds. Springer-Verlag. 29-48.

[66] Ponce, J., Herbet, M., Schmid, C. and Zisserman, A. eds. 2006. *In Toward Category-Level Object Recognition. Springer-Verlag Lecture Notes in Computer Science*. Springer-Verlag.

[67] Posner, M.I. 2011. Imaging attention networks. *NeuroImage*. (2011).

[68] Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlagger, B.L. and Peterson, S.E. 2011. Functional network organization of the human brain. *Neuron*. 72, 4 (2011), 665-678.

[69] Rao, R.P. and Ballard, D.H. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*. 2, 1 (1999), 79-87.

[70] Rizzolatti, G. 2007. *Mirrors in the brain*. Oxford University Press.

[71] Rolls, E.T. 1991. Neural organization of higher visual functions. *Current Opinion in Neurobiology*. 1, 2 (1991), 274-278.

[72] Rolls, E.T. and Deco, G. 2002. *Computational Neuroscience of Vision*. Oxford University Press.

[73] Sandini, G., Metta, G. and Vernon, D. 2007. The iCub Cognitive Humanoid Robot: An Open-System Research Platform for Enactive Cognition. *50 Years of Artificial Intelligence*. M. Lungarella, F. Iida, J. Bongard, and PfeiferR., eds. Springer. 358-369.

[74] Saxe, R., Tenenbaum, J.B. and Carey, S. 2005. Secret agents: Inferences about hidden causes by 10- and 12-month-old infants. *Psychological Science*. 16, 12 (2005), 995-1001.

[75] Sejnowski, T.J., Koch, C. and Churchland, P.S. 1988. Computational neuroscience. *Science*. 241, 4871 (1988), 1299-1306.

[76] Serre, T., Kreiman, G., Kouh, M., Cadieu, C., Knoblich, U. and Poggio, T. 2007. A quantitative theory of immediate visual recognition. *Progress in Brain Research*. 165, 06 (2007), 33-56.

[77] Seung, H.S., Sompolinsky, H. and Tishby, N. 1992. Statistical mechanics of learning from examples. *Physical Review A*. APS.

[78] Shepherd, S.V., Steckenfinger, S.A., Hasson, U. and Ghazanfar, A.A. 2010. Human-monkey gaze correlations reveal convergent and divergent patterns of movie viewing. *Current Biology*. 20, 7 (2010), 649-656.

[79] Siri: *http://sri.com/about/siri.html*.

[80] Smale, S., Rosasco, L., Bouvrie, J., Caponnetto, A. and Poggio, T. 2009. Mathematics of the Neural Response. *Foundations of Computational Mathematics*. 10, 1 (Jun. 2009), 67-91.

[81] Spelke, E.S., Breinlinger, K., Macomber, J. and Jacobson, K. 1992. Origins of knowledge. *Psychological Review*. 99, 4 (1992), 605-632.

[82] Steinwart, I. and Christmann, A. 2008. *Support Vector Machines*. Springer.

[83] Stringer, S.M., Perry, G., Rolls, E.T. and Proske, J.H. 2006. Learning invariant object recognition in the visual system with continuous transformations. *Biological Cybernetics*. 94, 2 (2006), 128-142.

[84] Sussillo, D. and Abbott, L.F. 2009. Generating coherent patterns of activity from chaotic neural networks. *Neuron*. 63, 4 (2009), 544-557.

[85] Tanaka, K. 1996. Inferotemporal cortex and object vision. *Annual Review of Neuroscience*. 19, 1 (1996), 109-139.

[86] Tenenbaum, J.B., Kemp, C., Griffiths, T.L. and Goodman, N.D. 2011. How to grow a mind: statistics, structure, and abstraction. *Science*. 331, 6022 (2011), 1279-1285.

[87] Tomasello, M. 2009. *Why We Cooperate*. MIT Press.

[88] Tsagarakis, N.G., Metta, G., Sandini, G., Vernon, D., Beira, R., Becchi, F., Righetti, L., Santos-Victor, J., Ijspeert, A.J., Carrozza, M.C. and Caldwell, D.G. 2007. iCub - The Design and Realization of an Open Humanoid Platform for Cognitive and Neuroscience Research. *Journal of Advanced Robotics, Special Issue on Robotic platforms for Research in Neuroscience*. 21, 10 (2007), 1151-1175.

[89] Tsao, D.Y., Freiwald, W.A., Tootell, R.B.H. and Livingstone, M.S. 2006. A cortical region consisting entirely of face-selective cells. *Science*. 311, 5761 (2006), 670-674.

[90] Tsotsos, J.K. 1990. Analyzing vision at the complexity level. *Behavioral and Brain Sciences*. Cambridge University Press.

[91] Téglás, E., Vul, E., Girotto, V., Gonzalez, M., Tenenbaum, J.B. and Bonatti, L.L. 2011. Pure reasoning in 12-month-old infants as probabilistic inference. *Science*. 332, 6033 (2011), 1054-1059.

[92] Ullman, S. 1996. *High-Level Vision*. MIT Press.

[93] Ullman, S. 1984. Visual routines. *Cognition*. 18, 1-3 (1984), 97-159.

[94] Ullman, T., Baker, C., Macindoe, O., Evans, O., Goodman, N. and Tenenbaum, J. 2009. Help or Hinder: Bayesian Models of Social Goal Inference. *Advances in Neural Information Processing Systems*. (2009), 1874-1882.

[95] Ullman, T.D., Goodman, N.D. and Tenenbaum, J.B. 2010. Theory Acquisition as Stochastic Search. *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society* (2010).

[96] Valiant, L.G. 1984. A theory of learnable. *Proceedings of the sixteenth annual ACM symposium on Theory of computing STOC 84*. 27, 11 (1984), 436-445.

[97] Valiant, L.G. 2000. Robust logics. *Artificial Intelligence*. 117, 2 (2000), 231-253.

[98] Vapnik, V.N. 1998. *Statistical Learning Theory*. Wiley-Interscience.

[99] Vapnik, V.N. 1995. *The Nature of Statistical Learning Theory*. Springer.

[100] Vernon, D., Von Hofsten, C. and Fadiga, L. 2011. *A Roadmap for Cognitive Development in Humanoid Robots*. Springer.

[101] Vernon, D., Metta, G. and Sandini, G. 2007. A Survey of Artificial Cognitive Systems: Implications for the Autonomous Development of Mental Capabilities in Computational Agents. *IEEE Transactions on Evolutionary Computation*. 11, 2 (2007), 151-180.

[102] Vul, E., Goodman, N.D., Griffiths, T.L. and Tenenbaum, J.B. 2009. One and Done? Optimal Decisions From Very Few Samples. *Proceedings of the 31st Annual Meeting of the Cognitive Science Society* (2009), 66-72.

[103] Wang, X.J. 2010. Neurophysiological and computational principles of cortical rhythms in cognition. *Physiology Review*. 90, 3 (2010), 1195 - 1268.

[104] Weigelt, S., Koldewyn, K. and Kanwisher, N. (in preparation). Domain-specific development of face memory, not face perception.

[105] Weisbuch, M. and Ambady, N. 2011. Thin-Slice Vision. *The Science of Social Vision*. R.B. Adams, N. Ambady, K. Nakayama, and S. Shimojo, eds. Oxford University Press.

[106] Zorzos, A.N., Boyden, E.S. and Fonstad, C.G. 2010. Multiwaveguide implantable probe for light delivery to sets of distributed brain targets. *Optics Letters*. 35, 24 (2010), 4133-4135.

11. Ethics Plan

All students, PDAs, staff and faculty in the BMM Center will be required to participate in a program of instruction on the Responsible Conduct of Research and IP Rights. In setting up our program, we pattern it after a curriculum developed by Prof Matt Wilson for Instruction in the Responsible Conduct of Research in the Department of Brain and Cognitive Science at MIT for its graduate program. The instruction will include students, postdocs and faculty at Harvard and MIT and consists of a twelve-hour course meeting for two and a half hours daily for one week. It is designed to provide instruction and dialogue on practical ethical issues relating to the responsible conduct of theoretical, human and animal research in the brain and cognitive sciences. Specific emphasis is placed on topics relevant to young researchers including data handling, animal and human subjects, misconduct, mentoring, intellectual property, and publication.

The 7 core instructional areas are:

- Data acquisition, management, sharing and ownership
- Mentor/trainee relationships
- Publication practices and responsible authorship
- Collaborative science
- Research with biological materials
- Research misconduct
- Intellectual property

Assigned readings are taken from the text “Scientific Integrity” by Francis L. Macrina to be completed prior to each class. Sessions begin with a lecture by faculty to introduce each topic. The class is then divided into smaller discussion groups of four to five students each. Case studies prepared for each class are evaluated and two groups are selected to present each case for discussion by the entire group. Faculty members are available to facilitate and guide discussion. Each student submits a short written summary of the discussions at the end of each class. A short quiz on readings from the text is given at the end of the final class. These topics will be raised in several different settings: pizza luncheons at the beginning of each semester; annual January events at which a guest lecturer from the MIT Technology Licensing Office will speak (with Skype connection to Rutgers) for an hour on the subject of Intellectual Property; and at our annual summer retreat. Some of these topics can be addressed via webinars in the video-conferencing facilities used for the graduate teaching consortium, providing yet another opportunity for the three campuses to interact. The BMM Center intellectual property rights will, by necessity, follow those of each individual institution. In the particular case of MIT, the policy is as follows:

“The prompt and open dissemination of the results of M.I.T. research and the free exchange of information among scholars are essential to the fulfillment of M.I.T.'s obligations as an institution committed to excellence in education and research. Matters of ownership, distribution, and commercial development, nonetheless, arise in the context of technology transfer, which is an important aspect of M.I.T.'s commitment to public service. Technology transfer is, however, subordinate to education and research; and the dissemination of information must, therefore, not be delayed beyond the minimal period necessary to define and protect the rights of the parties.”
http://web.mit.edu/tlo/www/community/guide2.html

12. Shared Experimental Facilities

The Center for Brains, Minds and Machines will have access to a unique combination of facilities for computational, cognitive and neuroscience research. The facilities are all located in Cambridge in the Harvard and MIT campuses which are within two T stops of each other.

MIT: Shared Facilities & Equipment:
CBMM will have use of the existing shared facilities housed within the Brain and Cognitive Sciences Department and the McGovern Institute for Brain Research at MIT (MIBR.)

Animal: Experimental animals (fruit flies, frogs, mice, rats, hamsters, ferrets, cats and monkeys) are kept in animal colonies in Building 46 operated by the MIT Division of Comparative Medicine. A wide range of veterinary services are available to researchers, including surgical assistance and post-perative care. The facilities are maintained in accord with federal guidelines and equipped for he varied studies involved. Animal treatment in monitored by the Committee on Animal Care. ll members of a laboratory working with animals are certified by the Institute for handling of animals and awareness of health-related factors.

Confocal Microscopy Facility: A state-of-the-art confocal laser scanning microscope (Biorad 1024 ES) is available at a shared facility. Equipment for video-microscopy and neuroanatomical tracing is also available.

Neurophysiology: Fully equipped laboratory setups for behavioral and electrophysiological work; equipment for multi-electrode recording from awake behaving animals including primates, rats and mice; visual stimulation systems; data acquisition and analysis systems; eye tracking devices. Optical imaging of intrinsic signals in vivo from visual and somatosensory cortex is carried out in individual laboratories. Intracellular and patch recording rigs are available in several laboratories. Imaging rigs for calcium dye imaging are available, together with two-photon microscopes for calcium imaging and FRET imaging of functional activity and dynamics in single cells and their processes in vitro and in vivo.

Tetrode Array Recording Facility: This state of the art facility provides the capability to monitor behavior while recording the neural activity of large ensembles of individual cells (up to 150 simultaneously) from animals engaged in active behavior. Six separate experimental setups currently enable 28 to 72-channel recording from multiple brain regions of freely behaving wild type and genetic mutant mice and rats. The adjustable tetrode array microdrive developed by MIT's Prof. Matthew Wilson produces the maximum single cell activity yield of any currently available in vivo recording technique. Coupled with the availability of genetically engineered learning deficient mutant mice from several collaborating laboratories, it provides a truly unique opportunity to study the mechanisms and function of specific cell populations and brain regions involved in brain function.

Neuroimaging: The Martinos Imaging Center at the McGovern Institute The imaging center is run as a collaboration between MIT and Harvard, with major contributions from the McGovern Institute, the Martinos Center, and HST (the Health Sciences and Technology program of MIT and Harvard). The imaging center contains three sunken bays for the magnets used in fMRI. One bay holds a a 3T Siemens Trio magnet (with an 18 channel digital RF receiver system) for humans and primates. The second bay has a 9.4T Brucker MRI for animal studies. This machine provides higher resolution images, which can then provide insights into areas to be explored in human studies. The third bay houses the new a magnetoencephalography (MEG) scanner for human brain imaging. The principal sources of MEG (magnetoencephalogram) and EEG (electroencephalogram) are synchronous synaptic currents in the cerebral cortex. Under special circumstances, activity in the cerebellum as well as in certain subcortical structures can be detected as well. MEG and EEG provide an instantaneous view of neural activity and thus are complementary to fMRI. The combination of MEG/EEG with fMRI provides images of brain activity with high temporal as well as spatial resolution. Our comprehensive suite of analysis software allows smooth integration of MEG, EEG, MRI, and fMRI data.

Mouse behavior lab: there is a mouse behavioral laboratory that is housed within the McGovern Institute and operated by the Broad Institute. It is used for collaborative projects between the Broad Institute and the MIT neuroscience community.

Machine and Electronics Shop: A machine shop and electronics shop have been established in the Brain and Cognitive Sciences Complex, targeted to the needs of neuroscience researchers and funded by a core grant from National Eye Institute (NEI). The machine shop is staffed by a full-time machinist and the electronics shop is staffed by a full-time electronics technician. Its services are provided free to NEI grant holders and on a cost-recovery basis to other researchers.

Computers: The MIT department of Information Services and Technology operates the campus-wide computer network and provides IT support and consulting services to all MIT personnel. MIT is among the leading centers for computer science research, and access to advanced computing technology and expertise is available especially through our affiliation with CSAIL. CSAIL offers among other facilities a Cluster Computing infrastructure to support both a public batch queued system using Condor and supporting various levels of customized group clustering. Users can dynamically join their own machines to donate resources to the cluster. Presently, the Condor cluster has around 700 CPU's available. In addition to the cluster, Quanta Computers has recently (summer 2011) donated a large computing resource based around Eucalyptus with some Quanta specific software enhancements for defining multiple system application groups and demand based scaling. The hardware consists of 16 enclosures, each with 4 servers. This gives the department a total of 768 physical CPU cores (shown as 1536 CPU's due to Hyper-Threading), and 3072G of RAM.

Harvard University: Shared Facilities & Equipment

Existing core facilities at Harvard will provide researchers in the Center for Brains, Minds, and Machines with unique, state-of-the-art equipment. In addition, Harvard provides academic support for the participating faculty and space in which to perform the research and knowledge transfer activities. These shared facilities are available for use by participants in the proposed STC. The operational costs, including technical staff payroll and equipment maintenance, are funded from other sources. These directly applicable organizational resources are clearly adequate to support the proposed Harvard activities of the CBBM.

Neuroengineering. At Harvard, the Center for Brain Science (CBS) has a Neuroengineering core facility in the Northwest Building in Cambridge that provides customized engineering solutions to neuroscience problems faced by its members. Assistance with experimental design, electronics, machining, and software development, are all services provided.

Magnetic Resonance Imaging. The CBS MRI facility in the Northwest Building has a 3-Tesla magnetic resonance imaging scanner for non-invasive human brain imaging. It also anticipates the addition of a 4.7-Tesla small-bore scanner for technique development.
Light Microscopy. CBS hosts an optical imaging facility in the Northwest Building that houses advanced commercial devices and is developing the next generation of optical techniques. This core facility provides: (a) laser scanning microscopes with motorized stages for high throughput reconstructions of the nervous system; (b) stereo fluorescence macroscopes; (c) super-resolution optical microscopes; and (d) ultra-fast optical scanning microscopes. The Harvard Center for Biological Imaging (CBI) in the Biolabs building at Harvard hosts the full line of Zeiss light microscopes, with an evergreen policy to keep pace with the newest microscopes.

Center for Nanoscale Systems. The Center for Nanoscale Systems (CNS) at Harvard provides core facilities for nanofabrication and imaging. CNS accomplishes this mission by purchasing, operating, and maintaining large, centralized scientific facilities for use by users. CNS also provides training and assistance to users. CNS is located steps away from the Northwest Building.

13. Data Management Plan

1. Project Information

This STC on Brains, Minds and Machines (CBMM) will address the grand challenge of understanding and engineering key aspects of intelligence. The Center's research will integrate cognitive science, neuroscience, computer science and artificial intelligence.

2. General Data Management Plan Information

This DMP describes how the proposed STC will conform to NSF policy on dissemination and sharing of research results.

3. Data description

This project will generate multiple data sets and resources that will be shared with researchers, teachers and students within and outside the Center. There will be 4 main types of data/resources: (1) The research efforts in the 4 main thrusts will generate experimental data (including neurophysiological recordings, functional imaging data, image and movie databases); mathematical theories (equations, proofs); computational models (equations, software); and publications, all of which will be shared. (2) The "Enabling Theory" thrust will develop a theoretical framework instantiated through computational algorithms in Thrust 3. An important core output of this proposal will be the open-source platform created in relation to the validated vision system. (3) Thrusts 1-3 will also generate data from model simulations and comparisons between models and neural circuit studies. (4) The education program will provide material

shared with the community. The table below summarizes the characteristics of the different data sets and resources and the mechanisms for management and sharing.

Title	**Thrust**	**Description**
(1) Neuroscience data (including fMRI, MEG and neurophysiology)	1-4	This resource includes functional imaging data, MEG data and neurophysiological recordings from neuronal ensembles. During experimentation, the data will reside on MIT website: http://cbcl.mit.edu/software-datasets/index.html and then be migrated to the server of the new BMM Center, where it will be accessible to others on the team. PIs: Wilson and Kanwisher.
(2) Algorithms	Theory and 2-3	The mathematical theory will be implemented in computational algorithms. The code and software will reside in http://cbcl.mit.edu/software-datasets/index.html and will be shared through the BMM Center web site. PIs: Rosasco.
(3) Model data	3	These datasets will reside in http://cbcl.mit.edu/software-datasets/index.html The data sets will include model simulations and also image datasets, movie datasets and annotations used in the research efforts. PIs: Rosasco.
(4) Education, Knowledge Transfer	Education	We will combine our online graduate and undergraduate courses developed with MITx together with the associated material (teaching outlines, student handouts, supplemental documentation, videos). We will also host the videos made to promote careers in research in high schools coming from the participation of the BMM Center in the MIT *Highlights for High School initiative.* Preliminary versions will be kept on http://cbcl.mit.edu/courses/index.html and then moved to the new Center servers. PIs: Hirsch and Mahadevan.

4. Access and Sharing

We will use two hosts for the management of data including active data sharing, hosting/dissemination of educational materials, support of the final platform, and long-term storage. The raw data and updated versions of code and software will be available to the different members of the team. The research efforts will be disseminated (see Outreach plan) and the data will be available to the research community. Data will also be available in electronic version in scientific journals. Members of our team have a strong tradition of sharing data and resources. All software tools will be available at no cost to the general public, including software packages already available at CBCL (http://cbcl.mit.edu). After publication, MRI and physiology data will be available upon request. For optogenetics tools and data, Dr. Boyden's lab has consistently made their published reagents freely available via the rapid reagent sharing non-profit Addgene (http://www.addgene.org/edward_boyden), which distributes plasmids rapidly and will continue to do so in the future. The BMM Center will also include video tutorials and white papers on how to build closed-loop optical control hardware.

5. Formats

Standard formats will be used for all the data ensuring easy read/write and communication. Metadata associated with each data set is automatically generated and stored together with the raw data files. Novel data formats will be documented.

6. Storage, backup, replication, and versioning

The host system will consist of a shared and regularly updated repository. The repository provides automatic version (revision) control over all deposited materials. All systems providing on-line storage are contained in a physically secured facility that is continually monitored. System backups are made on a daily basis. Replicas of data are held by independent archives, regularly updated, and regularly validated.

7. Security

The depository complies with MIT and Harvard requirements for good computer use practices. Both have developed extensive technical and administrative procedures to ensure consistent and systematic information security (e.g. http://security.harvard.edu/). "Good practice" requirements include system security requirements (e.g., disabling of generic accounts; inhibiting password guessing) operational requirements (e.g., breach reporting; patching; password complexity; logging); and regular auditing and review. When dealing with human data, strict confidentiality will be maintained. As stipulated in detail in the approved IRB protocols, no identifying information is kept and privacy is ensured by using alphanumeric codes.

8. Responsibility

Individuals responsible for data management are spelled out in the table above.

9. Budget

The cost of preparing data and documentation will be borne by the project and is reflected in the budget.

10. Intellectual Property Rights

The data will not be encumbered with intellectual property rights (including copyright, database rights, license restrictions, trade secret, patent or trademark) by any party (including the investigators, investigators' institutions, and data providers.); nor is subject to any additional legal requirements. Depositing data will grant permission to re-disseminate and transform the data as necessary for studies, preservation and access.

11. Adherence

Adherence to this plan will be monitored regularly. Adherence checks will include review of the DVN content, number of studies released, availability for each study of subsettable/preservation friendly data formats; availability of documentation (public); and correctness of data citation, including UNF integrity check.