Strategic Plan (updated 2015)

Center for Brains, Minds, and Machines

An NSF Science and Technology Center

Vision

Accumulated knowledge and technology that is now in place has set the stage for rapid advances in our scientific understanding of intelligence and our ability to replicate intelligence in engineered systems.

Mission

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

Goals

Our scientific goal is to discover how intelligence is grounded in computation, how these computations are implemented in neural systems, how they develop during childhood, and how social interaction amplifies the power of these computations. As we progress, we will aggressively pursue opportunities to discover and develop unifying mathematical theories.

To foster collaboration across disciplines, we will jointly develop top-to-bottom computational models powerful enough to explain visually perceived situations the way humans do. The models will emerge from fundamental questions about visually perceived situations: *who, what, why, where, how, with what motives, with what purpose, and with what expectations*. Models of visual understanding will be further advanced by developing computational models of what children know and learn about physical objects and intentional agents, and how they learn so much so rapidly. We will develop computational models of learning, memory, reasoning, and concept formation that are consistent with behavior, neural systems, and neural circuits. We will also develop computational models that enable computers to think new thoughts, imagine new scenes, form hypotheses, propose interventions, and compose narratives. Through these collaborative efforts, we will develop new methodologies and new technologies that will help to reach our goals.

Our diversity goal is to ensure that the field of Science and Engineering of Intelligence is broadly inclusive.

Our education goal is to ensure that our new knowledge is packaged in accessible ways, including model subjects at graduate and undergraduate levels.

Our knowledge transfer goal is to ensure that new knowledge is quickly and broadly disseminated and brought to bear on the great challenges of the 21st century, so as to serve the people of the nation and the world.

Focus on Turing-like Challenges

Our scientific goal is to discover how intelligence is grounded in computation, how these computations are implemented in neural systems, how they develop during childhood, and how social interaction amplifies the power of these computations. As we progress, we will pursue opportunities to discover and develop unifying mathematical theories.

Because “Intelligence” encompasses a large set of topics, we have chosen Visual Intelligence in human and non-human primates as a primary focus. Our approach to Visual Intelligence includes connections to some developmental, spatial, linguistic, and social questions. To further sharpen our focus, we are emphasizing challenges, described in more detail below, that might be viewed as inspired by the Turing test. We have dubbed these Turing^{++} Questions. Computational models we develop will be capable of responding to queries about visual scenes and movies – who, what, why, where, how, with what motives, with what purpose, and with what expectations. Unlike a conventional engineering enterprise that tests only absolute (computational) performance, we will require that our models exhibit consistency with human performance/behavior, with human and primate physiology, and with human development. The term Turing^{++} refers to these additional requirements our models should satisfy. The reason behind this additional requirement is our belief that human intelligence must be understood at several different Marr-like levels — from the level of the circuits to the levels of algorithms and computations. It is even possible that only understanding at all levels is necessary to avoid the potential problem of a possibly fake intelligence such as Ava (a shadow of millions of people behaviors captured by the Blue Book system in Ex-Machina).

One of the key challenges of a Center such as ours is to ensure added value beyond increased funding to individual investigators. We have, from the outset, insisted on a funding model that, we believe, increases that super-linear impact. We do not fund individual Center faculty. The funding actually goes to individual students and postdocs who are doing the research, under the supervision of more than a single Center faculty member. Thus, we fund collaborative students and postdocs on projects with acknowledged significance for the Center, not faculty pursuing business as usual.

Our choice of Questions follows from our understanding of human intelligence grounded in the neuroscience of the brain. We believe that <u>intelligence is one word but many problems.</u> Each Question roughly corresponds to a distinct neural module in the brain. We have begun defining an initial set of such problems/questions about visual intelligence, since vision is our entry point into the problem of intelligence. We call such questions Turing^{++} Questions because they are inspired by the classical Turing test but go well beyond it. Traditional Turing Tests permit counterfeiting and require matching only a narrowly defined level of human performance. Successfully answering Turing^{++} Questions will require us not only to build systems that emulate human performance, but also to ensure that such systems are consistent with our evolving understanding of human behavior, brains, neural systems, and development. This open-ended set of Turing^{++} Questions spans all of our thrusts and effectively measures our progress in understanding the brain-based intelligence needed to understand images and video.

Let's take some concrete examples. Consider an image as the one shown below. A deep learning network might locate faces and people. One could not interrogate such a network, however, with a list of Turing^{++} Questions such as these:

- What is there?
- Who is there?
- What are they doing?
- How, in detail, are they performing actions?
- Are they friends or enemies or strangers?
- Why are they there? What will they do next?
- Have you seen anything like this before?

We effortlessly recognize objects, agents, and events in this scene. We, but not a computer program, could recognize that this is an amusement park; several people are walking; there is a stroller in front of the fence; two women are carrying bags; very few people, if any, are riding the carousel. We, but not a computer program, could generate a narrative about the scene. It's a fairly warm, sunny day at the amusement park. A blonde young mother or caregiver in rolled-up blue jeans is waiting, presumably with a baby, by the carousel. One or two friends may be walking up to meet her.

Consider the second image. What current machine vision program could analyze this scene and determine that it shows many people, primarily young men, probably in the Philippines, escaping a flood with a few belongings, some of them by attempting to walk on electric cables?

Our brains effectively answer Turing++ Questions when interpreting the third scene, below. What is the man in the white hat holding? What is he looking at? What has just happened? We see what might be a celebrated, venerable artist in a piece of performance art—destroying a large, gilt-framed mirror in an art gallery—to the delight of his audience. We know who and what are there, and we can anticipate what will happen next. These days, we might even generate counter-factual, internal images of how the scene would differ if the location were to shift to the streets of Ferguson, Missouri, or Baltimore, Maryland.

We would assess the performance of a model built to answer questions like these by evaluating a) how similarly to humans our neural models of the brain answer the questions, and b) how well their implied physiology correlates with human and primate data obtained by using the same stimuli.

Our Turing^{++} Questions require more than a good imitation of human behavior; our computer models should also be human-like at the level of the implied physiology and development. Thus the CBMM test of models uses Turing-like questions to check for human-like performance/behavior, human-like physiology, and human-like development. Because we aim to understand the brain and the mind and to replicate human intelligence, the challenge intrinsic to the testing is not to achieve best absolute performance, but performance that correlates strongly with human intelligence measured in terms of behavior and physiology. We will compare models and theories with fMRI and MEG recordings, and will use data from the latter to inform our models. Physiological recordings in human patients and monkeys will allow us to probe neural circuitry during some of the tests at the level of individual neurons. We will carry out some of the tests in babies to study the development of intelligence.

The series of tests is open-ended; we will rely on a set of databases and add to them during the life of the Center. The initial ones, e.g. face identification, are tasks that computers are beginning to do and where we can begin to develop models and theories of how the brain performs the task. The later ones, e.g. generating stories explaining what may have been going on in the videos and answering questions about previous answers, are goals for the next few years of the Center and beyond.

The modeling and algorithm development will be guided by scientific concerns, incorporating constraints and findings from our work in cognitive development, human cognitive neuroscience, and systems neuroscience. Each Research Thrust is expected to contribute to the development of these models and algorithms and will be evaluated at the appropriate level of its contribution. For instance, the challenge of the development thrust is to create systems that represent objects the way a 3-month-old baby does; an analogous challenge in the neural circuit thrust is to develop models and theories that fit the physiological data, e.g. some part of the ventral stream. These efforts likely would not produce the most effective AI programs today (measuring success against objectively correct performance); the core assumption behind this challenge is that by developing such programs and letting them learn and interact, we will get systems that are ultimately intelligent at the human level.

Collaborative projects in the Center will have to pass the litmus test of contributing directly to demonstrable progress.

Given the focus on measuring performance and measuring progress over time, we decided that a top priority for the Center is to develop databases for measuring research progress. The same stimuli can be used to measure how well our models and our computer systems perform in absolute and relative terms. We will use data and tools to evaluate progress of our work and our theories over the years.

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

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

Expectations and Evaluation Principles

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

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

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

- **Contribution** to the Center's objectives. We aim to better understand human intelligence, to make smarter machines, and to establish a new Science and Engineering of Intelligence. Thus, participants are expected to advance our understanding of how intelligence develops in early life, how it grounds out in neural hardware, how it works at a computational level, how it rests on social interaction, and how our understanding can be magnified via unifying mathematical theories.
- **Collaboration** within and among the thrusts. We believe that seminal contributions are most likely to emerge from collaborative efforts. To increase the likelihood of success, we will operate exclusively in terms of collaborative projects between and among participants, rather than in terms of efforts limited to the research group of an individual participant.
- **Centerness.** Our Center funds only collaborative projects that cannot be done in a single lab with typical single investigator grants. All projects -- as a general rule -- should be a key

component of a thrust and pass the litmus test provided by the CBMM challenge. No single PI will have students or postdocs funded by CBMM. Instead, collaborative projects will (projects have to be collaborations between two or more PIs). Thrust leaders have the responsibility of hiring postdoc/students for the cooperative projects in their thrust -- with the help of thrust members, the Research Coordinator and the director.

- **Community** growing. We believe that our common objectives are best reached by establishing a new field of study. To further this end, we will work to encourage Center participants, and especially our students, to have broad interests and participate energetically in Center activities.
- **Commitment** to outreach. We believe that diversity is intrinsically valuable. To seize opportunities for bringing diversity into our new field, all Center faculty have committed to contribute to at least two CBMM outreach activities per year.

Respect for contribution, collaboration, centerness, commitment and community will be at the focus of the Research Coordinator thinking as he supervises the decisions made by the management team through our changing landscape.

Integrated Optimal Outcomes

Research

Outcome 1:

A computational account of the mind that interprets and describes visual scenes, and answers questions about them, the way humans do, at the *psychophysical and neural level,* thus answering a set of ***Turing^{++} Questions*** for vision.

Outcome 2:

A computational model of the mind, grounded in models of child development, that constructs intuitive theories of physical objects and intentional agents as effectively as a child, using the same kind of information that is available to a child.

Outcome 3:

Neural models of how memory, learning, and reasoning are implemented in the brain.

Education

Outcome 1:

Students well-prepared to become future research and education leaders in the new field of the Science and Engineering of Intelligence, with integrated knowledge and skills in computation, neuroscience, and cognitive science.

Outcome 2:

A model framework for education in the new field of Science and Engineering of Intelligence, including curriculum frameworks for interdisciplinary undergraduate and graduate training that are disseminated and adopted at a range of educational institutions.

Outcome 3:

Interdisciplinary courses that integrate multiple approaches to the study of intelligence that are available to students at all CBMM partner institutions, and ultimately to the broader academic community, offered by faculty at partner institutions or through an online teaching consortium.

Outcome 4:

Evidence that CBMM research, education, and HR/diversity programs have led to broader participation of women, underrepresented minorities, and other underserved groups in the new field of Science and Engineering of Intelligence.

Outcome 5:

A dedicated community of educators across CBMM partner institutions engaged in the collaborative development of courses, learning materials, and curricula related to the interdisciplinary study of intelligence.

HR/Diversity

Outcome 1:

Help faculty from computer science and neuroscience (or psychology) departments at CBMM partner institutions for the broader participation of women and minorities (BPWM) collaborate to create interdisciplinary course and joint concentration to attract increasing number of students to the field of intelligence.

Outcome 2:

Establish a robust program of workshops, faculty summer sabbaticals, and summer research experiences for undergraduates from BPWM institutions to introduce necessary skill sets, to provide intellectual and practical training in research methods, and to prepare students to succeed in graduate school.

Outcome 3:

Increase number of women and under-represented minorities who major/minor in neuroscience, Computer science or computational neuroscience.

Outcome 4:

Increase number of women and under-represented minorities who enroll in PhD programs in neuroscience, computer science or computational neuroscience. Increase number of women and under-represented minorities who enter a career in computational neuroscience.

Knowledge Transfer

Outcome 1:

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

Outcome 2:

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

Outcome 3:

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

Metrics and Milestones

Milestones for Thrust 1: Core Knowledge

Near term, 1-3 years:

- Characterize the core cognition abilities and basic learning mechanisms required for intuitive physics and intuitive psychology. Perform experiments with adults to test key claims. Develop tests for the CBMM challenge.
- Create developmental models of age appropriate behaviors for the visual CBMM challenge tests for 0-12 months
- Develop stimulus databases and experimental methods to validate models.

Mid to long term, 4-10 years:

- Towards the "what happens next" question: test intuitive physics and intuitive psychology models increasingly integrated with language, sensitive to increasingly complex and abstract physical properties and mental states.
- Create developmental models of age appropriate visual Turing test for 24-36 months.
- Develop stimulus databases and experimental methods to validate models for 24-36 months

Milestones for Thrust 2: Circuits for Intelligence

Near term, 1-3 years:

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

Mid to long term, 4-10 years:

- Compare neural circuit data (neurophysiological recordings in rodents, monkeys, humans) with behavioral data and computational models (Thrust 5) in invariant recognition of actions, objects, people and interactions among them, in making intelligent predictions about future behavior including where monkeys/humans will saccade next in the context of cluttered scenes and/or natural videos; in rodent/human navigation; and in evaluation of social interactions (monkeys/humans) (Thrust 4).
- Compare neural circuit data (neurophysiological recordings in rodents, monkeys, humans) with psychophysical data and computational models to constrain and inspire computational models (Thrust 5) in tasks that involve answering the Central Challenge questions including:
    - o “What is there?” -- Neural circuits for invariant representation of objects and people
    - o “What is the person doing?” – Neural circuits for invariant representation of actions
    - o “What will happen next?” – Neural circuits involved in predictive coding
    - o “What happened before?” – Investigate how neural circuits can support inference of causal relationships including elements of intuitive physics (Thrust 1) and social interactions (Thrust 4)
    - o “Who is doing what to whom and when and why?” – Combine elements of the above questions into a mechanistic understanding of how such intelligent inferences can be instantiated in neural hardware

Milestones for Thrust 3: Vision and Language

Near term, 1-3 years:

- Demonstrate models that can recognize objects and their parts, integrating vision and language in the domain of objects, and their properties and spatial relations, and develop algorithms for action recognition involving one agent and one object, answering the question "what is happening?" Interact with thrust 1's work on modeling developmental and learning trajectories leading to these capabilities.
- Organize regular language/vision workshops to encourage work on vision/language Turing++ questions. These workshops will feature invited speakers and poster presentations on the topic of joint vision/language tasks

Mid to long term, 4-10 years:

- Demonstrate bidirectional cooperation between vision and language in answering questions about, for example, social interactions involving multiple agents via questions such as "What is the person doing?" and "Who is doing what to whom?" and "What will happen next?" and "What do the people think of each other?". Interact with thrust 4's work on modeling of brain mechanisms involved in inferring information about social interactions from visual perception. Interact with thrust 2's work on modeling computations that use top-down neuronal circuits in these tasks.

Milestones for Thrust 4: Social Intelligence

Near term, 1-3 years:

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

Mid to long term, 4-10 years:

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

Milestones for Thrust 5: Theory for Intelligence

Near term, 1-3 years:

- Develop a theory of invariant recognition in hierarchical architectures, and develop associated neural model of the ventral stream. Test theory with respect to the "what" and "who" CBMM challenge questions on various image databases. Test theory (with thrust 2, 4 and 1) with respect to physiology and psychophysical constraints.
    - From the perspective of computer science applications of the theories above refine and test a machine learning framework for learning from very few labeled examples via unsupervised/weakly supervised learning of symmetries and other constraints from the environment. Generate open-source code.
    - Organize the first workshop on one of the Turing++ Questions. The workshop will be on "Who is there?" that is on face recognition. It will take place on Sept 3-5, 2005 at MIT.
    - Organize additional workshops on other Turing++ Questions, such as action recognition and general object recognition.

Mid to long term, 4-10 years:

- Develop a theory of visual understanding incorporating the physiology of attention and the anatomy of the back projections. Such a theory will be based on formal frameworks such as Bayesian reasoning and visual routines and will build upon our work on eccentricity dependent resolution of human vision and its relation to scale and position invariance. Demonstrate the feasibility of the associated models to answer Turing++ Question such as "What happens next?"
- Characterize and test <u>neural</u> models of inference and reasoning and of models capable of representing intentional agents and their interactions.

Milestones for Thrust 6: Seed Projects

Milestones for Education

Near term, 1-3 years:

- Develop graduate and undergraduate versions of an introductory course on the interdisciplinary science of intelligence, using material drawn from the annual summer course at the Marine Biological Laboratory (MBL) in Woods Hole (see Knowledge Transfer).

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

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

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

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

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

Mid to long term, 4-10 years:

- Offer additional CBMM courses through the online teaching consortium. Transition courses from the Small Private Online Course (SPOC) model with access limited to students and faculty at CBMM partner institutions, to the Massive Open Online Course (MOOC) model with broad access

- Adopt the graduate and undergraduate curriculum frameworks to the full range of CBMM partner institutions, and disseminate these frameworks through publications in educational journals in neuroscience and computer science

- Develop innovative simulation tools to support collaborative and integrative problem-based learning in the science of intelligence

Milestones for HR/Diversity

Near term, 1-3 years:

- Establish an annual summer program for 10 undergraduates from minority serving institutions (MSI), and non-research intensive institutions, to help prepare them for graduate school.
- Post up to eight videotaped lectures given by CBMM faculty to undergraduates from institution for BPWM on the CBMM website for broad access.
- Establish an annual 6-day workshop for 25 students and faculty from MSI and non-research intensive institutions.
- Establish on average one research collaboration per year between CBMM PIs and CBMM faculty at BPWM partner institutions.
- Work towards submitting two grants to support collaborative research or educational endeavors between CBMM faculty and faculty at BPWM partner institutions.
- Submit three co-authored publications that include summer students or faculty from BPWM partner institutions.
- Summer students will have presented up to 10 posters at national (for example ABRCMS, SACNAS, AAAS) or international meetings (for example Society for Neuroscience).
- At least three CBMM faculty members per year will visit institutions for BPWM and meet with undergraduate and graduate students.
- Two to three faculty members from BPWM partner institutions will be invited to give seminars to CBMM faculty at MIT or Harvard each year (up to nine by year 2017).
- Three faculty members from BPWM partner institutions will have spent a summer sabbatical in a CBMM PI's lab.

Mid to long term, 4-10 years:

- At least 50% of undergraduates who will have participated in the annual 6-day workshop or the summer research program and graduated from college will have applied to PhD programs.
- Three successful grant submissions involving BPWM partner institutions.
- At least one BPWM partner institution will have created a computational neuroscience concentration.
- All BPWM partner institutions will have integrated modules developed by CBMM into their curriculum.

- BPWM partner institutions will have developed strong relationships with CBMM faculty to the point of co-teaching courses and hosting student exchanges between collaborating labs.
- Increase the number of slots for summer undergraduate internships by up to 5 with support from industrial partners.

Milestones for Knowledge Transfer

Near term, 1-3 years:

- Introduce 20-25 young scientists each year to the science of intelligence
  - Host 2-week summer course at MBL every year
  - Host 1-2 scientific workshops every year

- Establish relationships with 2-3 industry partners with AI focus
  - Host one workshop with a set of companies/groups with an AI focus to explore significance and direction of AI in their industry

- Deepen relationships with 2 significant AI industrial partners
  - Host two workshops (one with a big company, one with a small company), with partners who have significant AI focus, to explore deeper relationships

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

- Strengthen academic exchange with BPWM partner institutions
  - Year 3: Each CBMM faculty member will have contributed to the Outreach program.

- Website fully functional by end of Year 1

- 1-2 public talks per year, beginning Year 2

Mid to long term, 4-10 years:

- Maintain program to introduce 20-25 young scientists each year to the science and engineering of intelligence
  - Continue to host 2-week summer course at MBL

- Expand Scientific workshop program
  - Maintain 1-2 scientific workshops each year
  - Host 1-2 international workshops per year

- Maintain program to develop relationships with 2-3 industry partners

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

- Expand program to deepen relationships with 3-4 significant AI industrial partners
    - Host two workshops (one with a big company, one with a small company), to strengthen relationships
- Maintain centeredness of CBMM
    - Host one retreat per year
        - Maintain 80+ participants
    - CBMM Weekly Research Meetings alternating between MIT and Harvard
- Maintain program to strengthen academic exchange with BPWM partner institutions
    - CBMM faculty members will make regular visits to a BPWM partner institutions
- Expand Website to include video of research, expand shared publications, databases & methodologies
- Organize symposia and workshops at COSYNE, NIPS, and other large meetings
- Stream and archive MIT-Harvard seminars
- Develop faculty exchange program within CBMM as well as industrial partners
- Promote spawning international conferences from CBMM workshops
- Promote strong industrial and government presence at annual retreat
- Maintain 1-2 public talks per year