This report is the first of a series of workshops held at the National Science Foundation that will inform possible initiatives integrating research on physical systems with research on brains. The report represents the collective views of the workshop participants and not of the National Science Foundation. The report is available on the websites of the Directorates of Biological Sciences and Social, Behavioral & Economic Sciences at the National Science Foundation and is available for public comment there.

A Report on Grand Challenges of Mind and Brain

Steering Committee
Sheila Blumstein, Brown University, Chair
Thomas Carew, University of California, Irvine
Nancy Kanwisher, Massachusetts Institute of Technology
Terry Sejnowski, Salk Institute and University of California, San Diego

Workshop Panel
Susan Carey, Harvard University
Michael Davis, Emory University
Peter Dayan, UCL
Marc Hauser, Harvard University
Michael Mauk, University of Texas, Houston
Elissa Newport, University of Rochester
Elizabeth Phelps, New York University
Michael Shadlen, University of Washington
Michael Tanenhaus, University of Rochester
Joshua Tenenbaum, Massachusetts Institute of Technology
Anne Treisman, Princeton University
David Van Essen, Washington University

Executive Summary

The time is ripe to tackle some of the most fundamental questions (i.e., “Grand Challenges”) of mind and brain, where a concerted effort in several problem domains can be expected to yield major progress in the next few years. The workshop participants focused on both conceptual questions and experimental approaches in addressing the general issue of the relationship between mind and brain.

The Scope of the Initiative:

The panel felt that the scope of the initiative could best be captured by identifying a set of large-scale domains that serve as exemplars of the kinds of questions and themes that are ripe for exploration. The intention of these examples is to illustrate the scope of the questions the panel would encourage; they are not meant to be restrictive. In each domain, a series of questions that bear on its theme are articulated. The domains include:

- Adaptive Plasticity
- Conflict and Cooperation
- Spatial Knowledge
- Time
- Language
- Causal Understanding

In addition, several fundamental and overarching questions were identified that cut across the example domains.

Experimental Approaches that should be applied:

It was the unanimous view of the panel that the broad scope of the exemplar domains requires a wide variety of experimental approaches coupled with multiple levels of analysis. Thus, this effort should involve the application of research methods across an extensive range of disciplines including chemistry, physics, mathematics, psychology, cognitive science, neuroscience, linguistics, engineering, and computer science.

The panel also discussed new tools that would be needed to increase our understanding of the mind and brain. These include the mapping of human brain circuitry, mathematical innovations, the development of information databases, the use of molecular tools, and the development of a new cyberinfrastructure.

Structure of the Initiative:

In evaluating proposals for this initiative, the panel encourages the NSF to create a reward structure that promotes:

- **Multidisciplinarity**. Since the questions of mind and brain span many levels of analysis (behavioral, computational, systems, neurophysiological, and molecular), the initiative should support innovative combinations of techniques and levels of analysis. Moreover, modern technology permits interdisciplinary initiatives unbounded by physical proximity (e.g. *virtual universities*). The multidisciplinary approach taken in the study of the mind and brain also promises to radically transform neighboring areas of study. For example, the behavioral methodologies of cognitive science can provide critical contributions to the work of molecular biologists as they endeavor to understand the behavioral consequences of gene expression.

- **The importance of young investigators.** The panel placed a high value on the inclusion of young investigators in all aspects of the initiative. Thus, we would encourage every application to include young investigators who will be able to take conceptual and experimental risks not typical at early stages of a career.
- **Broad questions, restricted clusters.** The panel felt that this initiative could be best served by many small-scale partnerships rather than by a few large-scale clusters.
- **The role of education.** A unique feature of the NSF is its commitment to quality and innovative science while at the same time placing that scientific inquiry in an educational context. The panel encourages the NSF to incorporate this value in the initiative.

On July 18-19, a workshop on **Mind and Brain: Strategies and Directions for Future Research** was held at the National Science Foundation. This workshop brought together scientists across many disciplines to assist in formulating a niche for funding new, promising work in cognition and neuroscience.

The question of the relationship of mind and brain has captured the imagination of philosophers, psychologists, and biologists for centuries. The 21st century has witnessed an unprecedented growth in both technical and conceptual advances in the study of the mind and the brain. We now find ourselves in an unparalleled position to make major progress in answering longstanding and fundamental questions (i.e., "Grand Challenges") abound mind and brain. This workshop reflects the first step in the commitment to expanding the horizons and pushing the frontiers of this critical domain of scientific inquiry.

The workshop focused both on scientific questions where a concerted effort can be expected to yield major advances in the next few years, and on the approaches and technical tools that need to be used in the investigation of these questions. The goals of the workshop were to identify topics and questions where strong inroads can be made through specific investigator-initiated projects. By inspiring ambitious new directions of scientific inquiry, this initiative will promote a deeper understanding of ourselves and the way we interact with the world. Our report describes the directions that we envision as optimizing this research agenda and outlines examples of questions that deserve further consideration.

The Scope of the Initiative

The challenge of this initiative is to promote inquiry that pushes the frontiers of understanding in ways that cannot be achieved with a single level of analysis or a single experimental approach. After lively discussion, we agreed that this approach could be best advanced by discussing large-scale questions that serve as exemplars of the kinds of questions that can be embraced by the initiative. To this end, we include in the report examples of questions that capture key elements of the overall challenge of the relationship between mind and brain. The intention of these examples is to illustrate the scope of the questions we would encourage. They are not meant to be restrictive.

The domains of research described below include:

- **Adaptive Plasticity**
- **Conflict and Cooperation**
- **Spatial Knowledge**
- **Time**
- **Language**
- **Causal Understanding**

Several fundamental and overarching questions cut across the example domains outlined in this report. How do we solve the fundamental problem of induction in perception and thought to make inferences when insufficient information is available? How are mental representations such as percepts, concepts, and verbal ideas related to each other and how do different parts of the brain communicate to integrate information from different modalities to reach decisions and to order our behavior into coherent sequences? How are concepts and thoughts coded in neural activity? How do these cognitive and neural representations change dynamically as we learn? These broad questions will be illuminated by the more specific investigations undertaken in each of the specific research domains funded by this initiative.

The Experimental Approaches That Should Be Applied

The scope of the challenge of understanding the human mind requires combining a wide variety of experimental approaches with multiple levels of analysis. The core method for understanding the mind is testing behavior. Behavioral work provides both a characterization of the fundamental mental phenomena to be explained and a set of paradigms that allow investigation of those phenomena in combination with other methods such as brain imaging and computational modeling. For example, the behavioral consequences of a gene knockout cannot be adequately described by simply using a handful of standard off-the-shelf paradigms (e.g. the Morris water maze). Studies of gene-behavior links require sophisticated behavioral methods that can provide a rich characterization of specific cognitive phenotypes. The goal of understanding the relation between mind and brain requires collaborations across a wide range of disciplines including chemistry, physics, mathematics, psychology, cognitive science, neuroscience, linguistics, engineering, and computer science. There is also a pressing need to develop new tools, methodologies, and computational models for studying mind and brain. These tools are discussed in a final section below.

The Structure of the Initiative

Several elements emerged as critical to the success of the initiative. Thus, we encourage the NSF to create a structure that rewards the inclusion of these elements:

- **Multidisciplinarity**. Since the questions of mind and brain span many levels of analysis (behavioral, computational, systems, neurophysiological, molecular), the initiative should support innovative combinations of techniques and levels of analysis. Moreover, modern technology permits interdisciplinary initiatives unbounded by physical proximity (e.g. *virtual universities*). We do not presume to prescribe any particular combination or level of analysis. Rather we would hope that NSF would encourage and support imaginative proposals that would capture the spirit of these recommendations.

- **The importance of young investigators.** The panel placed a high value on the inclusion of young investigators in all aspects of the initiative. Thus, we would encourage every application to include young investigators who will be able to take conceptual and experimental risks not typical at early stages of a career.

  Broad questions, restricted clusters. The panel felt that this initiative could be best served by many small-scale partnerships rather than by a few large-scale clusters. The research effort will be optimized by allowing a focused approach to each particular question.

- **The role of education.** A unique feature of the NSF is its commitment to quality and innovative science while at the same time placing that science in an educational context. The IGERT program exemplifies this integration of science and education. We encourage the NSF to incorporate this value in the initiative.

Exemplars of the Domains of Research

Adaptive plasticity

All behaviors reflect an interaction between an organism and its world: to optimize its success in an ever-changing environment, the organism must constantly fine-tune its behavior. This in turn requires that the brain be plastic, constantly shaping itself to meet the challenges of this changing world.

This plasticity is perhaps most striking as the brain first develops: in its initial formation, the brain is profoundly shaped by the nature of the animal's environment. But plasticity does not end there; it continues throughout the life of the animal, allowing it to constantly update its behavior in an adaptive fashion. This plasticity takes place at a variety of levels, from the synaptic interactions among single neurons and the circuits in which neurons interact, to the large-scale systems comprised of those circuits. The analysis of plasticity at any of these levels can inform the general issue of how we encode information about the world, how we store it, and how we access that information at a later time. This type of analysis can, in turn, allow us to explore a number of fascinating questions that capture critical features of the human experience. Why can two people witness exactly the same event, yet remember it so differently? Why can we sometimes remember what a word rhymes with before we can remember the word itself? How do children encode the experiences in their world? Do they use the same strategies as adults or engage a different set of mechanisms? Why are some memories fleeting while others endure a lifetime? Why are some memories encoded immediately, after a single experience, while others require extensive repetition over a prolonged period of time? When we try to remember something, but can't quite bring it to the surface, why does the memory suddenly pop into our minds at a later time? These are but a few examples of the rich diversity of questions and issues captured in the study of learning and memory.

A pressing question, then, is: What are the rules and principles of neural plasticity that implement these diverse forms of memory? Attacking a question of this magnitude requires an arsenal of techniques, from molecular and cellular approaches to systems-level analysis, to computational analyses, and ultimately to an understanding of how plasticity is played out in behavior. Moreover, there are clear biological constraints on these processes of plasticity; not all experiences result in changes in the brain. For example, species clearly differ in the stimuli that most readily provoke plastic changes. A complete understanding of plasticity thus requires not only a multi-level analysis of how plasticity occurs, but also an analysis of its constraints. Important insights have already been provided by exploring plasticity in a wide range of species, from fruit flies and mollusks, to birds and primates, and including, of course, the analysis of human memory. Finally, perhaps the greatest challenge of all in the general field of adaptive plasticity is directly relating any form of plasticity observed in any nervous system to *bone fide* cases of learning and memory expressed in behavior. This is a difficult task, but the combined approaches discussed above collectively provide exciting paths toward achieving this goal. In summary, the exploration of mechanisms of plasticity will serve to integrate diverse fields of scientific inquiry and will provide a vehicle to gain insights into our remarkable capacity to learn and remember.

Conflict and Cooperation

One of the greatest threats facing the world today is the conflict between people. In spite of the importance of this problem, the cognitive and neural systems underlying when and how we cooperate or engage in conflict are poorly understood. Although there have been psychological investigations of the influence of social forces on behavior, this research has rarely been informed by the sort of sophisticated and interdisciplinary models of neural function and cognitive representation and their interactions that are being developed to understand individual decision-making. An integrated scientific understanding of competing emotional systems as they guide behavior could well lead to new concepts to facilitate cooperation between individuals, and perhaps societies as a result of conflict reduction. For example, animal models can be informative about links between stress and aggression (e.g., fish crowding). Hence, an interdisciplinary approach including detailed animal models, engineering, statistics, economics, ethology, neuroscience, social and cognitive psychology needs to be brought to bear on this problem.

Such an approach allows for answering a broad range of questions. For example, what is the interaction of different emotions at behavioral, systems, cellular and molecular levels? That is, when and how does fear lead to aggression? When is aggressive behavior adaptive and when is it maladaptive? How is the range of social and emotional information that leads to choices integrated? For example, what are the mechanisms that lead an animal to choose to punish a conspecific at a personal cost? How do social forces change the neural and cognitive representation of response selection? For example, how does the presence of a social group alter an individual's sense of personal responsibility, and cause that individual to behave in ways that would never be done in isolation? What is the neural and cognitive representation of others, i.e. conspecifics, and how does this enable trust, empathy, learning through imitation, and social cooperation?

Spatial Knowledge

Where is Iowa? What is the shortest route to the post office? Where did I put my keys? How do I thread a needle? What is the relationship between our sense of space and our ability to solve the Pythagorean Theorem? We take for granted our ability to know where we are in the world and to interact with our surroundings. However not everyone has a good spatial sense, and there are some, such as those who are blind, who "see" the world in ways that are entirely different from ways used by those who are sighted. Rodents that burrow, bats that echolocate in the dark, electric fish that have a sense we cannot even imagine, all live in the same three dimensional world that we inhabit, and need to solve similar spatial and navigational problems.

Space is deeply embedded in the way that we use language. Objects are on top of, inside of, nearby, hidden by and illuminated by other objects. These are symbolic relationships that we use metaphorically when we say we are "on top of" a problem, or "near" to finding a solution. How does our brain represent the world around us? How do we use these representations to plan our actions and find our way to a goal? What is the relationship between space and language?

When we look into the brain, we find dozens of different types of maps of the world: In the early stages of visual processing the maps are based on retinal coordinates, but at later stages the maps become more abstract. For example, in the hippocampus, an area known to be important for long-term memory, single neurons called "place cells" represent where an individual is in the world. In rats, these cells fire when the animal is asleep, replaying the day's journey; they also fire when the animal is learning a task, replaying successes and failures. How do animal models scale up to humans who not only navigate by means of dead reckoning and landmarks, but use maps, verbal descriptions, and GPS devices?

Although we have discovered many cells that represent information about where we are in the world, it is a mystery how we use these neurons to think about space. Damage to these regions of the brain lead to bizarre symptoms, such as neglect of the part of the world these neurons once represented, including parts of one's own body. Considerable evidence suggests that that the brain systems that evolved to help us keep track of space and geometry made it possible for us to do abstract mathematics. Our sense of space gave rise to some of the greatest human achievements in architecture and art, from the engineering of the pyramids and the Eiffel Tower to the depth of Monet's landscapes and Picasso's cubism. The exploration of "inner space" in our brain is a frontier as mysterious and as important as our exploration of outer space.

Time

Over the next decade, the neuroscience community can provide a systematic understanding of how the brain senses the passage of time. Time connects us to the past, allows us to infer causes and distinguish them from effects. Time underlies the association of actions with outcomes; time allows us to learn the consequences of sensory stimuli. We all have a sense of the passage of time. Understanding the neurobiology of time is likely to revolutionize the study of learning and memory. The representation of elapsed time underlies anticipation, preparation and sequences of behavior. It establishes a lattice that allows us to parse information streams and produce them (e.g., language). It is at the heart of rhythm, music, dance, poetry. Elapsed time colors the value associated with behavior. Time mediates the tradeoff between speed and accuracy and is critical for understanding delayed gratification, impulsivity and addictive behavior.

Elucidating the neurobiology of time is an integrative program. It involves brain regions we do not understand at the integrative level: cerebellum, basal ganglia, and association cortex. It will inspire development of novel techniques to study functional connectivity—tools that will be widely useful in neuroscience. It taps into an existing computational and behavioral knowledge base, but it demands innovation: new paradigms; new theoretical insights. The neurobiology of time will require the development of new mathematical techniques: inference based on dynamical processes, stopping rules.

Language

Language sets us apart from all other creatures. Language allows us to say who and what we are, and is the primary way we communicate with others. Language enriches personal experience, helps us recall the past, plan the future and through written language, accumulate knowledge over generations.

New techniques for noninvasive brain imaging, new ways to analyze electrical and magnetic brain signals recorded from outside the brain, and recordings from the human brain during neurosurgery promise new insights into the biological basis of language. Language was a late-comer in evolution, capitalizing on preexisting systems. Many aspects of language, including the sensory, motor and learning components, can therefore be studied in animal models including birds, rodents and primates.

Speaking and understanding requires us to rapidly articulate and decode speech sounds with great precision. We can easily and quickly process the meaning of speech sounds, words, syntactic relationships, and disambiguate multiple possible meanings from context. Emerging insights into rapid auditory and visual processing, as well as motor control systems in a variety of species promise to help us understand how these abilities may have been adapted for human speech and language. How does the ability to speak develop during the first years of a child's life? Comparisons and contrasts with birdsong learning may prove illuminating in this regard. How has the auditory system adapted to represent the sounds of speech? How are words and their meanings, including abstract words such as 'liberty' and 'marriage', represented in the brain?

Languages can be learned by children without any formal education at an astonishing speed. Although there is a wide range of mutually unintelligible languages, they share some universal features. What is the biological basis of these features, and how do they arise? Adults can also learn new languages, although not with the same proficiency. What are the principles of language learning and plasticity? Are they domain specific or domain general?

Language is communicated across multiple modalities. What parts of the visual system are involved in learning to read? How is sign language learned by deaf signers? How is Braille represented in the somatosensory system of blind people who have learned to read with their fingers? What are the computational and neural systems underlying these systems?

The human and chimpanzee genomes have been sequenced, making it possible to begin to track down how language might have evolved. Human genomics can be used to help identify mutations that affect and underlie different aspects of language.

Probabilistic models of language and computation are providing important insights into how language is structured, learned, and processed. Recent models of cortical processing have converged on related probabilistic descriptions, which may for the first time allow us to develop a biologically based model of language acquisition and processing.

A targeted interdisciplinary study of language, drawing from engineering and computer science, linguistics, biology, psychology, neuroscience, and cognitive science, can result in breakthroughs within the next decade in our understanding of the biology of language, one of science's most enduring mysteries.

Causal Understanding

Causal understanding is central and pervasive in cognition. Causality is at the core of human concepts, intuitive theories and mental models, linguistic meaning, perception, judgment and decision-making. The importance of causality has long been recognized in psychology, but the topic has recently drawn much broader interest due to the development of new methods for studying causal inference and learning in computer science, statistics, animal behavior, and neuroscience. The most interesting questions remain open. We see and hear the world in terms of meaningful causal interactions -- a vase crashing to the ground, a father beckoning to his daughter – rather than raw lights and sounds. What is the role of causal representation and inference in perception, and in bridging perception to cognition? What is the role of causal inference in how we understand other human minds -- people's beliefs, goals, plans, and social interactions? Basic kinds of physical causal interactions can be directly perceived.

What is the connection between this perceptual causality and higher-level causal reasoning? Is there a single concept of causation that cross-cuts all domains of causal understanding – including intuitive theories of physics, psychology, biology -- or is causal understanding in each core domain essentially different, depending on domain-specific causal mechanisms? What are the origins of human causal understanding, over evolution and development? What is the neural basis of causal understanding? What is the relation between human causal reasoning and the basic mechanisms of conditioning, as studied in animals and in the neurophysiology of learning and memory? How can sophisticated computational frameworks for causal reasoning and learning be mapped onto -- and perhaps enriched by -- models of neural computation?

New Tools for a Science of Mind and Brain

Progress in the science of mind and brain will require the development of a new tool set for experimental and theoretical investigations. By promoting directed technical developments, the NSF can play a pivotal role in fostering brain research over the next decade with many potential dividends in related fields such as information technology. Investigations of mind and brain will benefit from and contribute to CyberComputing innovations. Here we suggest a few specific examples of experimental and theoretical tools.

Human Brain Circuitry. Information about the intricate circuitry of the human brain (its ‘wiring diagram’) is fundamental to understanding the higher brain functions that make us uniquely human – our thoughts, emotions, and communicative abilities. Moreover, elucidating how brain circuits differ from one individual to another is essential for discovering what makes our personality and our behavior patterns unique to each individual. Whereas much is known about neural circuits in laboratory animals, including nonhuman primates, comparable information about the human brain is woefully lacking. Novel methods for in vivo analysis of brain connectivity are currently under development and show significant promise. These include diffusion tensor imaging, functional connectivity analyses, in vivo manganese tracing, and transcranial magnetic stimulation (TMS) as well as human lesion studies combined with functional neuroimaging (fMRI). Additional investment in the development, validation, and utilization of such tools may lead to major advances in deciphering the wiring of the human brain. In turn, knowledge of human brain connectivity will be crucial for understanding how different modules of the brain interact dynamically to produce complex adaptive behavior.

The NSF can also facilitate progress on the human 'connectome' by promoting similar efforts in nonhuman species. It would be useful to develop web based tools and search engines for investigators to disseminate information to develop connectional analogies between species and to facilitate development of hypotheses and validation techniques in the human connectome.

Mathematical innovations. There is a need for new computational and mathematical tools that will inform a neuroscience of inference and theory construction. The NSF can capitalize on recent progress in machine learning, hierarchical Bayesian inference, estimation theory, "smart" search algorithms, and dynamical systems theory by promoting exchanges between communities of scientists active in information technology and applied mathematics and neuroscientists and behavioral psychologists. Innovation is likely to come from the field, but we anticipate developments in new areas of probability theory and statistical inference in dynamic systems with non-stationary distributional assumptions. These developments would have dividends for economic sciences, AI, information technology and prediction (e.g., meteorology). The NSF is uniquely positioned to define the vision and to create the incentives for bridging scientific communities that exploit prediction and inference.

Information databases. There is a need for expanded neuroscience-related databases, including those that catalog the information that the brain senses and manipulates. In the domains of language and vision, there is compelling evidence that organisms are remarkably sensitive to fine-grained properties of the input. In the limited domains where we have some knowledge of the input, computational techniques have changed the way we understand how the organism extracts structure and learns from the input. These fields would benefit immensely from rich databases that provide more veridical descriptions of the input to the organisms. The information would have a catalytic effect on the field of visual development, language acquisition, and concept acquisition to name few examples.

In addition, neuroscientists need improved access to the vast amount of information that is reported in scientific publications and to the even greater amount of useful experimental data that underlies these publications. Neuroscientists would benefit greatly from Google-like search capabilities for text and image data within publications and to improved and expanded databases that provide access to experimental data about brain structure, connectivity, function, development, and pharmacology. This would allow investigators to obtain the best available answers to questions such as: What are the inputs and outputs of brain area X and the receptor and protein distribution in area X? The neuroinformatics tools needed for this type of efficient data mining can best be generated by teams that combine computer science, neuroscience, and informatics expertise.

Integrating molecular tools into systems neuroscience. The NSF should promote the integration of molecular tools in the study of higher brain function. Animal models offer the most promise for conducting physiological measurements of brain function in animals during sophisticated behavioral tasks. However, two technical limitations curtail progress in this area. There are at present limited means to selectively manipulate brain circuit components, and there is no efficient technique to identify neurons for electrical recording based on their logical position in a circuit. For example, the field of cognitive neuroscience would be revolutionized by an ability to record from those neurons in the prefrontal cortex that project to the grasp region of the parietal lobe. Emerging technologies in molecular biology, virology and nanotechnology will permit introduction of genes, RNAi, and nanoparticles into groups of neurons with unique promoters or on the basis of their axonal projections. These techniques will be developed by molecular biologists and cellular neuroscientists through a variety of funding sources. The NSF can play a pivotal role in promoting the integration of these tools into systems level neuroscience. NSF can promote joint research programs to bring molecular biologists to cognitive neuroscience laboratories conducting neural recordings from awake animals.

Cyberinfrastructure. Computation is a vital resource for research on minds and brains. Neuroimaging experiments and recordings from thousands of neurons simultaneously are generating huge datasets that require large-scale computing to analyze the data. Behavioral experiments that use natural sensory and motor environments require real-time interactive computer systems. Petascale cyberinfrastructure is equally important for developing the next generation of neural models based on the detailed properties of neurons and their connectivity as well as high-level computational models of intelligence. The bridge between these two levels of description will require researchers with mathematical and computational skills to work closely together with experts with domain knowledge in the grand challenge problems.