

Steering Group Report:

**Brain Science as a Mutual Opportunity for the Physical and
Mathematical Sciences, Computer Science, and Engineering**

**National Science Foundation
August 21-22, 2006
Arlington, VA**

Any opinions, findings, conclusions, or recommendations expressed in this report are those of the participants and do not necessarily represent the official views, opinions, or policy of the National Science Foundation.

Steering Group Report: Brain Science as a Mutual Opportunity for the Physical and Mathematical Sciences, Computer Science, and Engineering

National Science Foundation, August 21-22, 2006, Arlington, VA

Summary

Humankind now stands at a special moment in its long history of thinking about the brain, a moment of revolutionary change in the kinds of questions that can be asked and the kinds of answers that can be achieved. Fundamental shifts include:

- **The Scope and Scale of Experimental Investigations:** Instead of one- or few-at-a-time measurements, it is becoming possible to measure brain structure, chemistry, and activity simultaneously at many locations with high specificity and spatial/temporal resolution.
- **The Character of Theoretical Understanding:** Instead of mainly bottom-up or top-down models and theories, it is becoming possible to formulate comprehensive multi-scale models that are both bottom-up and top-down and include relevant dynamics at different spatial and temporal scales.
- **The Ways in Which Knowledge Can Be Used:** Applications for the emerging multi-disciplinary knowledge about the brain abound: In large-scale neural simulations, in robots and other engineered systems that mimic biological systems, and in brain-computer interfaces that enable bi-directional communication for next-generation neural prostheses.

In this time of change there are significant unexploited opportunities for mutual scientific benefit between brain science and the physical and mathematical sciences, computer science, and engineering.

Four broad areas of opportunity were identified: Because of its strong record of leadership in the physical and mathematical sciences, computer science, and engineering, NSF is well-positioned to enable and exploit the following opportunities:

- **Opportunities in Instrumentation and Measurement:** New instruments, probes, and experimental tools are needed for comprehensive measurement of the structure, chemistry, and activity of individual nerve cells and neural populations in functioning neural systems. Such tools will permit vastly improved experimental studies of neural dynamics that accompany development, learning, cognition, and behavior.
- **Opportunities in Data Analysis, Statistical Modeling, and Informatics:** The availability of immense quantities of high-resolution data in turn will demand new statistical tools and models, and new informatics capabilities for storage, representation, and modeling of high-throughput multi-resolution data. New approaches for inferring association, linkage, and causality will be required.
- **Opportunities in Conceptual and Theoretical Approaches:** Advances in analysis and modeling of comprehensive multi-scale data will enable the exploration of much richer conceptual and theoretical approaches to understanding the brain at all levels. New mathematical approaches to understanding very high-dimensional, non-linear, non-stationary, multi-scale systems will be required.
- **Opportunities in Building Brain-like Devices and Systems:** Improved understanding of the brain, combined with advances in engineering capabilities, will permit revolutionary advances in neurally-inspired computing and information processing, in the design of robots and other engineered systems that mimic biological capabilities, and in brain-computer interfaces that enable bi-directional communication with the brain in real time.

These opportunities are discussed in detail in the following sections of this report, as are their implications for science education and for science organization.

Participants

C. C. Wood, Chair
Vice President
Santa Fe Institute
ccwood@santafe.edu
<http://www.santafe.edu>

Theodore W. Berger
David Packard Professor of Engineering
Professor of Biomedical Engineering and
Neuroscience
Director, Center for Neural Engineering
University of Southern California
berger@bmsr.usc.edu
<http://neural-prosthesis.com/>

William Bialek
John Archibald Wheeler/Battelle Professor in Physics
Princeton University
wbialek@princeton.edu
<http://www.princeton.edu/~wbialek/wbialek.html>

Kwabena Boahen
Associate Professor
Department of Bioengineering
Stanford University
boahen@stanford.edu
<http://www.stanford.edu/group/boahen/>

Emery N. Brown
Professor of Computational Neuroscience
Professor of Health Sciences and Technology
Massachusetts Institute of Technology
Associate Professor of Anaesthesia
Harvard Medical School
enbrown1@MIT.edu
<https://neurostat.mgh.harvard.edu/brown/emeryhomepage.htm>

Todd C. Holmes
Associate Professor
Department of Biology
New York University
todd.holmes@nyu.edu
<http://www.nyu.edu/fas/biology/faculty/holmes/index.html>

Nancy Kopell
William Goodwin Aurelio Professor
of Mathematics and Science

Boston University
nk@bu.edu
<http://cbd.bu.edu/members/nkopell.html>

Alan Leshner
Chief Executive Officer
American Association for the Advancement of Science
Executive Publisher, *Science*
aleshner@aaas.org
http://www.aaas.org/news/press_room/leshner-bio.shtml

Tom Mitchell
Fredkin Professor of AI and Machine Learning
Chair, Department of Machine Learning
Carnegie Mellon University
tom.mitchell@cs.cmu.edu
<http://www.cs.cmu.edu/~tom/>

Partha Mitra
Cold Spring Harbor Laboratories
mitra@cshl.edu
<http://mitralab.org/main/partha.html>

Lois Peters
Lally School of Management and Technology
Rensselaer Polytechnic Institute
peterl@rpi.edu
<http://www.rpi.edu/~peterl/>

Sara A. Solla
Department of Physics and Astronomy
Department of Physiology
Northwestern University
solla@northwestern.edu
<http://www.northwestern.edu/nuin/fac/solla.htm>

Charles F. Stevens
Molecular Neurobiology Laboratory
Salk Institute for Biological Studies
stevens@salk.edu
<http://www.salk.edu/faculty/faculty/details.php?id=50>

Jonathan V. Sweedler
Department of Chemistry and Biotechnology Center
Beckman Institute and Institute for Genomic Biology
University of Illinois, Urbana Champaign
jsweedle@uiuc.edu
<http://www.chemistry.uiuc.edu/gswee.htm>

NSF Organizing Committee

Semahat Demir, Program Officer
Division of Bioengineering and Environmental
Systems
Directorate for Engineering
sdemir@nsf.gov

Christopher Greer, Program Officer
Division of Biological Infrastructure
Directorate for Biological Sciences,
and Office of Cyberinfrastructure
cgreer@nsf.gov

Janice Hicks, Executive Officer
Division of Chemistry
Directorate for Mathematical and Physical Sciences
jhicks@nsf.gov

Mary Ann Horn, Program Officer
Division of Mathematical Sciences
Directorate for Mathematical and Physical Sciences
mhorn@nsf.gov

Rae Silver, Senior Advisor
Office of Integrative Activities
Office of the Director
rsilver@nsf.gov

Kenneth Whang, Program Officer
Division of Information and Intelligent Systems
Directorate for Computer and Information Science and
Engineering
kwhang@nsf.gov

Workshop Agenda

Monday August 21

8:00 - 8:15 AM

Welcome

Dr. Kathie Olsen, Deputy Director, NSF

8:15 - 8:30 AM

Goals of the Workshop

8:30 - 10:00 AM

10-min Presentations of “Opportunities Homework” Assignments by Steering Group members

10:00 - 10:30 AM

Break

10:30 - 11:30 AM

10-min Presentations of “Opportunities Homework” Assignments by Steering Group members

11:30 AM - 12:30 PM

Reactions, Discussion, Determine Structure of Subsequent Deliberations

12:30 - 1:30 PM

Working Lunch

1:30 - 3:30 PM

Continuation of General Discussion
Organization of Sub-Groups of 2-3 for Writing

1:30 - 3:30 PM

Sub-Group Writing

3:30 - 4:00 PM

Break

4:00 - 5:00 PM

Sub-Group Writing

5:00 - 6:00 PM

Status Report by Sub-Groups

Tuesday August 22

8:00 - 8:30 AM

Summarize and Prepare Discussion to be held with Drs. Bement and Olsen, NSF

8:30 - 9:30 AM

Discussion of Opportunities with Dr. Arden Bement, Director, NSF, and Dr. Kathie Olsen, Deputy Director, NSF

9:30 - 10:00 AM

Break

10:00 AM - 11:00 AM

Complete Sub-Group Writing

3

11:00 AM - Noon

Begin Integration of White Paper

Noon - 1:00 PM

Working Lunch, Continue Integration

1:00 - 2:00 PM

Complete Integration

2:00 - 3:00 PM

Read-Back, Critique, and Finalize White Paper

3:00 - 4:00 PM

Discuss Organization and Possible Participants or the Second Workshop

4:00

Adjourn

Contents

Summary	1
Participants	2
NSF Observers	3
Agenda	4
Contents	5
1. Introduction	6
2. Opportunities for Mutual Benefit	7
2.1 Opportunities in Instrumentation and Measurement	7
2.2 Opportunities in Data Analysis, Statistical Modeling, and Informatics	9
2.3 Opportunities in Conceptual and Theoretical Approaches	11
2.4 Opportunities in Building Brain-like Devices and Systems	13
3. Implications and Opportunities for Science Education	15
4. Considerations for Implementation: Science Organization and “Organizing for Repeated Innovation”	16
5. Conclusions	17
Appendices	19
A. Suggested Readings	19
B. Biographical Sketches of Participants	28

1. Introduction

The brain has long captured human curiosity. What are the origins of our perceptions, thoughts, intentions, and actions? How can we accomplish such complex tasks as recognizing the face of a friend across a crowded room, catching a ball, playing a musical instrument, learning our native language, writing a poem? Even creatures with much simpler brains can solve strikingly difficult problems: the acrobatic flights of birds and insects, fish finding their way home each spawning season, or bees communicating the location of a food source. Understanding how all of this is possible -- *how the brain generates meaningful behavior* -- remains one of the great frontiers of science.

Science now stands at a special moment in humankind's long history of thinking about the brain, a moment of revolutionary change in the kinds of questions that can be asked and the kinds of answers that can be achieved. This revolution is possible in part because of a change in the nature of collaboration across established scientific disciplines. Traditionally, psychologists and biologists have asked questions about brain function, relying on engineers and physical scientists to provide instruments to help answer them. This model of interaction remains strong and productive. However, physicists, chemists, computer scientists, mathematicians and engineers are increasingly asking their own questions about the brain, and in doing so, are reshaping the intellectual and scientific landscape. The goal to understand the brain is thus becoming a core challenge for many disciplines. The consequences for the physical sciences, mathematics, computational science, and engineering will be enormous, as will the implications for education and for economic competitiveness.

These changes in the scientific landscape reflect fundamental shifts in three broad areas:

A Shift in the Scope and Scale of Experimental Investigations: In the past, experiments typically focused on a single type of molecule in the brain, the electrical activity of a single neuron, or the connections from one cell to the next. Advances in chemistry, molecular biology, physics and engineering have allowed scientists to move beyond this "one at a time" approach. Thus, it is progressively becoming possible to catalog all the molecules involved in a particular signaling pathway, to record the activity of hundreds of neurons simultaneously, or to diagram a complex neural circuit completely. These increases in the scope of experimental measurements are paralleled by corresponding increases in the requirements for data acquisition and analysis, and in the scope and complexity of the mathematical/computational models required to organize and provide a preliminary understanding of the data collected. This shift toward a more complete view of the brain's internal workings is paralleled by a richer view of the behaviors to be explained. Rather than investigating limited sets of proscribed behaviors, new high-resolution measurement techniques make it possible to investigate complex behaviors over long periods of time as they occur naturally and spontaneously.

A Shift in the Character of Theoretical Understanding: Theories and models of the brain have been limited by the "one at a time" measurement constraints just discussed and by computational constraints that have prevented truly large-scale and comprehensive models of neural circuits and systems. Today, these methodological and computational constraints are beginning to be overcome, freeing scientists to focus on novel theoretical approaches to understanding the brain. In the past, models of brain function have tended toward either "bottom-up" or "top-down" strategies. Scientists can now envision models that are simultaneously both "bottom-up" and "top-down", and that can provide an integrated description across the many spatial and temporal scales on which brain function unfolds. From the "bottom up", scientists can ask how the complex functional behavior of the system emerges from its microscopic activity, providing profound challenges to dynamical systems modeling, statistical physics, and related disciplines. From the "top down", theorists from many different backgrounds are articulating global, functional principles from which one can hope to derive aspects of neural dynamics and architecture. The integration of these approaches offers the hope of a truly predictive theory of the brain.

A Shift in How Knowledge Can Be Used: Achieving “machine intelligence” is a longstanding ambition, but until recently the computing power that could compare even with the brain of a small insect was unavailable. Now, the remarkable increases in computing power in the last decade have opened unprecedented opportunities for neural simulation, emulation, and brain-based technologies. By exploiting and advancing the leading edge of computing technology, scientists can begin to simulate the structure and function of larger, more complex, and more comprehensive neural circuits and systems. An improved understanding of the fundamentals of neural information processing will rapidly advance progress toward the goal of genuine machine intelligence. Technological advances are also beginning to provide sufficient bandwidth and computational power to achieve interactive communication between brains and computers. These developments will have profound implications for neural prostheses, for robotics, and perhaps even for our everyday work environment. The continuing size decrease of semiconductor circuits means that electronic circuits will soon begin to exhibit irreducibly stochastic behaviors, much as those of ion channels and other biological phenomena. Understanding how the brain computes with fluctuating elements may help us understand how to work with comparably fluctuating transistors and logic gates in silicon. Finally, a more fundamental understanding of brain mechanisms of plasticity and learning, coupled with corresponding advances in cognitive science, could form the basis of a richer and more biologically based approach to teaching and learning.

Exciting as these opportunities are, significant challenges must be met if the potential for revolutionary scientific change is to be realized. Chief among them is the construction of a collaborative culture that combines the great depth of the traditional disciplines with the breadth, cross-fertilization, and innovation that comes from close cross-disciplinary collaboration, communication, and education. Strong interdisciplinary training programs at both the undergraduate and graduate levels, focused on and motivated by the opportunities identified in this report could have profound repercussions on both our scientific culture and in our fundamental understanding of ourselves. Because of its strong record of leadership in the physical and mathematical sciences, computer science, and engineering, NSF is optimally positioned to enable and exploit these opportunities.

The following sections consider in greater detail the opportunities for a genuinely multidisciplinary brain science that engages the physical and mathematical sciences, computer science, and engineering for mutual benefit.

2. Opportunities for Mutual Benefit

2.1 Opportunities in Instrumentation and Measurement

The shift in measurement scale from one or a few nerve cells to the study of networks and large populations requires new capabilities in instrumentation and measurement.

Overview: The development of new tools and instrumentation enables measurements that drive scientific advances, and those scientific results in turn motivate and direct the development of new tools. When coupled with well-chosen animal models that cross the scales of simple to complex, and include behavioral, genetic, and developmental approaches, answers to both new and longstanding questions on brain function become feasible. New tools and capabilities will have far-reaching benefits for chemistry, physics, mathematics, engineering and material science. While the capabilities in measurement science address questions about brain function, they also enable new information to be obtained from many complex dynamical systems. Examples of new tools and measurement capabilities needed to enable the next phase of brain studies are highlighted in the remainder of this section. These opportunities will provide a strong motivation for interdisciplinary training at the undergraduate and graduate levels.

Functional measurements in neurons and circuits: A major scientific goal is to characterize the neurons in complex neuronal networks and even in entire brains, and this requires enhancements to our measurement capabilities. Future measurements in neuroscience require greater sensitivity and improved information content, with a higher-throughput and massively-parallel character. Other goals are to develop multi-modality imaging techniques and to transform destructive measurement approaches into non-destructive, minimally invasive, and real-time. Future targets include: 1) functional imaging of multiple neurons in a complex circuit, 2) massively parallel electrical recordings from multiple neurons, control of neural activity using electrical stimulation, and related biophysical measurements, 3) monitoring the small molecule, transcriptional and proteomic changes in individual neurons within a network during behaviorally relevant processes such as learning.

Labeling in neurons and circuits: A century ago, Cajal revolutionized the prevailing view of the nervous system by using the Golgi method to show that the brain consists of individual nerve cells as opposed to a continuous plumbing-like network. Cajal's work was transforming because structure and function are so closely associated in the brain that knowing where some structure is located and what it looks like is intimately related to what it does and how it does it. The Golgi method transformed neuroscience, and it led to a wide variety of techniques for labeling individual neurons and parts of neural circuits. Recent advances in technology make possible a whole new range of labeling methods that have the potential to be as transformative as was the Golgi method in its time. The new labeling technologies will be able to tag structures at all spatial scales from molecules to entire brain regions, with temporal resolution that ranges from milliseconds to the animal's lifetime, and with a specificity that will range from molecular to functional features and historical properties (such as the connection history of neurons). Thus it will become possible to determine which synapses have been used and how much, which cells and brain regions have been active, and which new synapses have formed and which have been eliminated. Ultimately, it will be possible to measure the shifting patterns of functional connectivity between neurons as they adapt to demands of the environment. The ability to conduct comprehensive, multi-scale measurement is essential for understanding brain structure and function. Although this is a long-range and difficult program, the development of these new labeling methods has the potential to transform both neuroscience and the core physical, chemical, and biological science from which they derive.

Opportunities for controlling activity in neurons and circuits: Using molecular genetics, we can now target the expression of gene products that alter activity in specific groups of neurons in neural circuits of interest. There is still much work to be done to improve the temporal and spatial specificity of gene targeting for both invertebrate and vertebrate model organisms. Important targets include new genetically encoded tools to manipulate neuronal electrical activity, cell signal transduction, and synaptic communication in functionally specific groups of neurons. Chemical biology and protein engineering will be necessary to develop novel proteins that exogenously regulate targeted neural activity. Ideally, engineered proteins will be switched on or off by small molecules that do not disturb native neuronal signaling proteins. Control of cell activity can also be achieved via techniques such as the "dynamic clamp", which can inject prescribed currents at prescribed times into neurons with great accuracy. There will be a premium value on controlling neural activity with higher degrees of spatial and temporal specificity, as well as an important need for developing new molecular tools that map the native signaling and genetic machinery of neurons. All of this work should be focused on understanding the regulation of behavior in whole animals. These powerful developments will help advance the role of causal, as opposed to correlational, observations in neuroscience.

The importance of model organisms: Neuroscience research relies on a small number of model organisms that each offer specific advantages (e.g., large physiologically accessible neurons in *Aplysia* versus genetic amenability in *Drosophila* versus similarity of brain structures with humans in mice, rats,

and non-human primates). A reasonable question is whether information gathered from experiments using one model organism is applicable to understanding the nervous systems of other organisms including humans. Molecular biology has provided strong evidence of the unity of neurobiology across different species at the molecular level. This conclusion has been confirmed by decades of molecular neuroscience using vertebrate and non-vertebrate model organisms, showing that the basic molecular machinery of neurons is highly conserved, including neurotransmitter small molecules, key proteins such as ion channels, receptors, and the molecular machinery for synaptic release. Such evolutionary conservation suggests the value of the comparative approach and recommends further work using both invertebrate and vertebrate species, particularly those in which high-throughput comprehensive analysis can be implemented.

2.2 Opportunities in Data Analysis, Statistical Modeling, and Informatics

Increases in the scope of experimental measurements necessitate vastly greater capacities for data acquisition and analysis, and in scope and complexity of the mathematical/computational models employed.

Overview: Brain science is an increasingly data-intensive science that puts new demands on methods for data analysis and modeling. These demands create enormous opportunities for new research in statistics, machine learning, signal processing, stochastic process modeling and related fields. In effect, these demands from brain science constitute an important new motivation for fundamental research in mathematics and statistics, as well as a rich environment for testing new approaches. The resulting developments will be useful both for addressing fundamental questions in brain science and for accelerating research progress in other data-intensive fields of science. The overarching need is go beyond our current analysis methods to keep pace with new demands of the increasing variety and volume of data (neurophysiological, imaging, behavioral, genetic, molecular), increasing numbers of studies and types of experimental tools, and integration of multi-level models of brain function with diverse hypotheses and approaches (biophysical, physiological, behavioral, genetic, and cognitive). Examples of upcoming opportunities for data analysis, modeling, and informatics include:

Biophysical, Physiological, and Behavioral Models to Guide the Development of Data Analysis

Algorithms: While methods can be developed in an ad hoc manner, it is now clear that the increasing wealth of biophysical, physiological, behavioral, genetic, and cognitive information should be used to guide the development of statistics methods, signal processing algorithms and machine learning approaches. This approach has several advantages. First, it uses the specific subject knowledge to impose constraints and to identify assumptions that are most appropriate for model development. Second, it makes explicit use of available results in theory and modeling (See Section 2.3) to inform the development of analysis methods. Third, this approach helps to close the link between modeling, experimentation and data analysis in a principled manner. Analysis methods that use explicitly stated modeling assumptions can work cumulatively, using models inferred from previous experiments to constrain the analyses of subsequent experiments.

Methods to Integrate Diverse Data Sources: A common experimental paradigm in neuroscience is to make measurements on a neural system with different measurement tools simultaneously or in sequence. For example, brain imaging studies that record fMRI and EEG simultaneously are becoming more prevalent. The fMRI provides information on a fine spatial scale (millimeters), whereas the EEG provides information on a fine temporal scale (milliseconds). Therefore, optimal fusion of information from these two sources should be based on the design of the particular experiment, the known biophysics of fMRI imaging (pulse sequence, hemodynamic response, physiological noise, and scanner noise) and the known

biophysics of EEG (lead field model, physiological model and noise model). Solutions to this challenging dynamic inverse problem and to others like it in neuroscience will suggest ways to approach similar problems in other fields of science such as systems biology, climatology, ecology, geophysics and economics where simultaneous measurements are made at different scales on high-dimensional dynamical systems.

New Theory in Statistics, Signal Processing and Machine Learning: Attempts to apply and extend current statistics, signal processing and learning theory to the problems in neurosciences will require new fundamental theory in these and related areas, as well as new probability and data analysis models. For example, one set of theories attempts to bound the error in the learned model or statistical estimates, as a function of the volume and type of data, and the flexibility of the underlying modeling assumptions. How can this theory of statistical learning be extended to cover the type of multi-scale models discussed here? Simultaneous recordings of multiple neural spike trains present a new opportunity for developing multivariate, dynamic, point process probability models. The strong nonstationarity of neural signals provides an opportunity for new signal processing research, including problems of time-frequency analysis and nonstationary spectral analysis. Another important area for theoretical advances lies in developing cross-cutting tools that link the currently disparate approaches being used to analyze neuroscience data. More detailed state-space, hidden Markov or latent process models are needed at multiple temporal and spatial scales to infer the dynamics of neural systems beyond the observed activity of a few individual neurons. Advances in information theory are needed to describe more accurately how groups of neurons convey information about biological signal. These are compelling opportunities to develop a broader, unified conceptual and analytic framework to encompass statistics, signal processing and learning theory.

New Tools for Control Theory: Control and homeostasis (i.e. maintaining the state of a physiological system within an appropriate range given the current needs of the organism) are important principles governing the behavior of neural systems. Moreover, the design of prosthetic devices (e.g., prosthetic limbs, hippocampal prostheses) and brain machine interfaces (e.g., epilepsy implants) will require new control techniques. For example, classical stochastic control models often employ linear Gaussian observations and linear Gaussian state models (e.g., the Kalman filter). For neural systems, the state models are most likely to be non-linear (e.g., dynamics of limb movement) and the observations are likely to be high-dimensional point processes (neural spike trains from a motor area) or combination of high-dimensional point processes and continuous signals (EEG, EMG). Hence, there is a compelling need for new signal processing techniques, beyond linear and Gaussian methods, to properly study these control problems. Restated, applications of control theory to the brain involve new classes of observables and new classes of controllables, many of which challenge existing statistical and control theory.

Models to Analyze Processes Simultaneously at Multiple Levels and Spatial/Temporal Scales: Today there are statistical learning methods and signal processing algorithms that successfully operate at specific levels of scale and abstraction of neural systems. At the level of small groups of neurons, methods are available that can learn the spiking patterns in motor cortex and provide the program to move a prosthetic arm. At a more intermediate level exemplified by signal processing in the visual cortex, available algorithms can calculate optimal sparse codes for natural scenes. At the whole-brain level, methods are available that determine from fMRI data whether a human subject is reading words about tools or buildings. However, we lack methods for coupling these different levels of analysis in a way that allows them to mutually inform and constrain one another. A significant research opportunity is to develop new statistical learning methods and signal processing approaches to analyze simultaneously the variety of data across the broad range of temporal and spatial scales seen in neuroscience. These new methods will also have impact on many other fields examining complex, hierarchical systems, such as other fields of biology, economics, and geology.

Methods to Infer Causality in Neural Systems: While many of today’s data analysis methods focus on correlations or associations, studying the brain involves many questions about causality. How do environmental stimuli, the current state of a particular brain region and activity in other regions to which it is connected determine its subsequent activity? How does one area trigger activity in another? The opportunity here is to develop new methods for inferring causality (in contrast to correlation), driven specifically by studies of neural systems. Neuroscience is particularly well-suited for studies of causality because, unlike economics or the social sciences, it allows conducting multiple controlled experiments in which stimuli and experimental conditions are systematically varied and responses are recorded at multiple time and spatial scales. Advances in methods for inferring causality in neural systems will have significant impact throughout biology and many other domains.

New Approaches to Managing and Sharing Data and Computational Models: Because neuroscience is an increasingly data-intensive discipline, solving the problems of collecting, storing, indexing, retrieving, maintaining, and sharing data will be central to its progress. On one hand, neuroscience can benefit substantially from progress in other data-intensive sciences (e.g., integrating data collected across many laboratories, in a variety of modalities and formats, and under differing experimental conditions). On the other hand, the brain presents data management and sharing challenges that are intrinsically linked to their multi-scale character. These are ideal mutual opportunities with the computer science and artificial intelligence communities, where methods are under development to manage large, heterogeneous data sets using semantic web methods. Another opportunity is to link the models learned from a set of experimental studies to the data sets themselves, and subsequently, retesting and refining the models as relevant new data are acquired. Furthermore, it is critical that learned models themselves (which represent fragments of the evolving theory of the brain) be collected, maintained, indexed and retrieved by the scientific community in the same way as experimental data. The complexity and multi-scale nature of models in neuroscience offers an important opportunity to place learned models online, and to treat them as valuable contributions to be shared and refined by the scientific community.

2.3 Opportunities in Conceptual and Theoretical Approaches

Comprehensive measurements of the brain in space and time, combined with new approaches to the analysis and modeling of comprehensive multi-scale data, will enable the exploration of much richer conceptual and theoretical approaches to understanding the brain at all levels.

Overview: The goal of the theoretical approaches discussed here is to develop conceptual frameworks, mathematical approaches, and computational techniques that tie together data, analyses, and models across the multiple spatial and temporal hierarchical levels that characterize brain function. Ultimately, top-down ideas about how the brain accomplishes specific global functions must meet bottom-up ideas about how function emerges from molecular and cellular mechanisms. Comprehensive theories of brain function, properly constrained by experimental data, can provide novel bridges across these distinct levels. Theory can also provide a context for exploring the relevance of experimental data to specific questions, guide the selection of statistical tools for data analysis, and allow for the reevaluation and reformulation of foundational hypothesis that can be then tested experimentally. The continued and potentially explosive expansion of data, analyses, and theories, requires structured interdisciplinary collaborations among mathematicians, physicists, engineers and neuroscientists.

The Fundamental Role of Mathematics: Almost four hundred years ago, Galileo wrote (roughly translated) “the book of Nature is written in the language of mathematics.” For many scientists since Galileo, scientific understanding has come to mean providing a theory that includes a clear and concise mathematical description of the phenomenon. The strongest theories in the natural sciences transcend

formulation of abstract concepts to make precise quantitative predictions to be tested through similarly quantitative experiments. Although the intrinsic complexity of the problems addressed in the biological and social sciences have caused them to lag the physical sciences in this march toward mathematical understanding, the last few decades have seen major steps toward precisely formulated mathematical theories for particular aspects of brain function. Advancing from these first steps to something that could be legitimately be called a *theory of the brain* constitutes a great challenge to all of the mathematically oriented disciplines, including the physical sciences, computer science, engineering, and mathematics itself.

Dynamical Systems as a Framework for Understanding Brain Function: One major direction for understanding brain function is to show how the dynamics of microscopic components can be analyzed to explain the emergence of macroscopic features. The theory of dynamical systems provides useful tools for describing the time evolution of systems with many interacting degrees of freedom. Analytical methods for averaging, smoothing, and embedding allow for the identification of a relatively small set of relevant variables that dominate the dynamical behavior of the system. The resulting simplified models provide grounds for investigating the ‘how’ and the ‘why’ of brain function. An important example of the dynamical systems approach to the study of brain function arises in the investigation of collective rhythms, whose spectral bands are associated with different behavioral and/or cognitive states. Rhythms are an ideal subject for developing methods for integration across scales, since both experiments and theory at the local network level can be highly influenced by evolving knowledge about the underlying anatomy as well as the relationship between cognitive states (attention, arousal, response to rewards) and the neuromodulators that produce them. Conversely, the study of neural rhythms has enriched the field of dynamical systems. Dynamical systems theory has begun to be applied to questions of how different rhythms depend on different combinations of intrinsic and synaptic ionic currents, how the same network can switch among different rhythms in different modulatory contexts, and how the dynamics in a specific network can gate incoming signals and influence downstream effects. Understanding how large networks process their spatially and temporally patterned inputs requires new mathematical tools. These include techniques for locally reducing large dimensional systems, ways to understand switches in global dynamic behavior, and combinations of dynamical systems and probability/statistics that enable some features of the system to be treated probabilistically while others retain detailed characterization.

Statistical Physics Analysis of Systems With Many Interacting Degrees of Freedom: The conceptual framework and analytical tools of statistical physics are intrinsically well suited to study systems composed of many interacting degrees of freedom through the identification and investigation of a reduced set of relevant macroscopic variables. A model that includes every known microscopic detail is likely to come close to mimicking the actual behavior of the neural system under study, but this type of detailed description does not illuminate the emergence of complex collective behavior. A theoretical description that addresses fundamental questions of ‘how’ and ‘why’ must involve some degree of simplification, a modeling process that is well guided by the statistical physics concepts of scaling, invariance, symmetry, and level-dependent state variables. This type of approach has already yielded useful insights about brain function. In the area of learning and adaptation, statistical physics has provided a successful approach that complements that of machine learning in describing the gradual transformation that leads a neural network towards the implementation of a desired functionality. In an unexpected twist of reciprocity, the very study of such systems has opened a novel area of research in statistical physics, from considering the properties of a system given the interactions among its constituents to asking questions about the types of interactions that will give rise to specific properties.

Engineering Approaches: A fruitful and exciting direction is to ask engineering-style questions about brain function. For a brain system with identified function, one may ask: “Is the system at its optimal performance limits, given physical and functional constraints?” One might expect the answer to be positive if the system in question has been under strong selection pressure. An important example is

whether a sensory system converts its inputs (e.g., light, sound, or touch) into nerve impulses in an optimal manner given physical constraints. Similarly, wiring length minimization techniques that determine component layout in integrated circuits may help understand connectivity patterns in neural circuits. Questions such as these arise naturally in engineering approaches, such as control theory, communications and computation, providing a reciprocal opportunity with brain science. An important example is homeostatic regulation, which operates at multiple levels in the brain from individual neurons to whole systems. The principles that govern such regulation are the subject of study of feedback control theory. In addition, fundamental engineering principles such as nonlinearity and nonstationarity are critical to understanding the brain. Every known neurobiological mechanism is nonlinear; it is virtually impossible to understand molecular-, cellular-, or systems-level brain function without casting those phenomena in the theoretical framework of nonlinearity. New approaches to nonstationary, nonlinear systems are sorely needed to push the frontier of neuroscience, and likewise, to extend the boundaries of engineering.

Machine Learning Tools for the Investigation and Characterization of Adaptive Systems: Learning and adaptation are lifelong processes that not only control development and the acquisition of new skills and capabilities, but also underlie the robustness associated with the maintenance of acquired skills. The theory of machine learning has made rapid advances in recent years and is now providing data analysis techniques for a wide variety of neural data as discussed above, as well as a rich theoretical approach for understanding learning and adaptation in the nervous system. For example, machine learning algorithms have already been found useful in the application of reinforcement learning to the analysis of dopamine-controlled reward-based learning in primates. Computational cognitive models that capture aspects of the actual structural organization of the brain have successfully reproduced fMRI data across several brain regions for subjects solving simple algebraic and language processing problems.

Large-Scale Simulations as a Powerful Mutual Opportunity for Brain Science and Computer Science: This opportunity arises from the well-documented value of realistic simulations as a tool for the investigation of complex systems in many scientific fields. Currently, the simulation of realistic neural models even for small pieces of the brain strains the state of the art in computer science. This challenge is driving the development of new algorithms, new supercomputing hardware, and even new kinds of special purpose computing hardware. These developments have implications for the efficient computational solution of a wide range of complex problems, not limited to neuroscience. Such large-scale simulations also offer a new, neurally inspired, approach to the longstanding challenge of constructing artificial systems with “intelligent” capabilities.

2.4 Opportunities in Building Brain-like Devices and Systems

Developments in next-generation sensing and measurement tools, new statistical frameworks for data analysis, and novel theoretical formalisms for modeling and understanding of the nervous system, will together form the foundation for a new era of devices and systems based on brain-like principles.

Overview: “Biomimetic” systems, ones that mimic key features of biological systems in general and the brain in particular, are finding widespread applications in a number of important areas, including next-generation computing and simulation platforms, understanding neural coding and neural representations in the brain, the development of biomimetic systems that can interact with the nervous system in real-time, humanoid robotics, and biocompatible neural interfaces.

Analog Approaches to Brain-Like Computers and Large-Scale Simulations: Conventional digital simulations of the brain use transistors to perform binary arithmetic to approximate quantities of interest

in many simulations (synaptic currents, action potentials, etc.). In contrast, electrical engineers have begun to emulate the ionic current in a neuron's ion channel directly with the transistor's electronic current. While the current in a present-day transistor, which has a hundred-nanometer-wide channel, corresponds to the current in a small population of ion-channels and not a single channel, this analog approach provides an extremely efficient method to simulate the brain while at the same time laying the engineering groundwork for building brain-like computers out of the next decade's nanotransistors. Developing specialized neural simulation hardware based on such an analog approach promises the same kinds of performance-cost improvements achieved in astrophysical simulations, where a \$42K special-purpose computer has revolutionized the simulation of galaxies.

Stochastic Semiconductor Circuits: Exploiting the analogy between transistor currents and ion-channel currents may also have benefits for semiconductor circuits. The digital computers we use today face a serious problem within the next ten years, as transistors shrink to nanometer dimensions. At ten nanometers, the transistor's channel becomes so narrow that an electron trapped by a dangling bond at the surface (an unavoidable atomistic defect) can block electron flow, causing the current to turn on and off stochastically as trapping and detrapping occur randomly. This stochastic behavior, much like that of an ion channel, will undermine the foundation of digital computation, which counts on transistors behaving like switches with deterministic on/off states to perform binary arithmetic. The brain solved this problem when it evolved the capability to do computation with ion-channels, whose single-atom gates flip open and close randomly, agitated by thermal forces.

Neural Coding and Functional Biomimetic Systems: The advent of simultaneous multi-site recordings from spatially identified single neurons will for the first time provide the basis for capturing the system dynamics of specific neuronal populations. Characterizations of such system-level function is already emerging from studies of motor, sensory, hippocampal, and other cortical regions. Combined experimental and theoretical study of neuronal population dynamics in behaving animals will begin to unravel the "neural code"; that is, the identification and interpretation of spatio-temporal patterns related to specific environmental events, motor movements, or presumed cognitive functions. It is through this path that revolutionary progress can be made in understanding the neural basis of higher thought processes underlying essential brain functions such as perception, language understanding, and memory. Ultimately what will emerge are biomimetic models that capture the function of a neural system. When coupled with multi-site recording/stimulation arrays, such biomimetic systems will be capable of bi-directional communication with the brain, leading to a new generation of brain-computer interfaces (BCIs) that can both sense neural codes and respond with electrical stimulation to send biologically meaningful neural signals back into the brain. When miniaturized in silicon chips, such BCIs would be ideal as neural prostheses that can substitute for damaged neural systems through bi-directional interaction with the brain.

Brain-Like Robotics: Higher-level brain functions instantiated in silicon or other hardware platforms open the opportunity for utilizing biomimetic models in artificial systems such as robots. It has long been an objective in the field of robotics to develop machine vision systems having the perceptual and object recognition capabilities of the mammalian brain. Achieving system-level models of biological visual systems could help realize this goal. At an even higher level, developing an understanding of the neural basis of navigation could provide expanded capabilities in terms of autonomous guidance. Machine learning algorithms could be utilized to mimic other forms of intelligent behavior and adaptation. Similarly, incorporating neural strategies for the hierarchical control of arms, wrists, hands, and digits could provide a new level of agile reach, grasp, and manipulation functionality of robotic systems.

Biocompatible Neural Interfaces: Like other physiological systems, the brain responds to non-biological materials as foreign objects from which the brain should be protected. Thus, the electrophysiological, chemical, and other sensors used to measure brain activity, and the microfabricated

biomimetic devices we might attempt to interface with the brain, trigger a multi-phase, multi-stage “foreign body response” by brain cells that essentially insulates from each other the very neural and physical systems we are attempting to integrate. These issues are particularly critical for long-term (weeks to years) measurements and interactions between engineered and neural systems, which will be critical in the context of studying learning and memory functions, development and aging, and neural prostheses. Programmatic efforts by teams of material scientists, chemists, and biomedical engineers will develop novel molecular structures, perhaps variations of cell adhesion molecules, that can be applied to probes and silicon-based devices to envelop their surfaces in biocompatible materials. These efforts will require target-specific molecular design and synthesis, surface patterning methods to selectively apply the adhesion (or repulsion) compounds, and nanoscale-level coupling of materials and cells. Ultimately, families of designed, biocompatible neural interface systems will be developed to integrate different cell populations with multiple surfaces and materials.

3. Implications and Opportunities for Science Education

Overview: Achieving the promise of the opportunities identified above will require significant changes in undergraduate and graduate science education. In turn, a more systematic understanding of brain mechanisms for learning and memory, combined with parallel advances in cognitive science, may lead to new generations of teaching and learning strategies and to new educational technologies.

A deep understanding of how brains change in response to experience will benefit significantly from a stronger relationship of brain science to the physical and mathematical sciences, computer science, and engineering, and may in turn have profound implications for science education and our education system in general. Certainly the science of learning and memory, both at the behavioral and neurophysiological/neurochemical levels, has already contributed significantly to the design of educational systems and processes. However, those contributions are in general piecemeal, relatively isolated, and do not reflect a fully quantitative and predictive theory of learning and memory in the sense discussed in Section 2.3 above. The development of such theories, which is a major opportunity identified by the Steering Group, may lead to whole new generations of teaching and learning strategies and to new educational technologies.

There are also nearer term educational opportunities for increased involvement of the physical sciences, mathematics, computer science, and engineering in studies of the brain. The brain sciences of the future will demand an ever more broadly and quantitatively trained population of students, post-doctoral researchers, and faculty. The research and development required to achieve many of the opportunities identified above will constitute a rich training ground for the multi-disciplinary scientists of the future. As an example in theoretical areas, the Steering Group identified a number of very specific ways in which brain science could benefit from increased education and training in mathematical, statistical, and numerical methods:

- Develop continuing education programs such as workshops and short courses for neuroscientists to educate themselves in relevant quantitative methods from statistics, statistical physics, nonlinear dynamics, information theory, and machine learning (current examples include the Cold Spring Harbor Laboratory courses and the Marine Biology Laboratory Neuroinformatics and Methods in Computational Neuroscience courses).
- Create incentives to modify graduate and undergraduate curricula in order to raise the general level of quantitative and statistical sophistication (i.e. the ability to reason logically and under uncertainty).
- Identify exciting data analysis and modeling problems from brain science that can be used as part of a curriculum to foster greater interest among K-12 students in brain science in general and quantitative approaches in particular.

- Develop data analysis and modeling laboratory courses analogous to laboratory courses in physics, chemistry and engineering to teach conceptual principles, mathematical formulations, and numerical tools of relevance to brain science.
- Organize and support data analysis challenges (as was done for fMRI at the 2006 Human Brain Mapping meeting) to engage quantitative scientists in brain science problems.
- Provide specific incentives to engage more quantitative, theoretical, and statistical scientists in addressing brain science problems.

Similar strategies can be envisioned for other areas of the mathematical and physical sciences, computer science, and engineering, as well as for areas of the biological and social sciences.

4. Considerations for Implementation: Science Organization and “Organizing for Repeated Innovation”

Overview: Achieving the promise of the opportunities identified above may require new approaches to truly multi-disciplinary science organizations and mechanisms of science support.

Advancing interdisciplinary frontiers in brain sciences can be facilitated by organizational infrastructures supporting repeated discovery, invention and innovation.

Science Organization: The goal of achieving and enhancing multi- and inter-disciplinary collaboration appears with increasing frequency and prominence in the pronouncements of federal science agencies, leading foundations, major universities, and federal laboratories. However, the reality is that many key incentives for scientists, mathematicians, and engineers continue to favor the individual over the collaboration, and being the Principal Investigator (PI) over being an indispensable member of a research collaboration or team. These incentives range from hiring, promotion, and tenure decisions, to decisions about grants and other funding, to major scientific and technical awards. Although many factors contribute to the disparity between official pronouncements and organizational reality, one very important factor is what happens when institutional resources, whether they are university positions or granting agency research budgets, get divided into ever smaller and smaller pieces as they get allocated down an organizational chain. For example, while a university dean might sincerely empathize with the importance of multi-disciplinary and inter-departmental collaborations, when one of her departments needs to make a hiring decision with only one slot available and multiple teaching and research needs to be met, collaboration often takes second place to those more proximal and immediate needs. American universities and government agencies have developed reward systems that emphasize the role of individual PIs. For research programs that require large multi-disciplinary collaborations and teams, the development of necessary new approaches for crediting contribution and excellence will require insight and innovation. Funding approaches that specifically emphasize collaborative teams would provide a useful mechanism for addressing these institutional obstacles. These and other approaches merit further exploration.

Organizing for Repeated Innovation: Advancing interdisciplinary frontiers in brain sciences can be facilitated by organizational infrastructures supporting repeated discovery, invention and innovation. Successful implementation requires development of an initiative that addresses issues of collaboration and structuring a rich portfolio of funding strategies. A program for interdisciplinary innovation should interface agencies’ traditional disciplinary organization through developing vehicles for embedding and transitioning promising activities into continuing agency programs.

Organizing for repeated breakthrough innovation has been shown to involve a number of specific stages: concept development, program experimentation and acceleration, and building a critical mass to sustain interdisciplinary research and collaboration. Concept development was a feature of this steering group

workshop on brain science and of a previous workshop on grand challenges and cognition. Building and nurturing a successful effort from this point will involve experimenting with a diversity of activities and funding approaches. Examples of successful past approaches include: 1) development of knowledge bases for sharing theoretical perspectives, methods, data, results and ideas across disciplines; 2) mechanisms for fostering various scales of collaboration (PI dyads; small groups of investigators, centers -- an important question is whether the creation of necessary capabilities will require the development of new “centers of excellence” or can be adapted to existing laboratories); 3) encouraging different forms of engagement (workshops in which physicists talk to mathematicians; materials scientists talk to biologists, etc; interdisciplinary meetings/forums/debates); 4) diverse support mechanisms (e.g., open public challenges and contests where funding is used to support competitions such as in data analysis; innovative pacing and timing of funding where the next level of support is contingent on what one has learned about the acceptance and robustness of an idea, methodological approach, theoretical principle, or new tool, etc.) A key aspect of program incubation and development will be to build on ongoing domains of relevant expertise. Innovative program experimentation at this time should provide the initial basis for identifying key mechanisms to build a community for sustained discovery and innovation in brain science, the physical and mathematical sciences, computer science, and engineering.

5. Conclusions

The Steering Group concluded that there are indeed **significant unexploited opportunities for mutual scientific benefit** between brain science and the physical and mathematical sciences, computer science, and engineering.

- Questions and challenges posed by brain science are beginning to be adopted and incorporated into the physical and mathematical sciences, computer science, and engineering, and the brain continues to play an important role in the biological and social sciences.
- This breadth and depth of penetration by questions from brain science into the broad scientific landscape is well-matched to NSF’s scientific mission and organization, and NSF is well-positioned to help exploit the opportunities identified.
- The integrative character of the opportunities identified in this report require extensive coordination and collaboration across disciplines, as well as programmatic ingenuity and creativity to implement. For these reasons as well, NSF is ideally suited to lead such an initiative.
- While science is of course unpredictable in detail, the broad opportunities identified above are likely to be prominent in the next wave of major innovation and progress; the time is ripe to exploit them for major advances in fundamental science, for improvements in science education, and for advancing U.S. competitiveness.
- Advances in truly multi-disciplinary science education are required and pertinent implementation strategies will need to be considered.

A useful way to conceptualize the space of opportunities identified by the Steering Group is as a two-dimensional matrix, with the four broad areas of opportunity as rows:

- Instrumentation and Measurement
- Data Analysis, Statistical Modeling, and Informatics
- Conceptual and Theoretical Approaches
- Building Brain-like Devices and Systems

and the major disciplines (or NSF directorates) as columns. Each cell in such a matrix would contain the specific opportunities for that discipline (or directorate) in the corresponding broad area above. Just a moment’s thought will show that the vast majority of cells in such a matrix are filled, which serves to

emphasize the breadth of the collective opportunity for mutual scientific benefit between brain science and the physical and mathematical sciences, computer science, and engineering. NSF has the capacity to integrate broadly across disciplines, to integrate research with education, and to stimulate new innovation and discovery needed to exploit this matrix of opportunities.

This report has presented a framework, one that the Steering Group was excited to help identify and formulate. Further community input is needed to explore the opportunities in greater depth and to help craft an effective scientific agenda to realize them.

Appendices

Appendix 1: Suggested Readings

Readings for Section 2.1: Opportunities in Instrumentation and Measurement

Wightman, R. Probing Cellular Chemistry in Biological Systems with Microelectrodes. *Science*, 2006, 311: 1570-1574.

From the Abstract: “Over the past 20 years, the technological impediments to fabricating electrodes of micrometer dimensions have been largely overcome. These small electrodes can be readily applied to probe chemical events at the surface of tissues or individual biological cells; they can even be used to monitor concentration changes within intact animals. These measurements can be made on rapid time scales and with minimal perturbation of the system under study. Several recent applications have provided important insights into chemical processes at cells and in tissues. Examples include molecular flux measurements at the surface of single cells and through skin—which can offer insights into oxidative stress, exocytosis, and drug delivery—and real-time brain neurotransmitter monitoring in living rats, which reveals correlations between behavior and molecular events in the brain. Such findings can promote interdisciplinary collaborations and may lead to a broader understanding of the chemical aspects of biology.”

Becker, M., Schindler, J., Nothwang, H. Neuroproteomics - the tasks lying ahead. *Electrophoresis*, 2006, 27: 2819-29.

From the Abstract: “The brain is unquestionably the most fascinating organ. Despite tremendous progress, current knowledge falls short of being able to explain its function. An emerging approach toward improved understanding of the molecular mechanisms underlying brain function is neuroproteomics. Today's neuroscientists have access to a battery of versatile technologies both in transcriptomics and proteomics. The challenge is to choose the right strategy in order to generate new hypotheses on how the brain works. The goal of this review is therefore two-fold: first we recall the bewildering cellular, molecular, and functional complexity in the brain, as this knowledge is fundamental to any study design. In fact, an impressive complexity on the molecular level has recently re-emerged as a central theme in large-scale analyses. Then we review transcriptomics and proteomics technologies, as both are complementary. Finally, we comment on the most widely used proteomics techniques and their respective strengths and drawbacks. We conclude that for the time being, neuroproteomics should focus on its strengths, namely the identification of posttranslational modifications and protein-protein interactions, as well as the characterization of highly purified subproteomes. For global expression profiling, emphasis should be put on further development to significantly increase coverage.”

Segev, R., Goodhouse, J., Puchalla, J., and Berry, M.J. III. Recording Spikes from a Large Fraction of the Ganglion Cells in a Retinal Patch. *Nature Neuroscience*, 2004, 7: 1155-1162.

From the Abstract: “To understand a neural circuit completely requires simultaneous recording from most of the neurons in that circuit. Here we report recording and spike sorting techniques that enable us to record from all or nearly all of the ganglion cells in a patch of the retina. With a dense multi-electrode array, each ganglion cell produces a unique pattern of activity on many electrodes when it fires an action potential. Signals from all of the electrodes are combined with an iterative spike sorting algorithm to resolve ambiguities arising from overlapping spike waveforms. We verify that we are recording from a large fraction of ganglion cells over the array by labeling the ganglion cells with a retrogradely transported dye and by comparing the number of labeled and recorded cells. Using these methods, we show that about 60 receptive fields of ganglion cells cover each point in visual space in the salamander, consistent with anatomical findings.” This paper describes a novel approach to recording a large fraction of the neurons in a particular circuit, in this case the ganglion cells from a patch of retina, that form the

basis of the analyses and theoretical developments described by Schneidman et al. in the final “Putting It All Together” section below.

Prinz AA, Abbott LF and Marder E, The dynamic clamp comes of age, *Trends in Neurosciences*, 2004, 27: 218-224.

From the Abstract: “The dynamic clamp uses computer simulation to introduce artificial membrane or synaptic conductances into biological neurons and to create hybrid circuits of real and model neurons. In the ten years since it was first developed, the dynamic clamp has become a widely used tool for the study of neural systems at the cellular and circuit levels. This review describes recent state-of-the-art implementations of the dynamic clamp and summarizes insights gained through its use, ranging from the role of voltage-dependent conductances in shaping neuronal activity to the effects of synaptic dynamics on network behavior and the impact of in vivo-like input on neuronal information processing.”

Nitabach, M.N., et al. Electrical Silencing of *Drosophila* Pacemaker Neurons Stops the Free-Running Circadian Clock, *Cell*, 2002, 109: 485-495.

From the Abstract: “Electrical silencing of *Drosophila* circadian neurons through targeted expression of K⁺ channels causes severe deficits in free-running circadian locomotor rhythmicity in complete darkness. Pacemaker electrical silencing also stops the free-running oscillation of PERIOD (PER) and TIMELESS (TIM) proteins that constitutes the core of the cell-autonomous molecular clock. In contrast, electrical silencing fails to abolish PER and TIM oscillation in light-dark cycles, although it does impair rhythmic behavior. On the basis of these findings, we propose that electrical activity is an essential element of the free-running molecular clock of pacemaker neurons along with the transcription factors and regulatory enzymes that have been identified as required for clock function.”

Readings for Section 2.2: Opportunities in Data Analysis, Statistical Modeling, and Informatics

Brown E., Kass, R., and Mitra P. Multiple neural spike train data analysis: state-of-the-art and future challenges, *Nature Neuroscience*, 2004, 7(5): 456-61.

From the Abstract: “Multiple electrodes are now a standard tool in neuroscience research that make it possible to study the simultaneous activity of several neurons in a given brain region or across different regions. The data from multi-electrode studies present important analysis challenges that must be resolved for optimal use of these neurophysiological measurements to answer questions about how the brain works. Here we review statistical methods for the analysis of multiple neural spike-train data and discuss future challenges for methodology research.”

Victor, J. Analyzing receptive fields, classification images and functional images: challenges with opportunities for synergy. *Nature Neuroscience*, 2005, 8: 1651-1656.

From the Abstract: “In neurophysiology, psychophysics, optical imaging and functional imaging studies, the investigator seeks a relationship between a high-dimensional variable, such as an image, and a categorical variable, such as the presence or absence of a spike or a behavior. The usual analysis strategy is fundamentally identical across these contexts—it amounts to calculating the average value of the high-dimensional variable for each value of the categorical variable and comparing these results by subtraction. Though intuitive and straightforward, this procedure may be inaccurate or inefficient and may overlook important details. Sophisticated approaches have been developed within these several experimental contexts, but they are rarely applied beyond the context in which they were developed. Recognition of the relationships among these contexts has the potential to accelerate improvements in analytic methods and to increase the amount of information that can be gleaned from experiments.”

Eden U., Frank, L., Barbieri, R., Solo, V., Brown, E. Dynamic analyses of neural encoding by point process adaptive filtering, *Neural Computation*, 2004, 16(5): 971-998.

From the Abstract: “Neural receptive fields are dynamic in that with experience, neurons change their spiking responses to relevant stimuli. To understand how neural systems adapt their representations of biological information, analyses of receptive field plasticity from experimental measurements are crucial. Adaptive signal processing, the well-established engineering discipline for characterizing the temporal evolution of system parameters, suggests a framework for studying the plasticity of receptive fields. We use the Bayes’ rule Chapman-Kolmogorov paradigm with a linear state equation and point process observation models to derive adaptive filters appropriate for estimation from neural spike trains. We derive point process filter analogues of the Kalman filter, recursive least squares, and steepest-descent algorithms and describe the properties of these new filters. We illustrate our algorithms in two simulated data examples. The first is a study of slow and rapid evolution of spatial receptive fields in hippocampal neurons. The second is an adaptive decoding study in which a signal is decoded from ensemble neural spiking activity as the receptive fields of the neurons in the ensemble evolve. Our results provide a paradigm for adaptive estimation for point process observations and suggest a practical approach for constructing filtering algorithms to track neural receptive field dynamics on a millisecond timescale.”

Mitchell, T. The Discipline of Machine Learning. <http://www.cs.cmu.edu/~tom/pubs/MachineLearning.pdf>

From the Abstract: “Over the past 50 years the study of Machine Learning has grown from the efforts of a handful of computer engineers exploring whether computers could learn to play games, and a field of Statistics that largely ignored computational considerations, to a broad discipline that has produced fundamental statistical-computational theories of learning processes, has designed learning algorithms that are routinely used in commercial systems for speech recognition, computer vision, and a variety of other tasks, and has spun off an industry in data mining to discover hidden regularities in the growing volumes of online data. This document provides a brief and personal view of the discipline that has emerged as Machine Learning, the fundamental questions it addresses, its relationship to other sciences and society, and where it might be headed.

Mitchell, T. AI and the Impending Revolution in Brain Sciences. Presidential Address to the American Association of Artificial Intelligence, 2002. <http://www.cs.cmu.edu/~tom/pubs/AAAI-PresAddr.pdf>

From the Abstract: “The synergy between AI and Brain Sciences will yield profound advances in our understanding of intelligence over the coming decade, fundamentally changing the nature of our field.”

Pittendrih, S. and Jacobs, G. Neurosys: A Semistructured Laboratory Database, *Neuroinformatics*, 2003, 1: 167-178.

From the Abstract: “The inherent complexity of traditional relational database systems is a key obstacle to more widespread use of database technology in the neuroscience community. As an alternative to relational technology, we propose a simpler semistructured data model for documenting laboratory procedures and results. The semistructured data model allows researchers to document their data in an organized, regularly formatted, machine readable, and network accessible manner, without requiring the services of database professionals. We present proof-of-concept software, consisting of an HTML interface that communicates with a remotely located, semistructured database. We also discuss the importance of standardized terminology and the importance of building flexible data description systems that are more easily adapted and reconfigured to conform with standardized terminologies as they evolve.”

Mazziotta J, Toga A, Evans A, Fox P, et al. A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM). *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 2001, 356: 1293-1322.

From the Abstract: “Motivated by the vast amount of information that is rapidly accumulating about the human brain in digital form, we embarked upon a program in 1992 to develop a four-dimensional probabilistic atlas and reference system for the human brain. Through an International Consortium for Brain Mapping (ICBM) a dataset is being collected that includes 7000 subjects between the ages of eighteen and ninety years and including 342 mono- and dizygotic twins. Data on each subject includes detailed demographic, clinical, behavioural and imaging information. DNA has been collected for genotyping from 5800 subjects. A component of the programme uses post-mortem tissue to determine the probabilistic distribution of microscopic cyto- and chemoarchitectural regions in the human brain. This, combined with macroscopic information about structure and function derived from subjects in vivo, provides the first large scale opportunity to gain meaningful insights into the concordance or discordance in micro- and macroscopic structure and function. The philosophy, strategy, algorithm development, data acquisition techniques and validation methods are described in this report along with database structures. Examples of results are described for the normal adult human brain as well as examples in patients with Alzheimer's disease and multiple sclerosis. The ability to quantify the variance of the human brain as a function of age in a large population of subjects for whom data is also available about their genetic composition and behaviour will allow for the first assessment of cerebral genotype-phenotype-behavioural correlations in humans to take place in a population this large. This approach and its application should provide new insights and opportunities for investigators interested in basic neuroscience, clinical diagnostics and the evaluation of neuropsychiatric disorders in patients.”

Readings for Section 2.3: Opportunities in Conceptual and Theoretical Approaches

Modeling the Mind, Special Issue of Science, 2006, 314.

From the Editorial Introduction: “In areas ranging from molecules to the highest brain functions, scientists use mathematical models and computer simulations to study and predict the behavior of the nervous system. Simulations are essential because the present experimental systems are too complex to allow collection of all the data. Modeling has become so powerful these days that there is no longer a one-way flow of scientific information. There is considerable intellectual exchange between modelers and experimentalists. The results produced in the simulation lab often lead to testable predictions and thus challenge other researchers to design new experiments or reanalyze their data as they try to confirm or falsify the hypotheses put forward. For this issue of *Science*, [the editors] invited leading computational neuroscientists, each of whom works at a different organizational level, to review the latest attempts of mathematical and computational modeling and to give us an outlook on what the future might hold in store.”

Markram, H. The Blue Brain Project, *Nature Reviews Neuroscience*, 2006, 7: 153-160.

From the Abstract: “IBM's Blue Gene supercomputer allows a quantum leap in the level of detail at which the brain can be modelled. [Markram contends] that the time is right to begin assimilating the wealth of data that has been accumulated over the past century and start building biologically accurate models of the brain from first principles to aid our understanding of brain function and dysfunction. Detailed, biologically accurate brain simulations offer the opportunity to answer some fundamental questions about the brain that cannot be addressed with any current experimental or theoretical approaches.

Traub R, Contreras D, Cunningham M, Murray H, LeBeau F, Roopun A, Bibbig A, Wilent W, Higley M, and Whittington M. Single-column thalamocortical network model exhibiting gamma oscillations, spindles and epileptogenic bursts. *Journal of Neurophysiology*, 2005, 93: 2194-2232.

From the Abstract: “To better understand population phenomena in thalamocortical neuronal ensembles, we have constructed a preliminary network model with 3,560 multicompartiment neurons (containing soma, branching dendrites, and a portion of axon)... Our network model replicates several observed population phenomena, including 1) persistent gamma oscillations; 2) thalamocortical sleep spindles; 3) series of synchronized population bursts, resembling electrographic seizures; 4) isolated double population bursts with superimposed very fast oscillations; 5) spike-wave, polyspike-wave, and fast runs.”

Kopell, N. Does it have to be so complicated? Editorial Focus on: A single-column thalamo-cortical network model exhibiting gamma oscillations, spindles and epileptogenic bursts, *Journal of Neurophysiology*, 2005, 93: 1829-30.

This is a brief essay on different styles of modeling in neuroscience, and what one can expect to learn from each.

LeMasson, G., et al. Activity-dependent regulation of conductances in model neurons. *Science*, 1993, 259: 1915-1917.

Turrigiano, L.G., et al. Activity-dependent changes in the intrinsic properties of cultured neurons. *Science*, 1994, 264: 974-977.

The Hodgkin-Huxley model and its generalizations provides an essentially exact description for the electrical dynamics of neurons on fast time scales. Fitting these realistic versions of these models, with many different types of channels, to the behavior of particular neurons became something of an industry in quantitative neuroscience. A profound difficulty emerged though this effort, namely, that even the qualitative behavior of the cell could depend sensitively on the number of ion channels of each type that are present in the membrane. Thus, matching the observed properties of neurons would require an implausibly fine tuning of these parameters. The key idea of the first paper, and of subsequent work from Abbott and collaborators, is that if such fine tuning is a problem for us in finding the right model, it is also a problem for the cell in achieving its functional operating point. We have learned in physics that *ad hoc* fine tuning of parameters is not an acceptable explanation for natural phenomena, and in biological systems this argument is even more forceful because of the need for adaptation and evolution. The authors propose instead that functional operating points are stabilized by an additional layer of dynamics in which the electrical activity (perhaps through the intermediary of calcium concentration) feeds back onto the expression or membrane localization of the channels. They show that such a scheme can work, and propose experimental tests that were carried out within the year and reported in the second paper. This has led to a whole new field of experimental neuroscience, examining the mechanisms for homeostasis and the balance between stability and plasticity. Conceptually, the same issues of robustness against parameter variation have since arisen in problems as diverse as the neural circuits for integration and short-term memory and the biochemical circuits for bacterial chemotaxis and pattern formation in embryonic development.

Roxin, A., Rieke, H., and Solla, S.A. Self-Sustained Activity in a Small-World Network of Excitable Neurons, *Physical Review Letters*, 2004, 92, 198101.

This is an example in which techniques from statistical physics and nonlinear dynamics are combined to analyze a novel type of neural model that predicts a fundamental role for complex network connectivity in sustaining persistent activity and providing coherent rapid-processing capabilities. These predictions have been recently tested in both an anatomical analysis of the vertebrate reticular formation and a far from equilibrium photosensitive chaotic system, illustrating the links between brain science and the study of other complex nonlinear systems, as discussed in the workshop.

Vogels, T.P., et al. Neural Network Dynamics. *Annual Review of Neuroscience*, 2005, 28: 357-376.

From the Abstract: “Neural network modeling is often concerned with stimulus-driven responses, but most of the activity in the brain is internally generated. Here, we review network models of internally generated activity, focusing on three types of network dynamics: (a) sustained responses to transient stimuli, which provide a model of working memory; (b) oscillatory network activity; and (c) chaotic activity, which models complex patterns of background spiking in cortical and other circuits. We also review propagation of stimulus-driven activity through spontaneously active networks. Exploring these aspects of neural network dynamics is critical for understanding how neural circuits produce cognitive function.”

Brenner, N., et al. Adaptive rescaling optimizes information transmission. *Neuron*, 2000, 26: 695-702.

Fairhall, A.L., et al. Efficiency and ambiguity in an adaptive neural code. *Nature*, 2001, 412: 787-792 (2001).

It is an old idea that the brain should build representations of sensory data that are in some way matched to the statistical structure of those data, perhaps maximizing some information theoretic measure of efficiency. This matching could occur on evolutionary time scales, during development of the brain, as a part of learning, or as a more rapid adaptation. In particular, since we live in world with an intermittent statistical structure, real time adaptation to low-order statistics would provide the maximally efficient codes. These papers use the motion sensitive neurons of the fly visual system as an experimental testing ground for these ideas. The designs of the experiments use inputs drawn from probability distributions with different variances (but identical means) and correlation times. The first paper showed that (a) the apparent input/output relation of the neuron adapts to the changes in variance, (b) that this happens over a wide range of correlation times, (c) that the form of the adaptation corresponds to a simple rescaling which, in certain limits, is what we expect from a theory of optimal coding, and (d) that the precise scaling factor achieved by this adaptation maximizes the amount of information that the neuron transmits about the visual motion input. Along the way, the authors developed methods for characterizing the input/output relation with dynamic stimuli that have since been used to study many other systems. The second paper examines the dynamics of adaptation, showing that aspects of the process have access to many time scales in a nearly scale-invariant manner, ranging from 0.1 seconds to many minutes. Building on the idea that adaptation serves to maximize information transmission, the authors develop methods to track the transmitted information over time after a switch between two input distributions, in effect 'catching' the system using the wrong code and hence transmitting less information than it does at steady state. Remarkably, the time scale for recovery of the information rate is close to the minimum set by the need to gather statistics on the new distribution. This work reflects an interplay between theory (the ideas of optimizing information transmission), the design of new kinds of experiments motivated by theory (explicit manipulations of the distribution of inputs), and new methods of data analysis.

Readings for Section 2.4: Building Brain-Like Devices and Systems

Boahen, K. Neuromorphic Microchips. *Scientific American*, 2005, 292: 56-63.

From the Abstract: “When IBM's Deep Blue supercomputer edged out world chess champion Garry Kasparov during their celebrated match in 1997, it did so by means of sheer brute force. The machine evaluated some 200 million potential board moves a second, whereas its flesh-and-blood opponent considered only three each second, at most. But despite Deep Blue's victory, computers are no real competition for the human brain in areas such as vision, hearing, pattern recognition, and learning. Computers, for instance, cannot match our ability to recognize a friend from a distance merely by the way he walks. And when it comes to operational efficiency, there is no contest at all. A typical room-size supercomputer weighs roughly 1,000 times more, occupies 10,000 times more space and consumes a

millionfold more power than does the cantaloupe-size lump of neural tissue that makes up the brain. How does the brain--which transmits chemical signals between neurons in a relatively sluggish thousandth of a second--end up performing some tasks faster and more efficiently than the most powerful digital processors? The secret appears to reside in how the brain organizes its slow-acting electrical components.”

Hut, P. and Makino, J. Astrophysics on the GRAPE Family of Special-Purpose Computers. *Science*, 1999, 283: 501-505.

From the Abstract: The GRAPE-4, a special-purpose computer designed specifically for astrophysical applications “has produced some major scientific results through a wide diversity of large-scale simulations in astrophysics. Applications have included planetary formation, the evolution of star clusters and galactic nuclei, and the formation of galaxies and clusters of galaxies.” A new breed of such special-purpose computers could be developed for simulating the brain, with mutual benefits for neuroscience and information technology.

Roy, S. and Asenov, A. Where Do the Dopants Go? *Science*, 2005, 309: 388-390.

From the Abstract: “As the field-effect transistors used in modern electronic devices continue to shrink, scientists and engineers face new challenges. In this Perspective, Roy and Asenov discuss one such challenge: the problem that as device sizes shrink beyond a certain size, atomic-scale differences between devices result in different macroscopic properties. In particular, the locations and numbers of dopant atoms, introduced to alter the electrical properties of regions of the transistor, differ from device to device. The authors discuss recent successes in modeling the dopant distributions and their effects on transistor properties. Such simulations may help researchers to design devices that are resistant to fluctuations in dopant distributions.”

Berger, T.W., et al. Brain-Implantable Biomimetic Electronics as the Next Era in Neural Prosthetics. Invited Paper: *Proceedings of the IEEE*, 2001, 89: 993-1012.

From the Abstract: “An interdisciplinary multilaboratory effort to develop an implantable neural prosthetic that can coexist and bidirectionally communicate with living brain tissue is described. Although the final achievement of such a goal is many years in the future, it is proposed that the path to an implantable prosthetic is now definable, allowing the problem to be solved in a rational, incremental manner. Outlined in this report is our collective progress in developing the underlying science and technology that will enable the functions of specific brain damaged regions to be replaced by multichip modules consisting of novel hybrid analog/digital microchips. The component microchips are “neurocomputational” incorporating experimentally based mathematical models of the nonlinear dynamic and adaptive properties of biological neurons and neural networks. The hardware developed to date, although limited in capacity, can perform computations supporting cognitive functions such as pattern recognition, but more generally will support any brain function for which there is sufficient experimental information. To allow the “neurocomputational” multichip module to communicate with existing brain tissue, another novel microcircuitry element has been developed: silicon-based multielectrode arrays that are “neuromorphic,” i.e., designed to conform to the region-specific cytoarchitecture of the brain. When the “neurocomputational” and “neuromorphic” components are fully integrated, our vision is that the resulting prosthetic, after intracranial implantation, will receive electrical impulses from targeted subregions of the brain, process the information using the hardware model of that brain region, and communicate back to the functioning brain. The proposed prosthetic microchips also have been designed with parameters that can be optimized after implantation, allowing each prosthetic to adapt to a particular user/patient.”

Musallam, S., et al. Cognitive Control Signals for Neural Prosthetics, *Science* 2004, 305: 258-262.

From the Abstract: “Recent development of neural prosthetics for assisting paralyzed patients has focused on decoding intended hand trajectories from motor cortical neurons and using this signal to control external devices. In this study, higher level signals related to the goals of movements were decoded from

three monkeys and used to position cursors on a computer screen without the animals emitting any behavior. Their performance in this task improved over a period of weeks. Expected value signals related to fluid preference, the expected magnitude, or probability of reward were decoded simultaneously with the intended goal. For neural prosthetic applications, the goal signals can be used to operate computers, robots, and vehicles, whereas the expected value signals can be used to continuously monitor a paralyzed patient's preferences and motivation.

Berger, T.W. et al. Restoring Lost Cognitive Function: Hippocampal-Cortical Neural Prostheses. *IEEE Engineering in Medicine and Biology Magazine*, 2005, 30-44.

From the Article: The authors propose and describe “a prosthetic device that functions in a biomimetic manner to replace information transmission between cortical brain regions. In such a prosthesis, damaged CNS neurons would be replaced with a biomimetic system comprised of silicon neurons. The replacement silicon neurons would have functional properties specific to those of the damaged neurons and would both receive as inputs and send as outputs electrical activity to regions of the brain with which the damaged region previously communicated. Thus, the class of prosthesis being proposed is one that would replace the computational function of the damaged brain and restore the transmission of that computational result to other regions of the nervous system. Such a new generation of neural prostheses would have a profound impact on the quality of life throughout society; it would offer a biomedical remedy for the cognitive and memory loss accompanying Alzheimer's disease, the speech and language deficits resulting from stroke, and the impaired ability to execute skilled movements following trauma to brain regions responsible for motor control.”

Hochberg, L.R. et al. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 2006, 442: 164-171.

From the Abstract: “Neuromotor prostheses (NMPs) aim to replace or restore lost motor functions in paralysed humans by routing movement-related signals from the brain, around damaged parts of the nervous system, to external effectors. To translate preclinical results from intact animals to a clinically useful NMP, movement signals must persist in cortex after spinal cord injury and be engaged by movement intent when sensory inputs and limb movement are long absent. Furthermore, NMPs would require that intention-driven neuronal activity be converted into a control signal that enables useful tasks. Here we show initial results for a tetraplegic human (MN) using a pilot NMP. Neuronal ensemble activity recorded through a 96-microelectrode array implanted in primary motor cortex demonstrated that intended hand motion modulates cortical spiking patterns three years after spinal cord injury. Decoders were created, providing a ‘neural cursor’ with which MN opened simulated e-mail and operated devices such as a television, even while conversing. Furthermore, MN used neural control to open and close a prosthetic hand, and perform rudimentary actions with a multijointed robotic arm. These early results suggest that NMPs based upon intracortical neuronal ensemble spiking activity could provide a valuable new neurotechnology to restore independence for humans with paralysis.”

Readings That Cut Across Sections and Integrate Multiple Opportunities:

Womelsdorf T, Fries P, Mitra P, Desimone R. Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature*, 2006, 439: 733-6.

From the Abstract: “In visual areas, attended stimuli induce enhanced responses and an improved synchronization of rhythmic neuronal activity in the gamma frequency band (40–70 Hz)... Here we show that behavioural response times to a stimulus change can be predicted specifically by the degree of gamma-band synchronization among those neurons in monkey visual area V4 that are activated by the behaviourally relevant stimulus... Enhanced neuronal gamma-band synchronization and shortened neuronal response latencies to an attended stimulus seem to have direct effects on visually triggered behaviour, reflecting an early neuronal correlate of efficient visuo-motor integration.” Thus, this paper

illustrates how multi-electrode recordings (Section 2.1) during an attention task, coupled with advanced signal processing methods (Section 2.2) allowed for characterization of collective neural dynamics associated with attention. These results call for theoretical analysis of brain rhythms using dynamical systems methods (Section 2.3), and have implications for reading out cognitive states for a neural prosthetic device (Section 2.4).

Schneidman, E., Berry, M.J., Segev, R., and Bialek, W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 2006, 440: 1007-1012.

From the Abstract: “Biological networks have so many possible states that exhaustive sampling is impossible. Successful analysis thus depends on simplifying hypotheses, but experiments on many systems hint that complicated, higher-order interactions among large groups of elements have an important role. Here we show, in the vertebrate retina, that weak correlations between pairs of neurons coexist with strongly collective behaviour in the responses of ten or more neurons. We find that this collective behaviour is described quantitatively by models that capture the observed pairwise correlations but assume no higher-order interactions. These maximum entropy models are equivalent to Ising models, and predict that larger networks are completely dominated by correlation effects. This suggests that the neural code has associative or error-correcting properties, and we provide preliminary evidence for such behaviour.” This is a particularly clear example of how new, more comprehensive measurements (Section 2.1) allow novel forms of statistical analysis and modeling (Section 2.2), which in turn suggest richer theoretical interpretations (Section 2.3) that make strong ties to other work in statistical physics.

Towards 2020 Science. Workshop Report from Microsoft Research (http://research.microsoft.com/towards2020science/background_overview.htm)

Based on a workshop organized by Microsoft Research “to develop a new vision and roadmap of the evolution, challenges and potential of computer science and computing in scientific research in the next fifteen years. *Towards 2020 Science* sets out the challenges and opportunities arising from the increasing synthesis of computing and the sciences. It seeks to identify the requirements necessary to accelerate scientific advances -- particularly those driven by computational sciences and the 'new kinds' of science the synthesis of computing and the sciences is creating.” Computational neuroscience is a key focus of opportunity. *Towards 2020 Science* also stimulated a special issue of *Nature: 2020 Computing* (*Nature*, 2006, 440: 383ff)

Special Issue on Computational and Systems Neuroscience. *Nature Neuroscience*, 2005, 8(12).

From the Introduction and Editorial: “To understand the brain, theoretical and experimental approaches must be integrated to make sense of the enormous amount of existing data, and to guide future experiments... In reality, theory is an integral part of all good neuroscience papers including experimental papers. Any good paper includes an intuitive framework for its results and why they came out the way they did. For example, a study identifying a new protein involved in long-term potentiation is nothing more than a disconnected data set without a mechanistic framework for how it interacts with other elements in the pathway and an intuition for the functional consequences of these interactions. Theoretical papers simply formalize and explore these intuitions and mechanisms, sometimes leading to the conclusion that our initial, hand-waving explanations do not provide a good fit to the data. Good theories can synthesize large quantities of empirical data, distilling them to a few simple notions, and can establish quantitative relationships between individual observations. They can generate predictions that can serve to validate current and future experiments. Given the vast number of empirical studies being generated by the field and the sheer complexity of the brain, it is clear that theoretical approaches have great potential for making sense of the problem... An increasing number of theorists and biologists are becoming more facile with the language of the complementary approach and are coming to appreciate the value of integrating the two disciplines. However, both fields have a long way to go before it will be commonplace for them to proceed hand in hand.”

Appendix 2: Biographical Sketches of Participants

Theodore W. Berger is the David Packard Professor of Engineering, Professor of Biomedical Engineering and Neuroscience, and Director of the Center for Neural Engineering at the University of Southern California. Dr. Berger received his Ph.D. from Harvard University in 1976; his thesis work received the James McKeen Cattell Award from the New York Academy of Sciences. He conducted postdoctoral research at the University of California, Irvine from 1977-1978, and was an Alfred P. Sloan Foundation Fellow at The Salk Institute from 1978-1979. Dr. Berger joined the Departments of Neuroscience and Psychiatry at the University of Pittsburgh in 1979, being promoted through to Full Professor in 1987. During that time, he received a McKnight Foundation Scholar Award, twice received an NIMH Research Scientist Development Award, and was elected a Fellow of the American Association for the Advancement of Science. Since 1992, he has been Professor of Biomedical Engineering and Neurobiology at the University of Southern California, and was appointed the David Packard Chair of Engineering in 2003. While at USC, Dr. Berger has received an NIMH Senior Scientist Award, was given the Lockheed Senior Research Award in 1997, was elected a Fellow of the American Institute for Medical and Biological Engineering in 1998, received a Person of the Year “Impact Award” by the AARP in 2004 for his work on neural prostheses, was a National Academy of Sciences International Scientist Lecturer in 2003, and an IEEE Distinguished Lecturer in 2004-2005. Dr. Berger was elected a Senior Member of the IEEE in 2005, received a “Great Minds, Great Ideas” award from the EE Times in the same year, and in 2006 was awarded USC’s Associates Award for Creativity in Research and Scholarship. Dr. Berger became Director of the Center for Neural Engineering in 1997, an organization which helps to unite USC faculty with cross-disciplinary interests in neuroscience, engineering, and medicine. Dr. Berger has published over 170 journal articles and book chapters, and is the co-editor of a book recently published by the MIT Press on Toward Replacement Parts for the Brain: Implantable Biomimetic Electronics as Neural Prostheses. Dr. Berger’s research interests are in (i) the development of biologically realistic, experimentally-based, mathematical models of higher brain (hippocampus) function, (ii) application of biologically realistic neural network models to real-world signal processing problems, (iii) VLSI-based implementations of biologically realistic models of higher brain function, (iv) neuron-silicon interfaces for bi-directional communication between brain and VLSI systems, and (v) next-generation brain-implantable, biomimetic signal processing devices for neural prosthetic replacement and/or enhancement of brain function.

William Bialek is John Archibald Wheeler/Battelle Professor in Physics at Princeton University. He is interested in the interface between physics and biology, broadly interpreted. A central theme in his research is an appreciation for how well things “work” in biological systems. It is, after all, some notion of functional behavior that distinguishes life from inanimate matter, and it is a challenge to quantify this functionality in a language that parallels our characterization of other physical systems. Strikingly, when this is done (and there are not so many cases where it has been done!), the performance of biological systems often approaches some limits set by basic physical principles. While it is popular to view biological mechanisms as an historical record of evolutionary and developmental compromises, these observations on functional performance point toward a very different view of life as having selected a set of near optimal mechanisms for its most crucial tasks. Even if this view is wrong, it suggests a theoretical physicist’s idealization; the construction of this idealization and the attempt to calibrate the performance of real biological systems against this ideal provides a productive route for the interaction of theory and experiment, and in several cases this effort has led to the discovery of new phenomena. The idea of performance near the physical limits crosses many levels of biological organization, from single molecules to cells to perception and learning in the brain, and Dr. Bialek has made contributions to this whole range of problems.

Kwabena Boahen is Associate Professor in the Department of Bioengineering at Stanford University.

His research interests include mixed-mode multichip VLSI models of biological sensory and perceptual systems, and their epigenetic development, and asynchronous digital communication for reconfigurable connectivity. He is a bioengineer who is using integrated circuits to understand the way neurons compute, linking the seemingly disparate fields of electronics and computer science with neurobiology and development. His ultimate goal is to build a neuromorphic computer by reverse engineering the nervous system.

Emery N. Brown is Professor of Computational Neuroscience and Professor of Health, Sciences and Technology at Massachusetts Institute of Technology and Associate Professor of Anaesthesia at Harvard Medical School. He is an anesthesiologist and the Director of the Neuroscience Statistics Research Laboratory in the Department of Anesthesia and Critical Care at Massachusetts General Hospital. Dr. Brown earned his BA degree in Applied Mathematics from Harvard College, his MA and Ph.D. degrees in statistics from Harvard University and his MD from Harvard Medical School. He served his internship in internal medicine at the Brigham and Women's Hospital and his residency in anesthesia at Massachusetts General Hospital. His methodology research focuses on use of dynamic estimation methods to analyze neurophysiological systems in three areas: signal processing algorithms to study how individuals and ensembles of neurons represent information; statistical methods for the analysis of functional neural imaging data; and statistical models to characterize human circadian and neuroendocrine rhythms. His experimental research uses combined fMRI and EEG to study the neurophysiological changes in brain regions associated with the states of general anesthesia. Dr. Brown is a fellow of the American Statistical Association, fellow of the American Institute of Medical and Biological Engineering, and he is the Co-Director of the Neuroinformatics Course at the Marine Biology Laboratory in Woods Hole, MA.

Todd C. Holmes is Associate Professor in the Department of Biology of New York University. His major interest is the interaction between biochemical signaling and electrical signaling as they influence neuronal circuits and animal behavior. Ionic flux across cell membranes is mediated by ion channel membrane proteins. The activity of membrane ion channels is highly plastic; their activity is regulated by a wide range of biochemical signaling molecules, including protein kinases. Dr. Holmes' laboratory is focused on unraveling the molecular mechanisms of ion channel regulation, and the physiological consequences of this regulation. Recently, he has begun to engineer ion channels that exhibit novel regulatory properties. These modified ion channels are being introduced into transgenic animals in order to determine how systematic changes in cellular electrical activity determines circadian behavior, neuronal physiology and development. His interests in protein engineering and neurobiology extend to studies of peptide-based biomaterials. He has identified a unique class of biomaterials that mimics many of the features of the extracellular matrix. These materials are being developed to serve as artificial scaffolds for tissue engineering and transplantation.

Nancy Kopell received her A.B. from Cornell University in 1963 in mathematics, and her Ph.D. from U.C. Berkeley in 1967, in mathematics with a specialization in dynamical systems. She has taught at MIT as a C.L.E. Moore Instructor, and is currently W.G. Aurelio Professor of Mathematics and Science at Boston University. At BU, she co-directs both the Center for BioDynamics and the Program in Mathematical and Computational Neuroscience, programs intended to train young scientists to work at the interfaces among different scientific disciplines. She is a member of the National Academy of Sciences, the American Academy of Arts and Science, and is a former John D. and Catherine T. MacArthur Fellow. Her major current interest is dynamics of the nervous system, especially rhythmic behavior in networks of neurons. Rhythms have been known in the nervous system for about three quarters of a century, but it is still mysterious what biophysical mechanisms produce them, and what functions they serve. In the last decade, there have been many papers linking rhythms at different frequencies to attention, perception, learning and recall, as well as motor behavior. Synchronous assemblies of neurons are thought to be important for distributed processing in the nervous system,

including "binding" of activity from different parts of the nervous system, gating incoming signals, potentiating outgoing signals and facilitating plasticity. Some of the specific projects to which she has contributed include work on the biophysical substrate of network coherence, creation and modulation of cell assemblies, and synchronization across distances. A long-range goal is to understand how the dynamical properties of local networks help to filter and transform the patterned input from other parts of the nervous system, to provide clues to the function of dynamics in the nervous system. She continues to be interested in Central Pattern Generators, networks of neurons that govern rhythmic motor behavior, as well as in geometric theory of singularly perturbed systems.

Alan Leshner is Chief Executive Officer of the American Association for the Advancement of Science and Executive Publisher of the journal *Science*. Prior to coming to AAAS, Dr. Leshner was Director of the National Institute on Drug Abuse (NIDA) from 1994-2001. One of the scientific institutes of the U.S. National Institutes of Health, NIDA supports over 85% of the world's research on the health aspects of drug abuse and addiction. Before becoming Director of NIDA, Dr. Leshner had been the Deputy Director and Acting Director of the National Institute of Mental Health. He went to NIMH from the National Science Foundation (NSF), where he held a variety of senior positions, focusing on basic research in the biological, behavioral and social sciences, science policy and science education. Dr. Leshner went to NSF after 10 years at Bucknell University, where he was Professor of Psychology. He has also held long-term appointments at the Postgraduate Medical School in Budapest, Hungary; at the Wisconsin Regional Primate Research Center; and as a Fulbright Scholar at the Weizmann Institute of Science in Israel. Dr. Leshner's research has focused on the biological bases of behavior. He is the author of a major textbook on the relationship between hormones and behavior, and numerous book chapters and papers in professional journals. He also has published extensively in the areas of science and technology policy, science education, and public engagement with science.

Tom Mitchell is Professor and Chair of the Department of Machine Learning, School of Computer Science, at Carnegie Mellon University. His research interests include Machine Learning, Computer Science, Cognitive Neuroscience, and questions such as: How can we make computers improve automatically from experience? How can computers learn to decode a person's mental state from their brain activity? How can computers learn to extract information from the web? What is machine learning all about? Where is the study of intelligence headed?

Partha Mitra is head of a laboratory at Cold Spring Harbor Laboratories. Following graduate work in theoretical physics (Harvard, 89-93) Partha worked in quantitative neuroscience and theoretical engineering, at Bell Laboratories ('93-'03). He works collaboratively with groups at Caltech ('96), NYU, Princeton and Cornell University Medical School. In the fall of 2003 Partha moved his research effort to Cold Spring Harbor Laboratories, where he is currently based. He co-directs the neuroinformatics summer course at the Marine Biological Laboratories. The Mitra laboratory is a distributed research effort based at the Cold Spring Harbor Laboratory, with close collaborative ties to several other research groups. The broad goal of Dr. Mitra's research effort is to achieve theoretical understanding of complex biological systems, particularly neural systems. The research projects are of three types: Informatics: development of algorithmic and software tools for data analysis and inference; Experimental: electrophysiological and behavioral studies; and Theoretical Engineering: development of a theoretical framework for an understanding of the design principles associated with the functioning of biological systems.

Lois Peters is Associate Professor and Director Center for Science and Technology Policy, Lally School of Management Technology, Rensselaer Polytechnic Institute (RPI), NY and is Principal Investigator at the RPI NSF-sponsored Nanoscale Engineering Research Center. She is past director of the Lally School PhD program. In addition she is past president of the International Trade and Finance Association and a member of the IEEE Engineering Management Society Board of Governors. As a member of the Board of

Governors, Institute of Electrical and Electronics Engineers (IEEE) Engineering Management Society, and past Vice President of Conferences, Peters has organized three international conferences related to the management of technology and innovation. Professor Peters' current areas of interest lie in R&D globalization, Breakthrough Innovation and practices and conditions shaping emerging technology commercialization. She is a co-author of the Harvard Business Press Book, *Radical Innovation: How mature companies can outsmart upstarts* and is currently writing a second book on corporate transformation through breakthrough innovation. Peters has conducted extensive research on university industry research connections and R&D technological networks and has participated in numerous conferences focusing on technology policy and management of innovation. She has been an invited speaker in Japan, the OECD, European Community (EC), Latin America, and Thailand among other places. For five months in 1992, Peters was an invited visiting professor at the Max-Planck-Institute für Gesellschaftsforschung, contributing to their studies on technological innovation and learning their approaches to network analysis. She teaches courses in business implications of emerging technologies, technological entrepreneurship, innovation organization and change, and technological change and international competition. Professor Peters has a PhD in Biology and Environmental Health Science from New York University.

Sara A. Solla is a theoretical physicist trained in statistical physics. Dr. Solla joined Northwestern University (NU) as a Professor in 1997 with academic appointments in the Department of Physiology and the Department of Physics and Astronomy. She is also a member of the NU Interdepartmental Neuroscience (NUIN) program. Before joining the NU faculty, Sara was for 11 years Staff Member at AT&T Bell Laboratories, where she established a reputation as a leading scientist in neural networks research, specifically in the theory of learning in adaptive systems. At NU, Sara has developed a research program in computational neuroscience, in close collaboration with experimental groups that conduct both in vivo and in vitro studies of visual and auditory sensory processing, motor control, neuromodulation, and working memory. Her work makes combined use of concepts and techniques from statistical physics, statistical inference, information theory, nonlinear dynamics and theoretical computer science, to model the neurobiological processes of encoding, storage, retrieval, and utilization of information, both at the cellular and the network levels. Her current interests include memory formation, the organization of the early stages of sensory processing, the planning and execution of movement, and pathological behaviors associated with network damage.

Charles F. Stevens directs the Molecular Neurobiology Laboratory at Salk Institute for Biological Studies. The work in Dr. Stevens' laboratory centers on mechanisms responsible for synaptic transmission. These problems are approached by a combination of molecular biological, electrophysiological, anatomical, and theoretical methods. Dr. Stevens' group studies neurons both in dissociated cell culture and in brain slices, and also investigates the function of individual membrane proteins of importance for synaptic transmission. One main current research focus is the various mechanisms used by the central nervous system for the short- and long-term regulation of synaptic strength. A second principal project uses a combination of methods to elucidate the molecular basis to neurotransmitter release at synapses.

Jonathan V. Sweedler is Lycan Professor at the University of Illinois at Urbana-Champaign (UIUC) in the analytical area of the Department of Chemistry, and serves as director of the UIUC Roy J. Carver Biotechnology Center (<http://www.biotech.uiuc.edu/>) and the UIUC Neuroproteomics Center on Cell to Cell Signaling (<http://neuroproteomics.scs.uiuc.edu/>). He is affiliated with the Departments of Physiology and Bioengineering, the Neuroscience program, and the Beckman Institute of Science and Technology. Prior to joining the faculty at Illinois, Professor Sweedler received his B.S. in Chemistry from the University of California at Davis in 1983 and his Ph.D. from the University of Arizona in 1989. Sweedler has authored or coauthored over 190 peer-reviewed publications, has thirteen patents issued or applied for, and has delivered over 250 invited lectures to universities, companies, and at scientific meetings.

Sweedler's research emphasizes analytical neurochemistry. Sweedler and his group have developed analytical methods to assay complex microenvironments including capillary electrophoresis separation and laser-based detection, mass spectrometric sampling, nanoliter volume nuclear magnetic resonance, and microfluidic/nanofluidic devices. A subset of Sweedler's research is designed to understand the molecular (chemical) nature of learning and memory. By advancing the instrumental capabilities in separation science, significant gains have been made in understanding the distribution and release of neurotransmitters from individual cells in several invertebrate model systems. Sweedler has received numerous awards including the ACS Analytical Division Arthur Findeis Award, the Benedetti-Pichler Award in Microanalysis, the Gill Prize in Instrumentation and Measurement Science, the Merck Prize, and the Instrumentation Award from the Analytical Division of the ACS. For 2007, he has been selected as the Theophilus Redwood Lecturer, Analytical Division, Royal Society of Chemistry, and will receive the Pittsburgh Analytical Chemistry Award, SACP. He has, or is currently serving on, the editorial boards of the *JACS*, *Analytical Chemistry*, *Analytical and Bioanalytical Chemistry*, *Electrophoresis*, *Analytica Chimica Acta*, *The Journal of Microcolumn Separations*, *The Analyst*, and *The Journal of Separation Science*.

C. C. Wood is Vice President of the Santa Fe Institute. Dr. Wood received his Ph.D. from Yale University in 1973. Following a postdoctoral appointment at Walter Reed Army Institute of Research in Washington DC, he returned to Yale as a faculty member with joint appointments in the Departments of Psychology, Neurology, and Neurosurgery. Dr. Wood left Yale in 1989 to lead the Biophysics Group at Los Alamos National Laboratory, a position he held until becoming the Santa Fe Institute's Vice President in 2005. At Los Alamos, Dr. Wood's group was responsible for a wide range of biophysical and physical research, including protein crystallography, quantum information, and human brain imaging. During 2000-2001, he served as interim director of the National Foundation for Functional Brain Imaging, a collaboration involving Harvard / Massachusetts General Hospital, University of Minnesota, and a number of academic and research institutions in New Mexico devoted to the development and application of advanced functional imaging techniques to mental disorders. Dr. Wood's research interests include imaging and modeling the human brain, computational neuroscience, and biological computation.