Brandeis University

The Volen National Center for Complex Systems

August 2000

The M.R. Bauer Foundation Colloquium Series, Distinguished Lecturer Series, and Scientific Retreat

The 2000 M.R. Bauer Foundation Colloquium Series, Distinguished Lecturer Series, and Scientific Retreat

Brandeis University

Benjamin and Mae Volen National Center for Complex Systems

Table of Contents		
Introduction	3	
The M. R. Bauer Colloquium Series		The M.R. Bauer Distinguished Guest Lecturer Series
	5	
Michael J. Tarr		Introduction
Department of Cognitive and		J. Anthony Movshon
Linguistics Sciences		Center for Neural Sciences
Brown University		and Psychology
Providence, Rhode Island	-	New York University
Martha Capatantina Datas	1	New York, New York
Department of Pieleav		Michael M. M.
Massachusette Institute of		Michael M. Merzenich
Technology		and Physiology
Boston Massachusetts		Keck Center
	9	University of California
Carol A. Barnes	•	San Francisco, California
Center for Neural Systems,		can hanolooo, camornia
Memory, and Aging		The 2000 Volen National
University of Arizona		Center for Complex Systems
Tucson, Arizona		Scientific Retreat
	10	
John H.R. Maunsell		
Howard Hughes Medical Institute		
Division of Neuroscience		
Baylor College of Medicine		
Houston, Texas	10	
Patricia Coldmon Dakia	12	
Section of Nourobiology		
Vale University		
School of Medicine		
New Haven, Connecticut		
e e e e e e e e e e e e e e e e e e e	14	
David W. Tank		
Biological Computation Research		
Department		
Bell Laboratories		
Lucent Technologies		
Murray Hill, New Jersey		

15

16

17

We are very pleased to present this year's proceedings of the M.R. Bauer Foundation Colloquium Series, Scientific Retreat, and Distinguished Guest Lecturer Series. Now in its sixth year, the generous support of the M.R. Bauer Foundation has permitted the Volen Center to make emerging knowledge in the guickly moving field of neuroscience available to a broad audience in the scientific community. Once again, the range of topics and importance of findings have proven to be impressive. We are very grateful to the Bauer Foundation for providing the resources to disseminate new information in a variety of useful forums. The colloquia, visiting faculty, and retreat have each helped to draw neuroscientists together from throughout the country and to facilitate the exchange of new ideas and methods. The positive effect of these programs on research and training is clear in a number of ways.

The 1999-2000 M.R. Bauer Colloquium Series again featured talks by some of the most outstanding neuroscientists at universities, medical schools, and in private industry. As in past years, the topics ranged widely, from the role of receptors during the development of connections between brain cells, to the molecular structure of memory in the brain and the mechanisms involved in continuing nerve cell activities. Professor Carol Barnes, who serves at the University of Arizona's Center for Neural Systems, Memory, and Aging, delivered a talk on "Aging and the Hippocampus: from Neural Plasticity to Ensemble Dynamics." Professor Martha Constantine-Paton, from MIT's Department of Biology, addressed the "Function and Regulation of the

NMDA Receptor during Synaptogenesis in the Superior Coliculus." Professor John Maunsell, a Howard Hughes Medical Institute Investigator at the Baylor College of Medicine's Division of Neuroscience. spoke about "Effect of Attention on Sensory Representations in the Monkey Visual Cortex." Dr. David W. Tank, from the Biological Computation Research Department at Lucent Technologies's Bell Laboratories, examined the topic, "Cellular and Network Mechanisms of Persistent Neural Activity." Professor Patricia Goldman-Rakic of Yale University's School of Medicine, presented a talk on "Microstructure of Working Memory." The Bauer Colloquium Series demonstrated the quickening activity in a number of important areas of neuroscience, as new methods and new information came out of laboratories in a variety of settings.

The M.R. Bauer Distinguished Guest Lecturer Series completed its second year and again served as a highlight of the Volen Center's programming. The program brought to campus two neuroscientists whose work has affected the entire field. Dr. J. Anthony Movshon, a Howard Hughes Medical Institute Investigator and professor of neural science and psychology at New York University, is a leader in characterizing the brain's visual system-understanding how the brain receives and interprets visual information and then uses this information to control behavior. Dr. Michael Merzenich, a professor of otolaryngology at the University of California, San Francisco, and a member of the National Academy of Sciences, discovered that the brain can be modified by experience, a plasticity now regarded as a general

property of the cerebral cortex. He then used this discovery to design an implant for alleviating deafness, a training method for dyslexic and languageimpaired children, and models for stroke recovery. Movshon and Merzenich share an unusual facility with interdisciplinary research and teaching. They have used a wide variety of specialties in neuroscience to make progress in understanding how the brain works. Their public lectures at Brandeis University, presented for a more general science audience, give an overview of the large problems on which they are working. Movshon's lecture, "Deconstructing Synchrony," describes the binding problem in visual perception-how the brain puts together a complete image from the pieces provided by each neuron-and offers some strategies for solving it that do not require an excessively elaborate neural code. Merzenich's lecture focuses on his work on dyslexia and speech problems. His large study involving more than 30,000 children has helped him to understand how to apply to dyslexia repetitive exercises that can serve to re-map the brain. He shows in this talk that the timing of acoustic signals is vital to understanding and ameliorating this problem.

Highlighting "Twenty-First Century Technologies," the 2000 Volen Center Retreat sponsored by the M.R. Bauer Foundation was held at the Marine Biological Laboratory at Woods Hole, Massachusetts, on April 25 and 26. This year's event featured talks by four scientists from outside Brandeis University. Each scientist specialized in presenting a different technology or method that is transforming neuroscience as it enters the new century. The keynote speaker was Professor Edward Farhi from MIT's Center for Theoretical Physics. He raised an intriguing possibility in "What You Could Do with a Quantum Computer If You Had One." In addition, Professor Bruce Birren, from the Whitehead Institute's Center for Genome Research at MIT, presented a talk on "The Human Genome Project: Not a Thousand People with Pipettes." Professor John Chapin from Hahnemann University's Department of Neurobiology and Anatomy spoke about "Using Neuronal Populations to Control External Devices." The final talk was given by Professor Bonnie Berger from MIT's Mathematics Department and Laboratory for Computer Science. Her talk was "Mathematical Challenges in Protein Motif Recognition." The Volen Center Retreat again provided an excellent opportunity for faculty and students to learn about cutting-edge developments in the field, to meet new colleagues, and to gain exposure for their work. It is an especially important occasion for graduate students and postdoctoral researchers as they begin to establish a professional presence in their fields. On top of that, the retreat encouraged cross-disciplinary discussions among basic researchers in the life sciences, chemistry, computer science, psychology, and physics, as well as with their colleagues in medical schools and industry.

The publication of these proceedings is a key part of the effort to make the Volen Center's work and mission widely accessible to the scientific community. We seek to encourage active discussions and collaborations among the many neuroscientists interested in the fundamental issues and emerging results addressed in these summaries. The M.R. Bauer Foundation Colloquium Series, Scientific Retreat, and Distinguished Guest Lecturer Series have successfully served to advance the ongoing conversation among neuroscientists about the most important new and ongoing issues in the field. On behalf of our colleagues and the many participants in these Foundation-sponsored programs, we would like to express our sincere appreciation to the M.R. Bauer Foundation.

Laurence F. Abbott Nancy Lurie Marks Professor of Neuroscience and Director, Volen National Center for Complex Systems

Leslie Griffith

Associate Professor of Biology and Acting Director, Volen National Center for Complex Systems Michael J. Tarr, Ph.D.

Professor of Cognitive Science Brown University Providence, Rhode Island October 4, 1999

Visual Object Recognition: Can A Single Mechanism Suffice?

How do humans recognize threedimensional objects? This simple question leads to surprisingly complex answers. At the heart of what makes visual recognition a difficult problem are two issues. First, we live in a world made up of three-dimensional objects, yet only receive twodimensional stimulation on our retinae as sense input. Second, we live in a highly variable world in which images of objects change constantly due to transformations in size, position, orientation, pose, color, lighting, and configuration. The challenge is to derive a consistent mapping from a potentially infinite set of images to a relatively small number of known objects and categories. It is a problem that the human visual system routinely and effortlessly solves.

To begin to understand how visual object recognition works, we must consider three factors that affect human recognition behavior: (1) the image geometry for objects and object classes and how it changes with changes in three-dimensional orientation, illumination, etc.; (2) the level of categorization required for a given task, varying from coarse "basic-level" categorization to fine item-specific recognition; and (3) the differing degrees of perceptual expertise that observers have with specific object classes and how visual experience fine-tunes the recognition system to attain such expertise.

The research my collaborators and I have pursued is aimed at elucidating how these factors interact to produce recognition competence across a wide

range of contexts. When considering image geometry the most salient example of how variability in the world affects image structure is a rotation in depth-even small rotations may produce dramatically different input patterns. Vision researchers have typically assumed that in order to recognize objects across such variability, object representations must be viewpoint independent. To test this hypothesis, we explored how perceivers generalized from familiar views of novel objects to unfamiliar views of the same objects. Surprisingly, recognition performance was not only viewpoint dependent, but, critically, was best predicted by the distance between the unfamiliar viewpoint and the nearest previously seen viewpoint. This finding led to the hypothesis that three-dimensional objects are represented as sets of viewpoint-specific descriptions, "multiple-views," rather than viewpoint-independent models.

In subsequent research we have considered how a viewpointdependent system can account for recognition at different categorical levels. My earlier work had established only that viewpointdependent mechanisms are implicated in more subordinate-level, within-category recognition tasks. Common wisdom associated such findings with image-specific templates that could not support the many-toone mapping required for object categorization. To test this claim, we investigated how perceivers recognized novel objects that corresponded to distinct perceptual classes. Recognition performance was found to be viewpoint-dependent despite the fact that individual objects were composed of relatively

discriminable features and parts. We also obtained evidence that perceivers were able to generalize from known instances of a perceptual class to unknown instances of that same class, but, again critically, only at familiar viewpoints. Thus, it appears that viewpoint-dependent recognition procedures are not limited to item-specific recognition, but rather include mechanisms for identification at more general categorical levels.

The fact that perceivers can recognize objects at multiple categorical levels has important implications for how we understand apparent dissociations based on object class. In particular, it has been argued that face and non-face object recognition is subserved by separable systems. Three sources of evidence are cited in favor of this stance: behavioral effects that are face-specific; brain imaging studies that show a specific neural substrate for face processing-the "face area"; and brain-injured prosopagnosic patients that appear disproportionately impaired at face recognition relative to non-face object recognition. Studies producing such evidence, however, have typically compared the recognition of faces to the recognition of common objects without controlling for either the level of categorization or the level of expertise across stimulus classes.

To address these confounds we designed a set of novel objects, "Greebles," that formed a hierarchically organized homogeneous class similar to that formed by faces. We then used these objects to test one of the most widely cited pieces of behavioral evidence for face-specific mechanisms— "configurable sensitivity."

In a part recognition task we found that Greeble experts, but not Greeble novices, showed a pattern of configurable sensitivity similar to that observed for faces. We also used brain-imaging studies (fMRI) to investigate the neural substrates mediating visual recognition. In a study using common non-face objects we found that the brain regions associated with face processing were similarly active for within-category recognition judgments. This common pattern indicates that it was the specificity of face recognition, not the class per se that led to enhanced activity in the face area.

In a second study we imaged Greeble novices and Greeble experts and found that the experts, but again not the novices, showed enhanced activity in the face area, even when this region was defined individually for each subject. Finally, we used Greebles and other non-face objects to explore prosopagnosic subjects' sensitivity to different levels of categorization.

Crucially, we used measures that took into account possible biases and differences in the effort spent recognizing different stimulus classes. Across several experiments we found that, independent of object category, prosopagnosic subjects were far more sensitive to the manipulation of the level of categorization as compared to control subjects. Thus, apparent face recognition deficits may be better explained as deficits in recognizing objects at more specific levels of discrimination.

Taken together, our behavioral, imaging, and neuropsychological work serves to implicate both the level of categorization and the level of perceptual expertise as important factors in visual recognition tasks. It is our hypothesis that the interaction of these two factors is sufficient to explain the apparent specialization of face recognition mechanisms. Overall, there appears to be little reason to posit separable recognition systems along the lines of viewpoint dependency, level of categorization. or stimulus class. Rather, humans appear to have a single highly adaptable visual recognition system that can be fine-tuned by experience to support a spectrum of recognition behaviors. Although there is clearly much work to be done, we have begun to illuminate some of the properties of this remarkable system.

Martha Constantine-Paton, Ph.D.

Professor of Biology Massachusetts Institute of Technology Cambridge, Massachusetts January 24, 2000

Function and Regulation of the NMDA Receptor during Synaptogenesis in the Superior Coliculus

Local control of protein synthesis through synaptic function is becoming recognized as an important signaling system for producing long-term changes in synaptic function. My talk focused on our studies of a mechanism that controls synthesis of the Ca++ Calmodulin dependent protein kinase (CaM KII). This enzyme has been implicated in increasing the synaptic response to glutamate during cellular models of learning and memory such as LTP. In our studies of developing synapses in the visual pathway we discovered that the normal expression of this enzyme is retarded by blocking the NMDA subtype of glutamate receptor. I discussed in my presentation our studies of a mechanism that might be responsible for this result. The content of this talk is briefly outlined below.

We discovered that eukaryote Elongation Factor 2 (EF2), a major control point in protein translation, became phosphorylated within minutes of activating NMDA receptors in the visual tectal lobes of young frogs. In collaboration with Angus Nairn of Rockefeller University, who developed an antibody specific for phospho-EF2, we demonstrated that EF2 indeed becomes phosphorylated within 30 seconds of NMDAR activation in tadpole tecta maintained in vitro. Furthermore, photopic stimulation of the retina for 30 seconds significantly increased phospho-EF2 in the retinotectal neuropil in intact tadpoles, and this increase required NMDAR function in that tectum. Using light and electron microscopy we documented that this phosphorylation occurs in the immediate post-synaptic process; that compared to levels in the tadpole, NMDA-induced phophorylation of EF2 is much reduced in the adult frog; and that in the frog, phospho-EF2 is retricted to the most distal segments of tectal neuron dendrites in layer 9A. This lamina receives a dense, indirect input from the ipsilateral eye via the nucleus isthmi and is likely to be the site of most of the structural rearrangement of synapses in frogs after metamorphosis. EF2 is phosphorylated by an EF2 specific Ca⁺⁺ Calmodulin dependent kinase, which is also localized in dendrites.

These experiments demonstrated a direct signaling pathway between the NMDARs and a major control point in protein translation. Furthermore, this control was exerted during physiological stimulation of the visual pathway. It was downregulated when plasticity was downregulated. It occurred with an exceptionally short latency. It could occur in isolated segments of tectal neuron dendrites. These findings in conjunction with our data suggesting that NMDAR function was necessary for the normal maturation of the sSC neuropil in the neonatal rat motivated a new set of studies on NMDAR-induced effects on protein translation in the rat.

Using isolated synaptic fractions (synaptoneurosomes) from P13 rat pup sSC, we found that the same AP5 blockable NMDAR stimulation used in the frog tectum caused a rapid < 1 min onset and short-lived five to 10 minute phosphorylation of EF2. Coordinated ³⁵S-methionine pulsechase labeling experiments showed that the known effect of eEF2 phosphorylation, namely a brief block

of protein synthesis, occurred in the synaptoneurosome preparations. However, while the synthesis of the majority of proteins was decreased, some proteins showed increased synthesis. At this point we knew that CaM KII expression and maturation were linked to NMDAR stimulation in the developing rat sSC neuropil; that CaM KII transcript was prominent in dendrites; and that the kinase is intimately associated with plasticity at glutamatergic synapses. We developed an immunoprecipitation assay for 35 S-labeled CaM KII and showed that within the short latency time window when phospho-EF2 levels are high following NMDAR stimulation, ³⁵S-labeled CaM KII significantly increased. Moreover, this brief (30 sec) period of NMDAR activation also increased total CaM KII protein in synaptoneurosomes by nearly 50 percent. We also showed that a number of proteins do not increase their synthesis in sSC synaptoneurosomes in response to NMDAR stimulation and that the immunoprecipitation of alpha Cam KII co-precipitates several other 35Slabeled proteins. The 35S-labeled proteins in the immunoprecipitated complex differ between P8 and P13, suggesting a developmental change in NMDAR-mediated dendritic protein synthesis. A means of specifically blocking EF2 kinase was not then available. Consequently, to test the prediction that it was the NMDARmediated phosphorylation of EF2 and the slowing of protein translation that phosphorylation of EF2 is known to produce that caused the upregulation of the synthesis of certain proteins. including CaM KII we used cvcloheximide. Cvcloheximide blocks protein translation through a mechanism independent of EF2. Low

doses of cycloheximide applied to sSC synaptoneurosomes increased CaM KII synthesis while reducing synthesis of total protein by 90 percent. The actual mechanism of this effect of slowing protein translation is unknown, though it has been seen before for other proteins in fibroblasts. Our favored hypothesis is that the slowing of translation shifts the rate limiting step in translation from initiation to elongation. This situation would favor a significant increase in translation of those transcripts that are highly abundant but poorly initiated over the translation of proteins from fewer transcripts with a high affinity for the initiating complex.

In short, we have documented an extremely rapid means through which activation of the NMDAR at young visual synapses may rapidly and transiently alter the protein content of the post-synaptic process. Although other means of dendritic control of protein synthesis exist, the NR/EF2 pathway has kinetics sufficiently rapid to be titrated by synaptic activity. Moreover, the NR/EF2 pathway exerts an important control over at least one highly significant post-synaptic density protein, CaM KII.

Carol A. Barnes, Ph.D.

Professor of Psychology and Neurology Research Scientist Arizona Research Laboratories Division of Neural Systems, Memory, and Aging University of Arizona Tucson, Arizona March 13, 2000

Aging and the Hippocampus: from Neural Plasticity to Ensemble Dynamics

Over the past two decades a number of myths about the aging brain and about cognition in normal aging have been shattered. The idea that there is mandatory widespread neuron loss or dramatic cognitive deterioration during aging in healthy individuals is clearly wrong. This is not to imply that there are no changes in neurobiology or behavior over the lifespan; rather, the neural alterations can be very selective, and the cognitive changes subtle.

The study of the hippocampus, and its role in certain forms of memory, has been particularly fruitful in facilitating the understanding of the neural mechanisms of memory in rats, monkeys, and humans. In all these mammals, an intact hippocampus is necessary for the ability to navigate in extended environments. Healthy, older humans, monkeys, and rats all show poorer spatial memory of this type, than do their younger counterparts.

A number of laboratories, including ours, have conducted studies of how the aging process affects cellular and molecular mechanisms of synaptic plasticity and spatial memory in rats. These experiments have provided a framework for understanding how the brain stores and retrieves information and what biological processes may underlie the cognitive changes that are observed in mammals as they age. Alterations in cell connectivity, and brain plasticity mechanisms during aging are reviewed. More recently our research group has developed methods for recording from many single neurons in freely behaving rats that has provided an unprecedented window into changes in neural population coding dynamics in the young and aged rodent hippocampus in relation to spatial learning and memory. With these methods we have discovered what appears to be a principal neuronal population correlate of memory retrieval failure in old rats. We propose that this age-related change in the dynamics of neural coding may provide a plausible explanation for why elderly people more frequently become spatially disoriented or lost.

Finally, a new anatomical method has just been discovered that is able to detect whether single cells have been recently active in a given behavioral experience. We believe that this method has the potential to provide a bridge between what is known about the activity characteristics of ensembles of cells recorded during behavior, and what we know about multiple genes that are activated during these behaviors. It is possible now to envision whole brain imaging of neuronal activity at the level of individual cells, using multiple genes as markers, with discrete temporal resolution of multiple experiences. This new cellular/molecular imaging approach should complement existing functional imaging methods, and should help achieve a more complete understanding of the systems responsible for both normal cognitive processing and the cognitive changes observed during aging.

John H.R. Maunsell, Ph.D.

Investigator, Howard Hughes Medical Institute and Professor of Neuroscience and Ophthalmology Baylor College of Medicine Houston, Texas April 10, 2000

Effect of Attention on Sensory Representations in Monkey Visual Cortex

Attention is a critical factor in determining what we perceive. At any moment we can give full attention to only a tiny fraction of the available sensory information. In a laboratory setting it is easy to show that attending to a particular spatial location improves thresholds for discrimination and speeds responses to stimuli at that location. In some situations attention can reliably make the difference between detecting a stimulus or missing it entirely. Such pronounced changes in perception must be associated with substantial changes in the way that the brain processes sensory information. Our laboratory is investigating how attention affects the way that individual neurons represent visual information.

We use microelectrodes to record the activity of neurons in the visual regions of the cerebral cortex of rhesus monkeys. Rhesus monkeys have excellent vision, comparable in many ways to that of humans. The monkey's visual system has been extensively studied, and much has been learned about its functional organization. The visual cerebral cortex, which lies at the back of the brain, contains dozens of discrete areas, each of which has its own representation of the visual scene. Each area contains neurons that are specialized for representing a particular type of visual information. For example, some areas are specialized to represent motion. Each

neuron within them responds selectively to stimuli moving in a particular direction, but is insensitive to the color, size or shape of stimuli. Neurons in other cortical areas have complementary properties, and are specialized to represent other information, such as the orientation of edges. Cortical areas also differ in the complexity of the sensory information that they represent. For example, while some cortical areas contain neurons that respond well to any contour or edge, neurons in other areas respond only to more elaborate shapes or patterns.

The responses of neurons within these specialized cortical areas are determined not only by inputs coming from the eyes, but also by top-down influences related to attention. Neurophysiological studies from many laboratories have shown effects of manipulating attention on the responses of individual neurons in the visual cortex of trained monkeys. Neuronal responses are generally stronger when the animal pays attention to that stimulus.

Some of our recent experiments have been directed at understanding how attention affects the quality of sensory signals in cerebral cortex. While it is known that attention makes sensory signals stronger, we are interested in learning whether attention makes neuronal responses more selective. Each cortical neuron is selective for particular stimulus dimensions, such as color, orientation, or direction of motion; it responds strongly only to a particular range of stimuli that matches its sensitivity. Highly selective neurons, which respond only to a narrow range of stimuli, provide the most precise information about the visual scene. We examined whether attention to stimulus orientation affects orientation-selective neurons by restricting their responses to a narrower range of orientations.

We trained monkeys to watch a display that contained two stimuli: a small grating pattern and a small patch of color. On some trials, the monkey had to report the orientation of the grating, while on others it had to report the color of the other stimulus. In either case the animal had to keep its gaze fixed on a small fixation spot at the center of the display, so that the retinal stimulation was the same in both cases. The grating was positioned to optimally activate the neuron that we recorded. By changing the orientation of the grating from trial to trial, we measured the range of orientations to which a neuron responded under two conditions, one when the animal was paying attention to the grating and the other when it was ignoring the grating and paying attention to the patch of color.

We measured the orientation selectivity of neurons in area V4, which is an important stage in cortical analysis of information about orientation and shape. As expected, responses of V4 neurons were stronger when the animal paid attention to the grating stimulus. Attention did not, however, change the selectivity of these neurons. Instead, responses to all orientations increased proportionately. Thus, attention effectively increases the gain of a neuron's response. This result suggests that what attention does to the cortical representation of the visual scene is roughly equivalent to adjusting the contrast on video display. However, this adjustment is not uniform across the whole representation; instead attention selectively enhances those parts of the scene that are of immediate importance.

This increase in the gain of neuronal responses suggests that the effect of attention is limited to making sensory responses stronger. Other experiments in our laboratory have been exploring whether all neuronal and behavioral effects of spatial attention might be explained in this way. These experiments have examined whether the benefits of attending to visual stimulus are quantitatively equivalent to making that stimulus stronger. We trained animals to do a task in which they had to release a lever when they detected that a stimulus started to move in a particular direction. We could adjust the stimulus so that the motion was easier or more difficult to detect. The visibility of the motion was varied from trial to trial, allowing us to measure behavioral performance and neuronal responses to a range of stimulus strengths. By presenting two stimuli and directing the animal's attention to one, we could measure the effect of attention on neuronal responses to different motion strengths, and the animal's ability to detect those

motions. With this design we could determine whether attention has effects that are quantitatively equivalent to presenting a stronger sensory signal. We asked: If a weak motion at the attended location produces the same neuronal response as a strong motion at the unattended location, will the animal's performance be the same in the two cases?

We recorded the responses motionsensitive neurons in the middle temporal visual area (MT), which is an important early stage in the processing of motion in visual cortex. Although neurons in MT responded more strongly when the animal attended to the stimulus, the effect of attention on MT neurons was too small to account for its effect on behavioral performance.

Because effects of attention are stronger in later stages of cortical processing, we examined the ventral intraparietal area (VIP) which represents a later stage of motion processing. Neurons in VIP are direction selective, like those in MT, but also include more elaborate response properties, such as sensitivity to proximity. In VIP, we found that many neurons were strongly modulated by whether the animal was paying attention to the motion stimulus in their receptive field. The average attentional modulation in VIP was too strong to match the behavioral effects of attention. That is, the changes in the neuronal responses suggested that the animal should have done much worse than it did when attention was directed to the incorrect location.

Although this result was unexpected, we believe it can be explained by the fact that attentional modulations grow stronger at successive levels of cortical processing. Because attentional modulations differ between levels of processing, there could be only one level that has a correspondence between attentional modulation of neuronal responses and behavior. We are currently exploring levels of motion processing between MT and VIP to see if this expectation holds up.

This work should lead to a more complete understanding of the role of attention in creating the representations that underlie sensation and perception, and the neuronal mechanisms involved. Our long-term goal is to extend the scope of this work to explore the mechanisms that transform these representations into decisions and actions.

Patricia Goldman-Rakic, Ph.D.

Professor of Neurobiology, Neurology, and Psychiatry Yale University School of Medicine New Haven, Connecticut March 17, 2000

Microstructure of Working Memory

Recent advances in integrative and cognitive neuroscience now allow penetration into the neural circuits and cellular processes that underlie human cognition and provide insights into the pathophysiology of diseases like schizophrenia. Nonhuman primates share with humans the capacity for working memory, the ability to hold items of information transiently in mind in order to meld together separate events and ideas into coherent lines of thought and communications. Metaphors for working memory include "blackboard of the mind," "mental sketch-pad" (Baddeley, 1989), and "on-line memory" (Goldman-Rakic, 1987). When we listen to human speech, we are using working memory to hold the segments of sentences "on-line" millisecond by millisecond. We employ working memory to carry forward, in real time, the subject of a sentence and associate it with verbs and objects in order to comprehend the sense and meaning of sentences. When we perform a mental arithmetic problem, recall a phone number, plan a hand of bridge or chess move, or follow a verbal instruction, we use working memory. In fact it is difficult to think of a cognitive function that does not engage the working memory systems of the brain.

Working memory differs from the traditional idea of short-term memory, which was thought of as a passageway for information to enter long-term memory. The modern concept is more dynamic and active-an information processing mechanism rather than a transitory but static state of information before it is permanently stored. In testimony to the conservation of mental processes in nonhuman primates and humans, behavioral tasks used in human imaging and in neuropsychological testing in patients are formally similar to those employed in nonhuman primate research. The cortical areas of the cerebral cortex, upon which working memory capacity depends, are well developed in the nonhuman primate (and poorly or not at all represented in lower mammals). These considerations make the macaque monkey unexcelled as an animal model of human mentation and human neurological and psychiatric disease.

Physiological and behavioral studies in the nonhuman primate are providing high resolution maps of the functional architecture of the prefrontal areas. These studies have revealed a striking modularity of function at the areal, cellular, and subcellular levels of neural organization. At the level of areal parcellation, different informational processing domains have been mapped to distinct anatomical subdivisions in prefrontal cortex such that spatial representations are processed in dorsolateral regions of the prefrontal cortex while object representations are processed in inferior regions of prefrontal cortex. These localizations have provided a framework for analysis of homologous working memory systems in the human brain by positron emission tomography and functional magnetic resonance imaging.

The areal modularity observed through physiological and behavioral analyses are supported by findings from anatomical tracing studies. These studies have revealed the intercortical networks that link distant sensory, motor, and limbic areas into domain constrained networks. Again such studies are critical to the interpretation of the landscape of cortical activations observed in noninvasive studies of human cognition.

At the cellular level, different neurons have been found to encode and process different items of information. indicating that information processing is carried out by neurons of remarkable specificity. Indeed, neurons that code only the identity of faces or only the location of an object and no other item of information have been observed in circumscribed areas of the prefrontal cortex. Some of these neurons, once triggered by a preferred stimulus, are specialized for maintaining their activity for many seconds after the stimulus has disappeared. This selective capacity is referred to as the "memory field" of the neuron, in analogy with receptive fields of sensory and motor neurons. Recent studies in this laboratory employing both single and dual unit recording reveal that neurons with shared memory fields, i.e., encode the same location of objects, are organized into microcolumns, as are

neurons which encode the features of objects or identity of faces. These studies suggest that the local organization of neurons in prefrontal cortex is similar to that in sensory systems, i.e., neurons coding the same item of information appear to be clustered and confined to a distinct location within the columnar structure of the cortex.

Electronmicroscopic,

electrophysiological, and biochemical studies on the localization of specific neurotransmitter receptors are establishing that modularity exists at the subcellular level, as they show that receptors of different subtype specificities are sequestered in different compartments of the same cell. For example, the dopamine D1 receptor is prominent in the distal portions of certain cortical neurons while the serotonergic 5-HT2 receptor is associated with more proximal elements of the same cell. Basic information on receptor localization and function is essential for understanding the signaling pathways and molecular mechanisms that are critical to optimal neurotransmission in prefrontal circuits. Thus, the study of neurotransmitters and neurotransmitter receptors in nonhuman primates are establishing a rational basis for development of drug therapies for the treatment of depression, age-related memory decline, Parkinson's disease, and schizophrenia, all of which exhibit some form of monoamine dysfunction.

Finally, as prefrontal areas are among those that have often been implicated in schizophrenia by imaging and postmortem neuropathological studies, experimental study of these areas in nonhuman primates provides a powerful approach to understanding the cellular mechanisms of psychopathology.

David W. Tank, Ph.D.

Researcher Biological Computation Research Department Bell Laboratories Lucent Technologies Murray Hill, New Jersey May 1, 2000

Cellular and Network Mechanisms of Persistent Neural Activity

Persistent neural activity is a sustained change in sodium action potential firing that has been observed in many brain areas involved in short term memory. We are studying the oculomotor neural integrator where persistent activity is a neural correlate of the short-term memory of eye position. The experimental preparation is the goldfish, which is particularly advantageous for a cellular and computational analysis of mechanisms. We find no evidence for plateau potentials or intrinsic oscillatory dynamics, two hypothesized cellular mechanisms for persistent neural activity. Partial pharmacological inactivation and analysis of changes in the rate of synaptic potentials are consistent with network mechanisms based on recurrent synaptic excitation. Finally, a visual training stimulus can profoundly alter the dynamics of persistent neural activity. The dependence on stimulus parameters suggests visual input is normally used to tune-up the neural integrator for better performance.

To test the reverberation hypothesis for persistent activity, the oculomotor integrator has been modeled as a network of conductance-based neurons interacting by recurrent excitation. The strength of synaptic feedback is tuned so that the network realizes "linear attractors" where the level of persistent activity varies linearly with the gaze position. The precision with which synaptic weights need to be fine-tuned depends on the time constant of the recurrent synapses, as it was emphasized previously by Seung. The issue of robustness for continuous attractors is discussed. These insights from computational modeling may be applied to short-term memory networks for a continuous quantity (in the form of continuous attractors) in general.

The M.R. Bauer Distinguished Guest Lecturer Series 2000

Introduction

One of our highlights again this year has been the M.R. Bauer Distinguished Guest Lecturer Series. This program, now in its second year, brought to campus two of the most outstanding neuroscientists in the world—J. Anthony Movshon, a Howard Hughes Medical Institute Investigator and professor of neural science and psychology at New York University and Michael Merzenich, a professor of otolaryngology at the University of California, San Francisco, and a member of the National Academy of Sciences.

Both guests visited Brandeis for a full week. They each gave a public talk (summarized on the following pages), spoke to classes, met formally with graduate students and postdoctoral fellows, presented at journal clubs, and met with many of the neuroscience laboratories. Their schedules were quite full, very educational, and enjoyed by all!

Anthony Movshon, Ph.D.

Professor Howard Hughes Medical Institute Center of Neural Sciences and Psychology New York University New York, New York January 31, 2000

Deconstructing Synchrony

In the early stages of visual processing, objects and scenes are represented by neurons with small visual receptive fields, which only "see" a small part of the retinal image. Each neuron can therefore provide information about local features of a scene, but to describe a scene in terms of objects requires that these features be combined. Objects can cover wide areas of visual space and be partially occluded by other objects, so the problem of binding the separate representations of object parts into coherent wholes is not a simple one. The computations involved in solving the binding problem probably take place in parallel at multiple levels of visual processing, and the problem may only be fully "solved" at a high level in the hierarchy of cortical visual areas.

Some theorists have advanced the view that binding is a special problem and requires a special solution because it seems necessary to "tag" each visual neuron to signify the object to which its activity relates. Each neuron therefore has to carry two distinct signals, one that indicates how effective a stimulus is falling on its receptive field, and a second that tags it as a member of a particular cell assembly. To make these signals distinct, von der Malsburg proposed that the "effectiveness" signal would be carried by a conventional code based simply on neural firing rates, while the "tag" signal would be created by synchronizing the spike activity of the neuron with spikes from other neurons in the same assembly. This theory has been elaborated by

Michael Merzenich, Ph.D.

Professor of Otolaryngology and Physiology Keck Center University of California San Francisco, California March 20, 2000

Cortical Plasticity: from Synapses to Maps

It has been clear for almost two decades that cortical representations in adult animals are not fixed entities, but rather, are dynamic and are continuously modified by experience. The cortex can preferentially allocate area to represent the particular peripheral input sources that are proportionally most used.

Alterations in cortical representations appear to underlie learning tasks dependent on the use of the behaviorally important peripheral inputs that they represent. The rules governing this cortical representational plasticity following manipulations of inputs, including learning, are increasingly well understood. In parallel with developments in the field of cortical map plasticity, studies of synaptic plasticity have characterized specific elementary forms of plasticity, including associative long-term potentiation and long-term depression of excitatory postsynaptic potentials.

Investigators have made many important strides toward understanding the molecular underpinnings of these fundamental plasticity processes and toward defining the learning rules that govern their induction. The fields of cortical synaptic plasticity and cortical map plasticity underlies cortical map reorganization. Recent experimental and theoretical work have provided increasingly stronger support for this hypothesis. The goal of this talk is to review the fields of synaptic and cortical map plasticity with an emphasis on the work that attempts to unite both fields. A second objective is to highlight the gaps in our understanding of synaptic and cellular mechanisms underlying cortical representational plasticity.

Singer and others, and supported by a variety of neurophysiological evidence that seems to show that neurons in the visual pathway, even at rather low levels of processing, tend to synchronize their firing under stimulus conditions that might favor perceptual binding.

First, I consider whether the theory is an a priori reasonable approach to solving the binding problem, and conclude that it is at best incomplete. Next, I ask whether spike synchrony can plausibly be used as an informational code, and conclude that encoding and decoding information in this way would be very difficult in the cerebral cortex because of the rich and massively parallel nature of the synaptic connections between neurons. I examine the experimental evidence adduced to support the synchrony hypothesis, and conclude that the evidence is largely indirect and has no proven relevance to the issue of binding per se. I next ask whether the binding problem is truly of unique difficulty and requires a unique solution, or is simply one of a number of "hard" problems in perception that have so far eluded our understanding. I finish by considering some strategies for solving the binding problem that do not require the creation of a special neural code.

The 2000 Volen National Center for Complex Systems Scientific Retreat "21st Century

Technologies"

Marine Biological Laboratory Woods Hole, Massachusetts April 25-26, 2000

On April 25 and 26, 2000, the Volen National Center for Complex Systems held its 12th annual scientific retreat. This year's retreat was titled "21st Century Technologies" and highlighted some of the exciting new technologies that are influencing current research. All four of the speakers did an excellent job at presenting their topics, which covered quantum computers, the human genome project, artificial intelligence, and robotics.

Approximately 75 people attended this year's retreat, which was held once again at the Marine Biological Laboratory (MBL) in Woods Hole, Massachusetts. The MBL facility includes lecture halls, function rooms, cafeteria-style dining, and overnight dorm room accomodations. Bringing the researchers (faculty, postdocs, and students) together off-campus for a 24-hour retreat was again tremendously successful. The MBL provides a stimulating environment for interactions as well as a scenic site for walking and relaxing. Tuesday, April 25, 2000

2:00-4:00 pm Arrival and check-in

4:30 pm Poster session and refreshments

6:00 pm Dinner

7:15 pm Keynote Speaker Edward Farhi Massachusetts Institute of Technology Center for Theoretical Physics "What You Could Do with a Quantum Computer If You Had One"

8:30 pm Music and dancing

Wednesday, April 26, 2000

7:00-8:30 am Breakfast

8:30 am

Bruce Birren Sequencing Center Whitehead Institute MIT Center for Genome Research "The Human Genome Project: Not a Thousand People with Pipettes"

9:30 am

John Chapin Medical College of Pennsylvania Hahnemann University Department of Neurobiology and Anatomy *"Using Neuronal Populations to Control External Devices"*

10:30 am Break

11:00 am Bonnie Berger Massachusetts Institute of Technology Mathematics Department and Laboratory for Computer Science *"Mathematical Challenges in Protein Motif Recognition"*

12:00 pm Lunch

1:00 pm Departure

Edward Farhi, Ph.D.

Professor, Center for Theoretical Physics Massachusetts Institute of Technology Cambridge, Massachusetts April 25, 2000

What You Could Do with a Quantum Computer If You Had One

There has been a lot of talk recently about guantum computers and here I am going to give a brief introduction to this trendy topic. At the outset I must say that nobody has yet built a quantum computer and it may be a long time before anybody builds one. The interest in this subject is in the fact that if you had a quantum computer it could do things that a conventional computer could not. A quantum computer could solve certain problems much faster than a conventional computer. This is not because the chips in a quantum computer run extremely fast. It is because a quantum computer uses a different logic than an ordinary computer and for certain problems using quantum logic this allows for tremendous algorithmic speedup.

This means that a quantum computer could execute many fewer operations to solve certain problems than a conventional computer would have to execute. This decrease in the required number of steps necessary to solve problems is what makes the study of quantum computation exciting.

We can idealize a conventional computer as a device that acts on strings of bits. Each bit takes the value 0 or 1. A conventional computer takes an input string of bits and produces an output string that is the answer to a problem. The computer acts by executing a program. The program results in a sequence of instructions that tell the computer how to change the bits. Each step of the running program changes a few bits. The harder the problem the more steps are required to go from the input to the output.

Although you may think that conventional computers are really very fast, there are problems that require so many steps that a conventional computer running since the Big Bang would not have had enough time to solve the problem.

The basic building block of a quantum computer is called a qubit. We can idealize a qubit as a spin-one-half particle such as the electron. You can think of a spin-one-half particle as spinning about an axis much like the earth spins. This is where, however, the analogy to a macroscopic spinning object ends.

The spinning electron carries angular momentum. If you pick any direction and measure the component of the angular momentum along that direction you get only one of two possible values +1/2 or -1/2 (in some units.) There is no classical spinning object that has the property that the component of its angular momentum along any direction must take one of two fixed values. This is called guantization of angular momentum and is a purely quantum effect. Now the qubit in a quantum computer plays the role of the bit in an ordinary computer. Spin +1/2 along some direction can be associated with bit value 0 and spin -1/2 can be associated with bit value 1. So far a gubit sounds like an ordinary bit. The difference is that an ordinary bit can

have the value 0 or 1 but the gubit can be a quantum superposition of the two values +1/2 and -1/2. These quantum superpositions have no classical analogues. A quantum computer consists of a string of qubits. A quantum algorithm is a set of instructions for manipulating these gubits a few at a time. The rules for how the gubits can be controlled are governed by guantum mechanical law (which is perfectly well understood) and these rules are completely different from the rules that govern the behavior of bits in an ordinary computer. Quantum algorithms take advantage of quantum superposition and the basic wave like nature of quantum sytems.

In a few key cases it has been shown that the number of qubit steps required to solve a problem is far fewer than the number of steps required in an ordinary computer acting with ordinary bits. The most striking example is the ability of a quantum computer to factor a large number with very few steps. Consider a large number, say one with 200 digits, and the goal is to write this number as a product of prime factors. It turns out that the fastest existing computers cannot factor a number this big in less than a century. The best known algorithm requires too many steps. If this problem is given to a quantum computer running carefully written quantum code, however, the factors could be found with very few steps. This dramatic speedup is

Bruce Birren, Ph.D.

Whitehead Institute Massachusetts Institute of Technology Center for Genome Research Cambridge, Massachusetts

The Human Genome Project: Not a Thousand People with Pipettes

important because it demonstrates that a quantum computer could do something a conventional computer could not do in a reasonable amount of time.

Also, factoring is not just an esoteric math problem. Almost all existing encryption systems are based on the inability of computers to quickly factor large numbers. If you had a quantum computer you could break all existing codes used by the military and financial institutions! For this reason the government is very interested in knowing if a functioning quantum computer could ever be built. On the theoretical side it is important to know if there are other problems that can be solved quickly on a quantum computer. We know of examples of problems where quantum speedup is not available and accordingly there is no advantage to using a quantum computer. We need to better understand the class of problems that are amenable to quantum speedup.

On the experimental side, so far only a few-qubit quantum computers have been built. By few I mean three or four. Experimentalists are trying hard to see if they could eventually build a three or four thousand qubit machine, which is what you would need to solve a difficult problem. Building a quantum computer is actually very technically challenging and probably a new idea is needed before a big one is ever built. But I am hoping that one is built because I am sure it will have applications beyond what has already been discovered.

Traditionally, the science of genetics concerned itself with the study of one gene at a time, or on occasion, the interaction of a very few genes. However, we increasingly understand that the regulation of gene function and expression is governed through the complex interaction of the entire set of proteins expressed in any given cell. Hence a single gene view necessarily limits our ability to understand complex biological systems. Despite a century's worth of research in genetics, it was only with the advent of high-through-put DNA sequencing that we were able to generate a complete list of the genes present in any organism. Our ability to sequence entire genomes thus represents a dramatic new turn in biology and the technical advances that make this possible, and the implications of this wealth of data bear some reflection.

The history of research involving the nematode C. elegans shows three distinct phases of work, each representing increasing power to understand basic biological processes. In the first, the genetic era, we identified mutants with readily observed phenotypes, including the large class with uncoordinated movement. In the molecular era. individual genes could be cloned and sequenced and transformation allowed the direct study of gene function. However, in the genomic era, when the entire genome sequence is known, we suddenly recognize that the complete cast of genetic players contains what were vet unimagined elements. Similarly, despite the intense study of the fruit

fly Drosophila over the past 70 years, by last year when the genome sequence became available, only 2.500 genes had been identified via traditional means. With the completion of the sequence we not only knew that there were 13.600 genes, but we could classify them into functional categories and immediately recognize genes capable of driving biological processes for which no evidence had ever been seen. In addition, from the sequence we also recognize that many functions that were thought to be performed by a single protein are carried out by a family of related proteins, which has obvious implications for our attempts to understand any one of them.

One of the most challenging technical aspects of the Human Genome Project has been to scale up what had originally been a complex series of laboratory steps performed by highly trained scientists. The first need was to deeply understand the laboratory processes, the causes of variation in the process, and the main drivers of cost and data quality. From this analysis we were able to design simple, robust procedures that could be performed in an automated fashion. In this way we were able to generate large amounts of high quality data in a highly predictable manner. In a large scale project that takes place over a sustained period, a great deal of work occurs once the learning curve has been climbed. Thus the energy invested in thoroughly understanding the process is paid back in increased efficiency and economies of scale. Automated data capture and analysis become equally essential to the success of a large project.

A central goal of the Human Genome Project has been to not only obtain the human genome sequence but in so doing develop a process for sequencing that will allow us to efficiently obtain sequence for any large genome. Increasing the ease with which we obtain further genome sequence will fundamentally change the way we approach biological research. For example, having the complete gene list of an organism allows us to study gene expression by simultaneously examining all genes, rather than a single gene at a time. The sequence of each new organism DNA sequence is beginning to reveal the full depth of the diversity of life on earth, a difficult problem given that only a tiny minority of the species can be cultured in a lab. Comparative genomics, in which sequence from one organism is compared to another, not only allows us to recognize the evolutionary history of species and their component biochemical pathways but provides an important new opportunity to identify the regulatory signals embedded in DNA sequence. Non-coding sequences that have preserved similarity over evolutionary time can imply conserved function associated with gene regulation. We continue to develop faster and cheaper ways to sequence DNA with the understanding that these data and associated technologies will form the foundation of biological research in the next few decades.

John Chapin, Ph.D.

Professor Department of Neurobiology and Anatomy Medical College of Pennsylvania Hahnemann University Philadelphia, Pennsylvania

Using Neuronal Populations to Control External Devices

It has long been thought that neuroprosthetic devices might be useful for restoration of motor functions in patients afflicted with paralyzing neurological diseases. In particular, neuronal activity recorded in the motor might be useful as a source of "motor command" information to drive a neuroprosthetic device. At its simplest level, such a device could be conceived as a robotic arm capable of being controlled from a patient's motor cortex to the extent that it could reproduce some normal limb functions of the arm. Further developments would include the ability to allow a patient to control his own muscles through functional neuromuscular stimulation (FNS). Such techniques might ultimately make it possible to electronically create posture and locomotion under the direct control of the brain. Finally, it might be possible to create a sensory neuroprosthesis capable of carrying cutaneous and proprioceptive feedback from the limbs, allowing ongoing movement to be more accurately controlled. Beyond this, the development of such neuroprosthetic technologies would provide new concepts and tools for the study of information processing in the cerebral cortex, which is of immense scientific importance.

To determine whether this basic idea is feasible, we developed the following experimental paradigm: up to 46 neurons were simultaneously recorded in the forelimb motor cortex, ventrolateral thalamus and/or cerebellum of rats initially trained to obtain water by pressing a lever to control movement of a robot arm. Neuronal population activity peaked just before onset of lever movement, but this peak encoded the trajectory of movement over the next 3-500 ms.

In order to use this multi-neuronal signal to control a robot arm it was first necessary to "decode" the information present in the neuronal population. In particular, we needed to define a mathematical transformation capable of taking the signals from up to 46 separate neural recordings and using them to create a single analog signal appropriate for controlling the position of the robot arm.

First, principal components analysis (PCA) was used to extract major sources of signal in this multichannel data set, while removing noise. To appropriately transform the time domain of this signal we utilized multivariate statistical techniques, or neural networks to yield output functions that closely matched the timing, magnitude, and direction of forelimb movements. After training, the control of the robot arm was switched from the lever press to the neuronal population. Animals with at least 25 recorded neurons successfully used the first principal component of the multi-neuron population signal to move the robot to the water source and return it to their mouths.

These results showed that the animals were to use the neural population signal as a surrogate for limb movement to move the robot arm and obtain reward, even though its tendency to peak in the premovement period resulted in the delivery of water reward before the onset of movement. This therefore provided a test of whether the animal could alter its operant behavior and/or its neural activity to take advantage of the changed situation. Typically, during the first few minutes of neurorobotic mode, the animals would continue for several trials their welltrained behaviors of pressing the lever down to the original threshold position for obtaining the water reward. In fact, this was necessary considering that the peak amplitude of pre-movement neural activity was still proportional to the ultimate magnitude of the movement.

Over continued trials, however, the normally high correlation between the neural signal and the movement magnitude tended to disappear such that the neural signal became independent of the movement. The correlation usually became insignificant within about 10-20 trials (over about five minutes). After that, it would often go back to full or partial movement, but exhibited many more movements where it only reached to

Bonnie Berger, Ph.D.

Professor Mathematics Department and Laboratory for Computer Science Massachusetts Institute of Technology Cambridge, Massachusetts

Mathematical Challenges in Protein Motif Recognition

One of the most important and challenging problems in computational biology is that of predicting the threedimensional structure or shape of a protein from its amino acid sequence. As a first step to tackling this problem, many researchers have focused on the structural motif recognition problem: given a known local threedimensional structure, or motif, determine whether this motif occurs in a given amino acid sequence.

In this talk, I will present algorithms that use probabilistic techniques to improve existing methods for recognizing protein structural motifs. These algorithms are particularly effective at eliminating false positives found by previous methods without introducing false negatives.

We have implemented these algorithms and have tested them on two-stranded and three-stranded coiled coils. The coiled coil motif has many important biological roles; for example, it is found in some DNA binding proteins and plays a role in the membrane fusion of viruses such as HIV.

Our algorithms have been codified in four programs: PairCoil (1995), which uses pairwise correlations to significantly improve upon existing methods for identifying two-stranded coiled coils; MultiCoil (1997), which uses multidimensional clustering to identify and distinguish between three-stranded and two-stranded coiled coils; LearnCoil-Histidine Kinase (1998) and LearnCoil-VMF (1999), which incorporate statistical

contact the lever but not to press it. Even so, the cortical signal reached a much higher level than it would have when performing similar movements during previous training in the neurorobotic mode. Thus the neural coding necessary to obtain the reward did not require that the associated reaching movement be an operant for the previously learned behavior, and instead, the neural signal itself could be rapidly replaced as the operant.

Overall, these findings have demonstrated the feasibility of using real time neuronal population activity to control external motion devices, and also suggest techniques for extracting multiple motor codes from these populations. Such techniques for large scale chronic recording of brain neurons, and for transformation of these recordings into appropriate output signals, could someday be used by paralyzed patients to control external movement devices or even their own muscles through functional electrical stimulation. learning techniques to identify histidine kinase linker domains and viral membrane fusion proteins, respectively, for which there are limited known solved structures.

Finally, I will talk about some of the biological implications of our work. In particular, our programs have been useful in identifying coiled-coil-like motifs in the envelope proteins of many viruses, such as the influenza virus, Moloney murine leukemia virus, HIV, SIV, and visna virus, whose structures have since been solved. This in turn has led to antiviral drug discovery by the Kim lab.

(Portions of this work are joint with Peter S. Kim, Ethan Wolf, Mona Singh, David Wilson, and Andrea Cochran.)