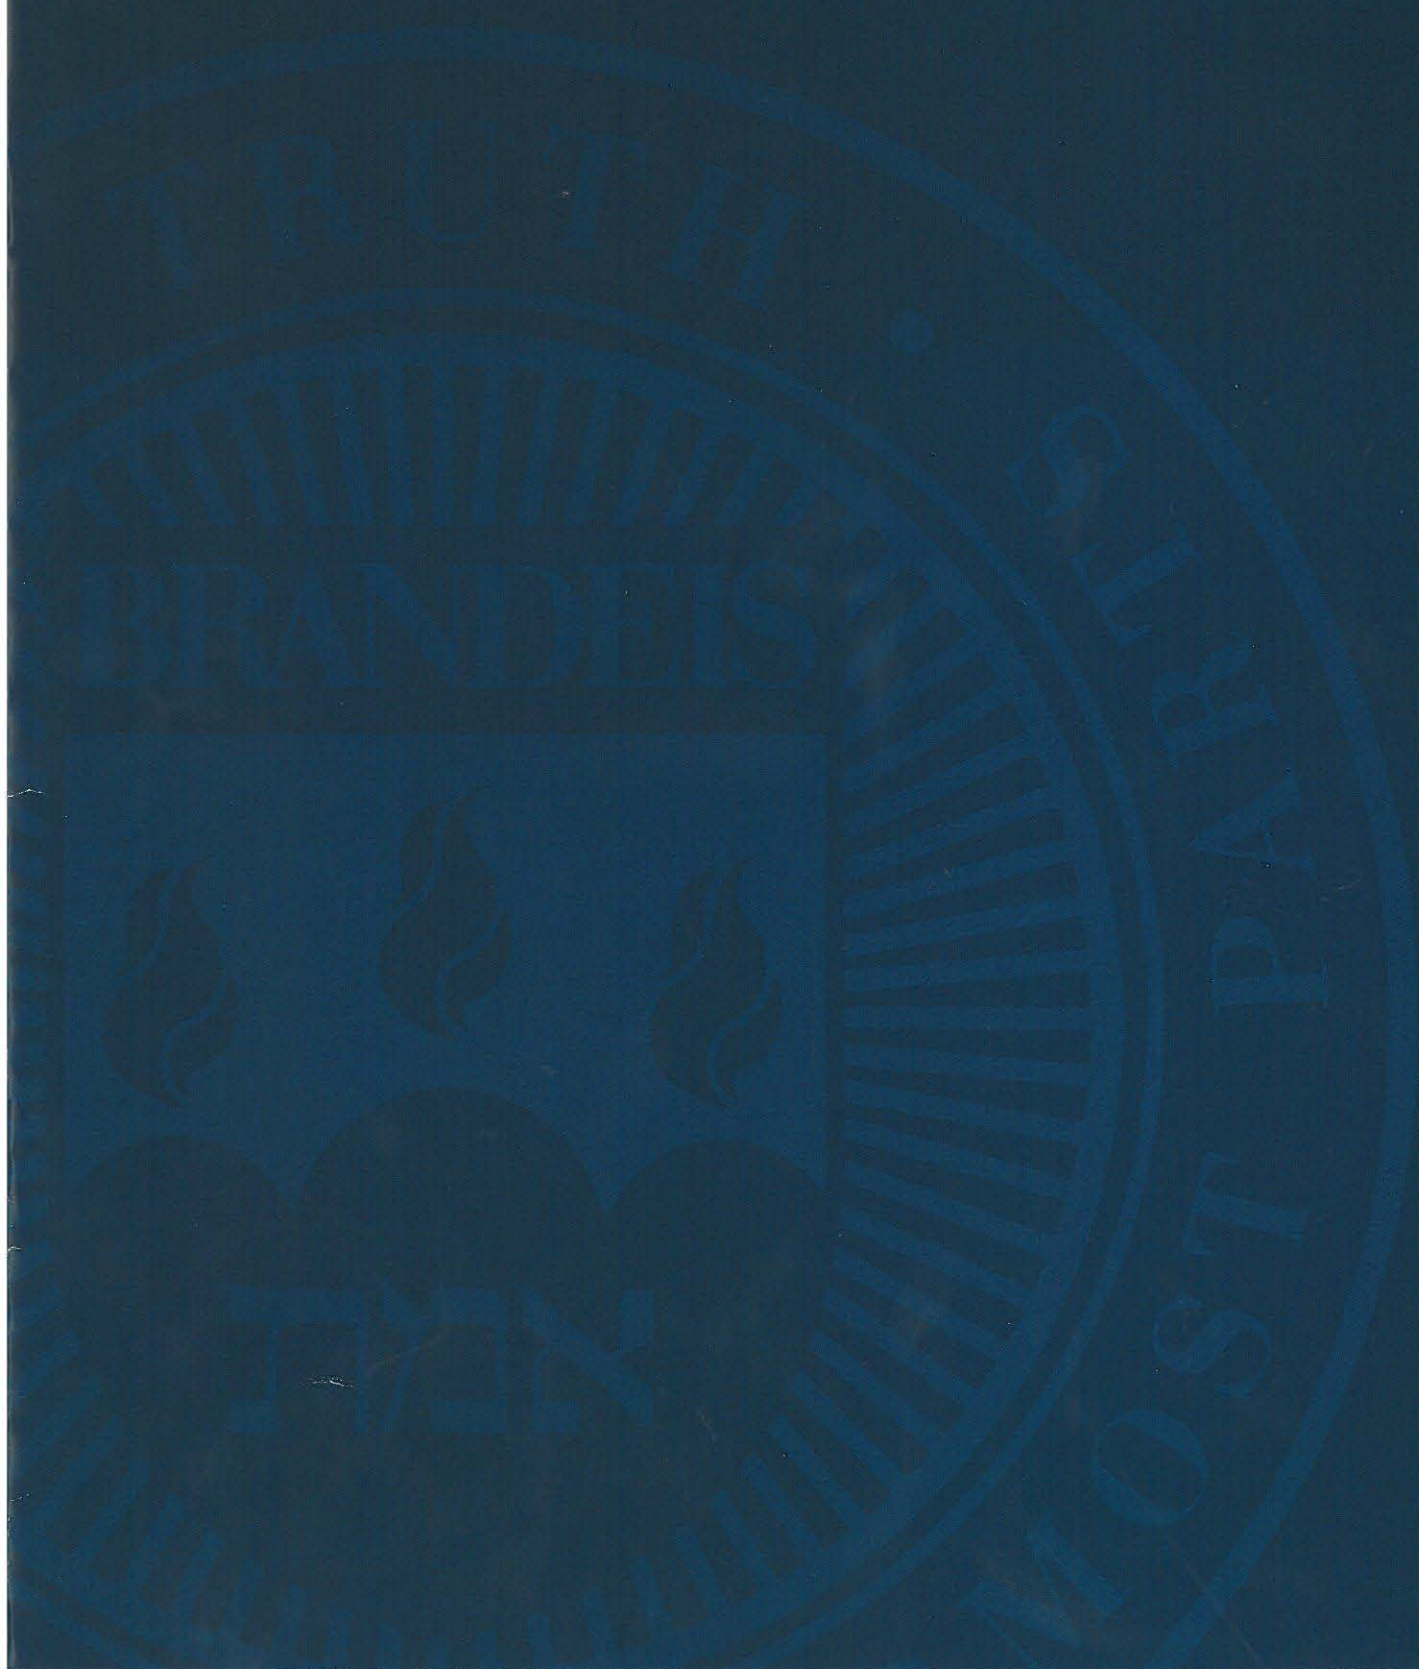Brandeis University

Benjamin and Mae Volen National
Center for Complex Systems

August 1995

The M.R. Bauer Foundation Colloquium Series and Scientific Retreat

Introduction

The 1995 M.R. Bauer Foundation Colloquium Series and Scientific Retreat

Brandeis University

Benjamin and Mae Volen National Center for Complex Systems

Through the generous support of the M.R. Bauer Foundation, the Benjamin and Mae Volen National Center for Complex Systems at Brandeis University has enhanced its conference and colloquium schedule. One of the most important duties of any academic center is to disseminate emerging information and to create a forum for the discussion of new ideas. These events create opportunities for Center faculty to share their work with the broader scientific community and to learn about the latest techniques and research projects from colleagues at other institutions.

The M.R. Bauer Colloquium Series was designed to bring to campus leading researchers in neuroscience, computer science, and cognitive science to interact with Brandeis faculty, researchers, and students. These scientists present their current projects, report their latest findings, and discuss the challenging issues that arise from their efforts. They are encouraged to visit Volen Center laboratories and exchange ideas about the work taking place there. Since funding for the Bauer Series was initiated in December 1994, five distinguished scientists have visited campus as part of this program. Topics range from molecular research to behavioral studies.

In May, the first Volen Center scientific retreat sponsored by the Bauer Foundation took place. This annual event provides Center faculty, researchers, and students with the opportunity to discuss their efforts, make research presentations, and learn more about their colleagues' work. The Volen Center was designed to bring scientists from various disciplines together in order to promote interdepartmental collaborations. The retreat presents an ideal occasion for Center scientists to become better acquainted with work in other disciplines and to discuss issues related to their research that may benefit from the perspective of their colleagues. This year's retreat was titled "The Center for Complex Systems: New Directions" and featured talks from four recently hired junior faculty members.

The M.R. Bauer Colloquium Series and the Volen Center retreat foster the exchange of ideas across traditional departmental boundaries, bring accomplished scientists to campus to share some of the most current research taking place in the field, and promote the flow of information within the scientific community. The publication of these proceedings is a key component in the Volen Center's effort to encourage scientific interaction and the sharing of knowledge.

Irwin B. Levitan
Nancy Lurie Marks Professor of Developmental Neuroscience
Director, Volen National Center for Complex Systems

Harvey Karten, Ph.D.

Professor
Department of Neuroscience
University of California, San Diego
January 5, 1995

"Evolutionary Origins of the Neocortex"

Karten began by raising the problem of the evolutionary origins of the neocortex in terms of comparative vertebrate neuroanatomy. He pointed out that most brain regions such as the cerebellum and spinal cord have a remarkably similar structure in the brains of fish, amphibia, reptiles, birds, and mammals. In contrast, a clearly identifiable neocortex is absent from nonmammalian species. In fact, the presence or absence of the neocortex is as sure a taxonomic quality of mammals as hair or mammary glands. This raises the question of which structures in nonmammalian brains gave rise to the neocortex.

Karten proceeded to review historically the neuroanatomical literature on homologies between the avian and mammalian forebrain. The avian forebrain consists of a thin pallium surrounding a very large mass of gray matter, in contrast to the large cortex and smaller central gray matter of mammalian brains. Early comparative neuroanatomists assumed that the avian central gray matter mass was homologous to the mammalian basal ganglia, because of their similar position relative to the ventricles.

Subsequent studies, however, revealed a great heterogeneity of these "striatum-like" structures. They can be roughly divided into a basal and dorsal ventricular ridge. On the basis of his own studies of the pigeon brain, using staining patterns for acetyl cholinesterase, dopamine, and substance P, and the pattern of specific sensory afferents, he concluded that it is only the basal ventricular ridge that is homologous to the mammalian basal ganglia.

Similar studies of connectivity patterns in the visual systems of birds and mammals suggested that the dorsal ventricular ridge (DVR) may bear homology to the primary input layers (layer four) of the extra striate regions of the mammalian visual cortex. Another structure, the visual wulst was found to be homologous to primary visual (striate) cortex. This suggested that the cortex may have originated not from a single structure as suggested by Allman, but from the integration of two separate structures. Similar studies of somatosensory and auditory pathways led to the same conclusion. A potential difficulty with this view is that while the wulst is like the cortex a laminated structure, the DVR is not. Karten then digressed to recount another example where clearly homologous neural structures are in one species highly laminated, while in another closely related species they are not. This example involves gustatory organs in fish. In the catfish the nucleus that receives the vagal gustatory afferents is crudely developed and non-laminated, while in the goldfish, the same region is highly developed and is fully laminated.

Karten then discussed the issue of how the DVR and visual wulst develop. Birth dating studies using bromodeoxy uridine suggested in mammals that early on there is DVR equivalent, the subventricular zone (SVZ). A common feature of cortical development is the "inside out" pattern in which deep layer cells are born earliest and superficial cells are born later. The deep layer cells contain efferent neurons that project out of the cortex, the middle layers contain recipient neurons that receive thalamic inputs, and the superficial layers contain, broadly speaking, interneurons that project within the cortex. A similar birth dating pattern of efferent, then recipient, then interneuronal populations was found in the avian brain, although here the separate populations were located in separate regions rather than in separate layers within the same region.

Karten speculated that the cell populations, which occupy different laminar positions in mammalian cortex and different regional positions in the avian wulst and DVR, may in fact correspond to neuromeres, first described by Ben Kalaine. He also raised the intriguing possibility that recently described homeobox genes, which have been shown to label neuromeric structures in other parts of the nervous system, could potentially be used to test his hypothesis of the dual origins of the mammalian neocortex. Preliminary results suggest that in fact there are HOX gene homologs that can be recognized within the avian DVR.

James A. Anderson, Ph.D.

Professor and Chair of Cognitive and Linguistic Sciences
Department of Cognitive and Linguistic Sciences
Brown University
January 26, 1995

"Teaching the Multiplication Tables to a Neural Network: Flexibility vs. Accuracy"

Abstract

There are only a few hundred well-defined facts in elementary arithmetic, but children find them hard to learn and hard to use. One reason for this difficulty is that the structure of elementary arithmetic lends itself to severe associative interference. If a neural network corresponds in any sense to brain-style computation, then we should expect similar difficulties teaching elementary arithmetic to a neural network. We find this observation is correct for a simple network that was taught the multiplication tables. We can enhance learning of arithmetic by forming a hybrid coding for the representation of a number that contains a powerful analog or "sensory" component as well as a more abstract component. When the simple network uses a hybrid representation, many of the effects seen in human arithmetic learning are reproduced, including overall error patterns and response time patterns for false products. An extension of the arithmetic network is capable of being flexibly programmed to correctly answer questions involving terms such as "bigger" or "smaller." Problems can be answered correctly, even if the particular comparisons involved had not been learned previously. Such a system is genuinely creative and flexible, though only in a limited domain. It remains to be seen if the computational limitations of this approach are coincident with the limitations of human cognition. Perhaps the ability to apply flexibly the information in a neural network to new problems is more important than overall accuracy. We will make a few speculations about how to build flexibility and programmability into a network.

Anderson's research concentrates on applications of neural networks to cognitive science. An appropriately designed network can do many pattern recognition functions in ways reminiscent of human performance. Neural networks have practical applications and can also serve as models for human behavior.

His group conducts research in several areas. Networks have been applied to models of human concept formation, to speech perception, and to models of low level vision, for example, the way local motion signals can be integrated to determine global object motion or the direction of self-motion. A current project involves the study of elementary arithmetic, a problem that is surprisingly hard for both humans and neural networks. Study of elementary mathematics also raises questions about the way a neural network can be designed to perform effectively more general mathematical operations.

Recent work has considered how intermediate level structure in the nervous system might be configured, and how it might be detected in experimental data, as well as what kind of computations it might perform.

A model using a network of local networks is being studied, in light of data from both multiple unit recordings and functional MRI.

Selected Publications

Anderson, James A. "The BSB Model: A simple nonlinear autoassociative neural network." In *Associative Neural Memories*. Edited by M. Hassoun. New York: Oxford University Press, 1993.

Anderson, James A. *An Introduction to Neural Networks*. Cambridge: MIT Press, 1995.

Anderson, James A., K.T. Spoehr, and D.J. Bennett. "A study in numerical perversity: Teaching arithmetic to a neural network." In *Neural Networks for Knowledge Representation and Inference*. Edited by D.S. Levine and M. Aparicio. Hillsdale: Erlbaum, 1994.

David S. Touretzky, Ph.D.
Computer Science Department, Center for the Neural Basis of Cognition
Carnegie Mellon University
March 23, 1995

"Multiple Representation of Space in Rats and Robots"

Abstract

There is a wealth of data on rodent performance in spatial learning and memory tasks, but as yet there is no comprehensive theory of how rodents navigate. I will describe a multiple-representation theory of navigation, developed in collaboration with David Redish and Hank Wan, that ties together behavioral and neurophysiological observations in a way that is quite explicit. The central thesis is that rodents maintain dual representations of space: one based on external cues tied to perception (the local view), the other based on internal metrical values maintained by path integration. The hippocampal system mediates between these two representations, allowing the animal to reconstruct one from the other.

Our theory allows us to reproduce a variety of rat and gerbil navigation experiments with a single computer model, called CRAWL. The model makes testable predictions about the behavior of hippocampal place cells and about the head direction cells in thalamus and postsubiculum. Portions of the model have also been implemented on CMU's mobile robot, Xavier.

David S. Touretzky is a senior research scientist in the computer science department and the Center for the Neural Basis of Cognition at Carnegie Mellon University. He received his B.A. from Rutgers University in 1978, and his Ph.D. from Carnegie Mellon in 1984. Touretzky's research focuses on the study of representations, both in computers and in nature. His earlier work was on symbol processing in connectionist networks using distributed representations. In the last few years his interests have shifted toward computational neuroscience. He is presently engaged in studies of animal learning and navigation, and the role of the hippocampus in spatial representation.

Summary

Landmark-based navigation is a rich domain for exploring issues of representation and processing in neural systems. At the behavioral level, there is a wealth of data on how animals use landmarks to locate food or return to their homes. At the neurophysiological level, the responses of hippocampal pyramidal cells, the well-known "place" cells, and of head direction cells in thalamus and the subicular comple provide striking neural correlates to behavioral variables.

Systems-level theories fill the gap between these two modes of description. To construct such a theory for navigation, we must first determine a set of computations at some reasonably abstract level that can reproduce the observed behavior, and then show how these computations could be realized in neural tissue. We are, of course a long way from this goal. However, in the present paper we describe a theory of landmark-based navigation in rodents that is constrained by both behavioral and neurophysiological data. The theory is embodied in a computer model called CRAWL, allowing us to replicate various experiments in the animal navigation literature and make predictions about place cell responses in novel situations. Portions of the model have also been implemented on a mobile robot.

Collett, Cartwright, and Smith (1986) trained gerbils to search for a food reward at a fixed position relative to an array of identical cylindrical landmarks. The array was translated but not rotated from trial to trial, and the animals were released from different starting points to ensure that the landmarks would provide the only reliable cues to the reward location. After training to criterion, probe trials were introduced in which the food was absent and the distribution of the animals' search efforts was measured.

For a single cylinder, the animals learned to search at the correct bearing and distance from that landmark. Bearing information was presumably measured with respect to the internal compass, because the arena was designed to minimize stimuli that could serve as directional cues. The room walls were painted black, and there was only a single overhead light, in the center of the ceiling. The light illuminated a circular region of the floor but left the walls in shadow.

Experiments with another group of animals using pairs of cylinders produced similar results: well-trained animals would go directly to the goal location. Now Collett et al. could test the animal's representation of the environment by introducing occasional probe trials with modifications to the landmark array.

When one landmark was missing on a probe trial the gerbils searched alternately in two locations, each at the correct bearing and distance from one of the landmarks they had observed during training. It was as if they were binding the cylinder to one and then the other of the two remembered landmarks. When the distance between the two cylinders was doubled on a probe trial (the "split landmark array" case, the animals also searched in two locations, each associated with one of the two landmarks. They did not search in the center of the expanded array.

Collett et al. proposed a mechanism that could account for these results, and several others involving more complex arrays of landmarks. We shall refer to it as the "vector voting hypothesis." According to this hypothesis, when the gerbils reach the goal location and find food there, they note the vector between each landmark and their present position. Then, at the beginning of a new trial, when they first emerge from the "start box" and see the landmark array, they apply every learned vector to every currently perceived landmark. The locations receiving the most votes are the ones they choose to search.

However, the vector voting hypothesis alone does not account for the split landmark result, where the gerbils searched in two locations instead of four. The four sites should receive one vote each when the landmark array is split. Collett et al. concluded that the gerbils must be using their perception of the entire array to distinguish the east from the west landmark, and binding remembered vectors to only the corresponding landmarks during the probe trials.

This explanation accounts for the data, but it introduces a binding mechanism whose characteristics are left unspecified. Our model reproduces all the effects in the cases above and a variety of similar experiments without resorting to explicit binding mechanisms to disambiguate landmarks.

There are several functional subsystems to our model, which should not be presumed to correspond to individual anatomical sites. VISUAL INPUT provides range and egocentric bearing information for a set of currently visible landmarks. The animal's HEAD DIRECTION, i.e., its "internal compass," is updated using vestibular cues; in the block labeled LOCAL VIEW, allocentric landmark bearings are calculated.

The animal is also assumed to maintain an estimate of its position in some internal coordinate system. This value is updated by the PATH INTEGRATOR as the animal moves, based on vestibular cues and an efferent copy of motor commands. There is evidence for path integration abilities in a wide variety of animals. Unfortunately, its precise neural substrate is presently unknown.

Finally, the role of the PLACE UNITS in our theory is to maintain an association between perceived landmark positions and path integrator coordinates, so that either can be reconstructed from the other. This explains several other questions raised by neural recording experiments using rats: How do rats self-localize when reintroduced into a familiar environment at a random spot? Our theory says they use visual

landmarks to activate a place code, which in turn evokes a set of coordinates used to reinitialize the path integrator.

How are place fields able to persist in the dark? The path integrator is updated with each motion the animal makes. Our theory says that the output of the path integrator may be used to drive place cells. Errors will eventually accumulate, but the system may be kept reasonably calibrated if other sorts of cues are available, such as tactile information.

How is drift in the path integration and head direction systems corrected? The place units representing a location keep track of the allocentric bearings and distances of landmarks visible from that location. If landmarks appear at the correct distances but their bearings are off by a consistent amount, this indicates drift in the internal compass. If the path integrator's output differs somewhat from the coordinates derived from the place units, this indicates drift in the path integrator.

For more information, see the related papers.

Collett, T. S., B. A. Cartwright, and B. A. Smith. "Landmark learning and visuo-spatial memories in gerbils." In *Journal of Comparative Physiology*. A, 158: 835-851, 1986.

Touretzky, David S., H. S. Wan, and A. David Redish. "Neural representation of space in rats and robots." In *Computational Intelligence: Imitating Life*. Edited by J. M. Zurada and R.J. Marks. IEEE Press, 1994.

Wan, H. S., David S. Touretzky, and A. David Redish. "Towards a computational theory of rat navigation." 16th Annual Conference of the Cognitive Science Society. Lawrence Erlbaum Associates, 1994.

Janet Metcalfe, Ph.D.

Professor of Psychology
Department of Psychology
Dartmouth University
April 13, 1995

"Binding in Episodic Memory"

Abstract

The construct of memory binding, whereby the features within an item or event are glued together in human memory, is investigated. It is proposed that the binding function characterizes explicit memory, whereas implicit memories may be unbound. In a series of explicit memory experiments, people studied several faces and then were tested for recognition with old faces, superimposed composites (comprised of two old faces photographically superimposed), conjunctions (comprised of the eyes and nose of one face set in the background of another old face), and new faces. Although all of the components of the superimposed composites and the conjunction faces were old, people were able to reject these lures. The probability of saying 'old' to the conjunction and superimposed composite faces was lower than it was to the studied old faces, though considerably higher than to the new face lures. Simulations of four computational models—CHARM, TODam, MINERVA, and a Back Propagation Network—showed that those models, including an operation that binds the features within an item together, produced these explicit-memory results, whereas those models without such a binding operation did not. Instead, the nonbinding models produced results like those found by Reinitz, Morrissey, and Demb (1994) in an analogous implicit-memory face-recognition task. The relation of binding and prototype extraction is discussed. It is concluded that both the construct of binding and the principle of superposition are needed to explain the human memory results.

Metcalfe's presentation was concerned with the concept of binding in memory. Binding is a term used by cognitive psychologists to describe the process of organizing lower order elements into higher order representations. Evidence from other studies has suggested that binding occurs in one type of memory but not in another. These two systems of the memory are referred to as explicit memory and implicit memory. Explicit memory is thought to involve conscious, controlled recollective processes, whereas implicit memory is thought to involve unconscious, automatic processes.

Metcalfe presented research on face recognition in normal subjects. She compared her empirical results with computer simulations of human memory. Earlier research suggested that human subjects have trouble with faces that are mixtures of previously studied, individual features such as eyes, nose, hair, and mouth. Subjects tend to recognize these types of faces, referred to as conjoined faces, as ones they had studied before when in fact they had not seen them at all. Apparently, changing the individual features of faces causes subjects to confuse these faces with previously studied faces that have not had their individual features changed. One reason may be that they invoke the higher order representations of faces that were studied previously.

What would happen if a face mimicked the higher order representation of a face even better than a conjoined face? This effect was produced by superimposing any two faces that subjects had studied previously to produce what Metcalfe referred to as a superimposed composite face. Metcalfe's research indicated that these faces, like conjoined faces, were easily mistaken for previously studied faces.

What happens when the two types of face-mixing, i.e., conjoined and composite faces, are compared? Which type of face invokes the higher order representation of a stored face in memory better? It would make the most sense for the composite faces to be most easily mistaken for unaltered faces studied before because, as discussed, they mimicked the higher order representation of a face the best. Metcalfe's research indicated that this was in fact true: subjects confuse the superimposed faces more frequently with previously studied faces than they do with conjoined faces and previously studied faces.

What processes underlie this process of binding the eyes, nose, mouth, and hair into higher order representations? Computational models of memory provide very specific equations for how memory works and therefore can possibly reveal some of the processes that underlie this ability. These models have proved successful in the past in accounting for a large variety of data on human memory and learning. After inputting the faces into the four different computational models, it was found that the models that best fit the data were models that had the ability to convolve, or autoassociate, features into a larger whole. The models that did not have this additional feature were unable to provide the results that were obtained with human subjects reported earlier here.

Another feature that proved essential in producing results similar to the results obtained with the human subjects was the capacity to store individual features all together in a single storage box, which computational memory designers refer to as a *memory vector*. Those

models that lacked this ability and designated that items be stored individually could not provide the data Metcalfe obtained. The models that fit the data obtained from research on human subjects best had the dual capacity to convolve, or autoassociate, individual features from objects together and also had the ability to store these combined features in a single storage box.

Earlier findings have suggested that binding is a process that occurs in explicit memory. This is the area from which Metcalfe's research took its lead. By studying the effect of randomly mixing features of a face such as the nose, eyes, mouth, and hair and comparing this with subject's recognition of faces that were superimposed on each other, it was found that binding is a process that involves the formation of higher order representations of faces such that the closer a face is to this higher order representation, the more easily it would be mistaken for this representation. Metcalfe's research also revealed, through the use of computational models of memory, that binding is a process that involves a combination of these features in a concrete manner and in a single storage box or memory vector.

Edward Adelson, Ph.D.

**Associate Professor of Vision Science
Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology**

"New Directions in Vision and Video"

"Compression" refers to techniques that reduce the amount of storage required to represent an image or sound. Entire families of compression techniques have been developed over the years in order to overcome particular limitations in central storage and transmission domains. The most important of these domains (from a commercial perspective) involve limited storage capacity of computer disk systems, and the difficulty of transmitting complex digital representations across transmission lines of limited bandwidth, e.g., voice-grade phone lines. One of the greatest challenges to image compression techniques is to transmit a sequence of full-motion digitized images at a sufficiently high rate (frames/second) to enable the receiver to reconstitute the image in real time and with the same motion characteristics as were present in the original.

Adelson divides image compression techniques into three families: low-level, mid-level, and high-level, in analogy with the three domains used by researchers to distinguish various levels of neural and cognitive processing in human vision. At present virtually all techniques in regular use exploit low-level regularities (redundancies) present in all images. These low-level compression techniques include computer algorithms such as those represented by TIFF and JPEG formats. The commercial importance of pyramidal compression schemes, such as those used by Kodak computer imagery, is beyond question. Most current image coding systems rely on signal processing concepts such as transforms, VQ, and motion compensation ("deblurring"). In order to achieve significantly lower bit rates (higher levels of compression), it will be necessary to devise encoding schemes that involve mid-level and high-level computer vision. Model-based systems have been described, but these are usually restricted to some special class of images such as head-and-shoulders sequences.

The most sophisticated imaginable compression schemes, for full motion images, would resemble the high-level, cognitive processing associated with human vision. For example, imagine that one has a video image sequence depicting a person who repeatedly opens and closes his fist. Instead of transmitting this sequence of images pixel by pixel, a high-level system would transmit the first image in the sequence and then a descriptive tag formally equivalent to "person opens and closes fist." Note that formal equivalence does not require linguistic equivalence or, for that matter, even that the tag be coded in natural language terms. The receiver of this transmission would decode the semantic instruction and apply it appropriately to the first image, thereby reconstituting this aspect of the entire sequence. Adelson notes that at present it is impossible to use such high-level compression techniques with any degree of fidelity (we lack the proper language and interpretive structures that would enable such strategies to work). However it is possible to make real progress on compression schemes that operate at an intermediate level by exploiting mid-level regularities in images, including depth information and properties of surfaces. Compression schemes that exploit surface and depth properties of images should be able to achieve far greater compression than currently achievable by use of low-level algorithms alone.

Adelson's research focuses on image sequences depicting simple, but real (not "toy") sequences of images. He treats such sequences as a three-dimensional volume, with the dimensions of x, y, and t (time). Motion analysis involves orientation-selective filtering within this volume. Standard approaches to motion analysis assume that the optic flow is smooth; such techniques have trouble dealing with occlusion boundaries. Note that occlusion may momentarily remove an object from the scene, but an effective compression scheme must continue to represent that object so that when the object is no longer occluded the scheme will treat that object as the same entity as before the disappearance.

The most popular solution to the occlusion problem is to allow discontinuity in the flow field, imposing the smoothness constraint in a piece wise fashion. But there is a sense in which the discontinuity in flow is artifactual, resulting from the attempt to capture the motion of multiple overlapping objects in a single flow field. So Adelson decomposes the image sequence into a set of overlapping layers, where each layer's motion is described by a smooth flow field. The discontinuity in the description is then attributed to object opacities rather than to the flow itself, mirroring the structure of the scene.

Adelson has been using mid-level vision concepts to achieve a decomposition that can be applied to many domains of image material. He described a coding scheme based on a set of overlapping layers, i.e., a scheme in which a scene was automatically segmented into layers, much as it is believed the human visual system does. The layers, which are ordered in depth and move over one another, are then composited in an animation as used by Walt Disney Studios and others.

Based on these ideas, Adelson demonstrated a set of techniques for segmenting images into coherently moving regions using a fine motion analysis and clustering techniques. This allowed him to decompose an image into a set of layers along with information about occlusion and depth ordering. Adelson applied the techniques to the "flower garden" sequence (an industry-wide standard set of images that are a benchmark for compression work). They analyzed the flower garden scene into four layers, and represented the entire 30-frame sequence with a single image of each layer, along with associated motion parameters. The next step is to develop early and mid-level vision mechanisms that emulate the processing that occurs in the primate visual cortex, and to design algorithms that apply such transformations with high computational efficiency. The candidate cortical mechanisms would be useful for edge detection, texture analysis, motion analysis, and image enhancement (i.e., de-convolution to eliminate blurring, contrast enhancement, and spatial frequency enhancement).

Two domains being explored are charting football plays and extracting choreography from a ballet sequence. These description schemes were demonstrated during Adelson's talk by means of videos in which real-life motion sequences were seen first compressed and then successfully uncompressed.

The 1995 Volen National Center for Complex Systems Scientific Retreat

Sponsored by the M.R. Bauer Foundation

The Center for Complex Systems: New Directions

MIT's Endicott House
Dedham, Massachusetts
Tuesday, May 30, 1995

On May 30, 1995, the Volen Center for Complex Systems held an all-day, off-campus retreat. The event was held at MIT's retreat center, Endicott House; a facility with lecture halls, beautiful grounds, and gardens.

The day was superb, with four stimulating lectures by Center junior faculty, posters by Center postdocs and graduate students, and gorgeous weather to accompany the events. A brief summary of each talk is on the following page.

The finale of the day was an after dinner talk by Professor Chris Miller, whose humorous and witty review of the history of the Center was outdone only by his musical talents, as he completed his comments with a song he had composed for the Center. The words to his composition are on page 12.

9:15 am
Registration and Coffee

9:45 am
Welcome
Irwin B. Levitan
Nancy Lurie Marks Professor of Developmental Neuroscience and Director, Volen National Center for Complex Systems

10:00 am
"Orientation Selectivity and Recurrent Excitation in the Visual Cortex"
Sacha Nelson
Assistant Professor of Biology and Volen National Center for Complex Systems

11:00 am
Coffee

11:15 am
"Modeling Adaptive Behavior and Learning—from Simple Agents to Humans"
Maja Mataric
Assistant Professor of Computer Science and Volen National Center for Complex Systems

12:15 pm
Lunch

2:00 pm
"Neurotrophins and the Control of Peripheral Neuronal Development"
Susan Birren
Assistant Professor of Biology and Volen National Center for Complex Systems

3:00 pm
Break

3:30 pm
"Temporal Coding in Human Memory"
Michael Kahana
Assistant Professor of Psychology and Volen National Center for Complex Systems

4:30 pm
Poster Session and Refreshments

6:00 pm
Dinner, followed by comments by Chris Miller
Professor of Biochemistry, Howard Hughes Medical Investigator, and Volen National Center for Complex Systems

Sacha Nelson, Ph.D.
Assistant Professor of Biology

Neurons in the visual cortex respond selectively to particular features or patterns in the visual world. The objective of our research is to understand the cellular and circuit-level properties that underlie this selectivity.

We employ two approaches to this problem: a "bottom-up" approach in which we try to understand the basic cellular and synaptic building blocks of cortical function in an *in vitro* slice preparation, and a "top-down" approach in which we study the responses of neurons to sensory stimuli *in vivo*.

This talk discussed our experiments concerning the way in which successive synaptic inputs are integrated over time, and the spatial aspects of synaptic integration.

Susan Birren, Ph.D.
Assistant Professor of Biology

The enormous complexity of nervous system function results from interactions between a vast number of different neuronal cell types. To begin addressing how this diversity is generated, we have been studying how sympathetic neurons develop from uncommitted embryonic precursor cells.

We have shown that the neurotrophic factors nerve growth factor (NGF) and neurotrophin-3 (NT-3) act sequentially on these developing neurons to support their survival during differentiation. Other factors such as fibroblast growth factor act as initial triggers of this differentiation process.

Taken together, these studies have demonstrated how precursor cells can develop under the influence of local factors (NT-3 and FGF), and eventually be supported by factors (NGF) produced by distant targets of sympathetic innervation. These developmental switches in factor responsiveness are likely to represent a general mechanism of nervous system development.

Maja Mataric, Ph.D.
Assistant Professor of Computer Science

This talk described two related research projects pursuing the understanding of complex interaction dynamics at two levels: macro and micro. The macro level is addressed by research into the synthesis and analysis of group behaviors with a collection of mobile robots.

The robots have been demonstrated to perform various cooperative group behaviors (including aggregation, dispersion, following, flocking, homing, and foraging) and group learning of foraging and social rules. The micro level is addressed by research into the mechanisms of learning through imitation, involving a collection of cognitive systems involved in perception, recognition, memory, and motor control.

Michael Kahana, Ph.D.
Assistant Professor of Psychology

Kahana presented his newly developed Temporal Coding Model of human memory. This model is designed to explain behavioral data on human memory for time of occurrence as well as long term recency effects. The basic idea is that summing the vectors representing prior memories can be mathematically modeled as a random walk in a high dimensional space. This random walk has mathematical properties that enable it to represent temporal information. This information is then concentrated with the vectors representing newly experienced memories and stored in an auto associative neural network. Kahana showed how this model overcomes a number of limitations in existing models of human memory.

Modern Complex Systems—
Dry and Wet

(To the score of *Modern Major General*)
Lyrics by Professor Chris Miller, Biochemistry
Volen Retreat, May 30, 1995

If you choose to do your work on systems with Complexity,
You'll surf the crest of fashion's wave as easily as x-y-z.
So many things in this wide world can oscillate chaotically,
You'll understand 'em well if you just treat 'em Mandelbroticaly.

You'll be called upon repeatedly to speak in words bombas-tic,
If your system seems random using models non-stochastic.
You'll double frequencies and bifurcate anytime, anywhere,
It worked 'cause the equations are egregiously nonlinear.

The subjects that we study are chosen to make us pensive. (Still,
We shy away from those that threaten to be comprehensible.)
Formal proofs of Lisman loops in memory and learning, yet,
Syntactics and semantics, nicely Gö-indeterminate.

There are neurons, glia, synapses, also primary astrocytes
And axons full of axoplasm and sensory dendrites,
And hippocampal bursters in which we store memory in bits,
And pattern generators in both you and Eve's invertebrates.

There are nerves of every size and shape in every kind of animal,
Cells horizontal, cells bipolar, even cells pyramidal,
And afferents, and efferents, and the stomatogastric ganglion, (long breath...)
And the spinal cord itself which all the other neurons dangle on.

From atmospheric chaos to the brain—hooray! It's all complex,
With properties emergent, neural nets catch butterfly effects,
From cunning linguistic distinctions Chomski-ish and Minski-esque,
To oscillating chemicals, reactions Zhabotinski-esque.

I'll give advice to you, my friends, your complex egos swollen,
Just keep working, but stay within the gorgeous walls of Volen.