

520
SCIENTIFIC
AMERICAN OFFPRINTS

T34

The Neurophysiology of Remembering

by Karl H. Pribram

SCIENTIFIC
AMERICAN
JANUARY 1969
VOL. 220, NO. 1 PP. 73-86



PUBLISHED BY W. H. FREEMAN AND COMPANY 660 MARKET STREET, SAN FRANCISCO, CALIFORNIA 94104

Copyright © 1968 by Scientific American, Inc. All rights reserved. Printed in the U.S.A. No part of this offprint may be reproduced by any mechanical, photographic or electronic process, or in the form of a phonographic recording, nor may it be stored in a retrieval system, transmitted or otherwise copied for public or private use without written permission of the publisher.

The Neurophysiology of Remembering

Experiments with monkeys have identified the brain areas involved in the recall of various learned tasks. Memory may take the form of interference patterns that resemble laser-produced holograms

by Karl H. Pribram

In 1950, toward the end of a busy life devoted to investigating the neurophysiology of memory, Karl S. Lashley wrote: "I sometimes feel, in reviewing the evidence on the localization of the memory trace, that the necessary conclusion is that learning just is not possible at all. Nevertheless, in spite of such evidence against it, learning does sometimes occur." That same year Edwin G. Boring, a leading psychologist of Lashley's generation, pointed out the deep impact that this failure to find physiological evidence for the memory trace had had on psychology. "Where or how," he asked, "does the brain store its memories? That is the great mystery. How can learning persist unreproduced, being affected by other learning while it waits? On the proper occasion what was learned reappears somewhat modified. Where was it in the meantime?... The physiology of memory has been so baffling a problem that most psychologists in facing it have gone positivistic, being content with hypothesized intervening variables or with empty correlations."

Hardly were these bleak observations in print before new research tools became available and were promptly applied in experiments on the neurophysiology of memory. As in all research that produces results important to workers in more than one discipline, however, dissemination across traditional boundaries is slowed by differences in vocabulary, in research technique and in the way a problem is subtly influenced by the subjects and materials employed by workers in different disciplines. As a result one finds even today that many psychologists (even those kindly disposed toward physiology) have the impression that little or no progress has been made in the effort to establish the neurophysiological basis of memory. This stems


from the fact that psychologists have addressed themselves primarily to questions about *process*, whereas neurophysiologists and neurochemists have addressed themselves primarily to the question of how the brain achieves short-term and long-term *storage*.

My own research has sought to answer more directly the questions posed by psychologists: What kinds of memory process must exist in the brain to allow remembering to take place? The results of this research have cast doubt on at least some of the assumptions about brain mechanisms (explicit and implicit) that are held by both psychologists and physiologists and that in my view have impeded any coming to grips with the problem of process.

Neurophysiologists had over several decades extensively mapped the brain with electrical recording devices and with weak electric currents to trace nerve pathways. As a result of such experiments on cats, monkeys and even men (performed during neurosurgery) physiologists could speak with some confidence of visual, auditory and somesthetic and motor areas in the cerebral cortex. Although they remained baffled by the "memory trace," they still felt they could describe the nerve pathways from a stimulus input (say the flash of a light) to a muscular response. The success of these studies often blinded the investigators to the fact that many of these presumed pathways could hardly be reconciled with Lashley's experiments dating back to the 1920's, which showed that rats could remember and could perform complex activities even after major nerve pathways in the brain had been cut and after as much as 90 percent of the primary visual cortex had been surgically removed.

As a neurosurgeon I had no reason to challenge the prevailing views of physiologists until I met Lashley and was convinced that we knew less than we thought. I soon resolved to continue his general line of investigation, working with monkeys rather than with rats, and in addition to make an effort to follow recordable changes of the electrical activity of the brain as the animals were trained to perform various tasks. Although this work has gone slowly at times (one experiment I shall describe took seven years), my co-workers and I have now gathered neurophysiological data from more than 950 monkeys. The results of these experiments are forcing many revisions in traditional concepts of how the brain works when tasks are learned and later remembered.

Beyond this I believe there is now available a hypothesis about the nature of the memory trace that satisfies the known physiological requirements and that can be tested by experiment. It is perhaps not surprising that the brain may exploit, among other things, the most sophisticated principle of information storage yet known: the principle of the hologram. In a hologram the information in a scene is recorded on a photographic plate in the form of a complex interference, or diffraction, pattern that appears meaningless. When the pattern is illuminated by coherent light, however, the original image is reconstructed. What makes the hologram unique as a storage device is that every element in the original image is distributed over the entire photographic plate. The hypothesis is attractive because remembering or recollecting literally implies a reconstructive process—the assembly of dismembered mnemonic events. In what follows, therefore, I shall give first the evidence for believing that




VISUAL-DISCRIMINATION TASK developed in the author's laboratory at Stanford University School of Medicine is depicted in this illustration. On the translucent panel in front of him the monkey sees either a circle or a series of vertical stripes, which have been projected from the rear. He is rewarded with a peanut, which drops into the receptacle at his left elbow, if he presses the right

half of the panel when he sees the circle or the left half when he sees the stripes. Electrodes record the wave forms that appear in the monkey's visual cortex as he develops skill at this task. Early in the experiments the wave forms show whether the monkey sees the circle or stripes. Eventually they reveal in advance which half of the panel the monkey will press (see illustration on page 6).

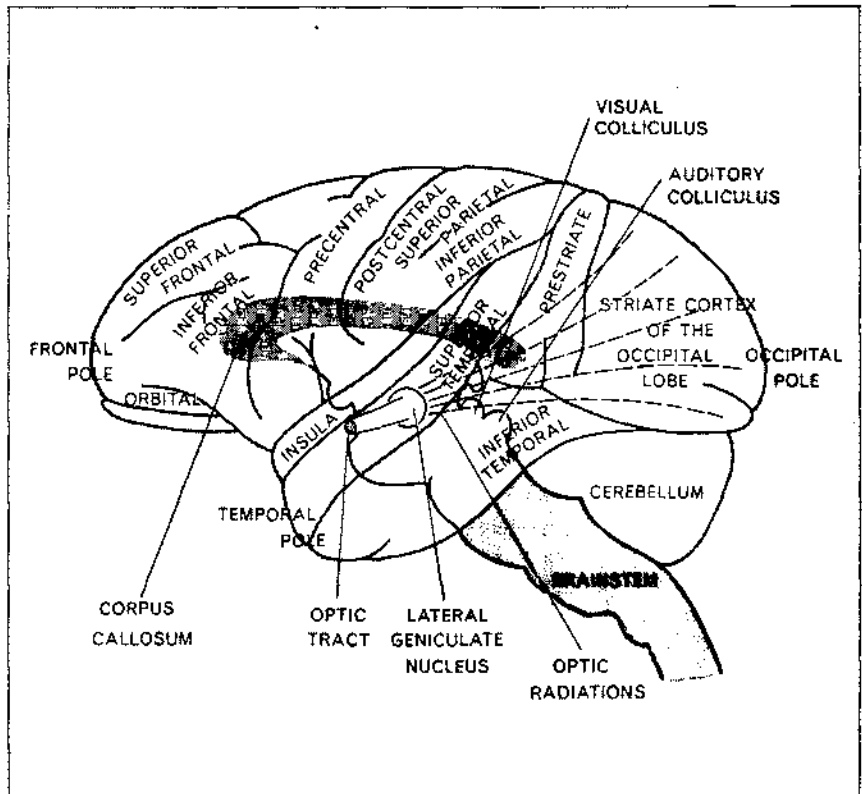
mnemonic events are distributed in the brain and then describe experiments that tell us something about the way these mnemonic events become re-collected into useful memory processes.

The abuses that the brain can survive and still function successfully have been documented many times since Lashley's pioneering experiments. Human testimony is provided daily in the neurological clinic of every large hospital when diseased or damaged brain tissue has to be removed. In the laboratory the brain seems to mock the ingenuity of the experimenter. Robert Calambos of the University of California at San Diego has severed up to 98 percent of the optic tract of cats without seriously impairing the cats' ability to perform skillfully on tests requiring them to differentiate between highly similar figures. Roger W. Sperry of the California Institute of Technology has surgically cross-hatched sensory receiving areas in the cortex of monkeys without disturbing the presumed organization of the input system. In other experiments the system continued to function even when Sperry inserted strips of mica in the cross-hatched troughs in an effort to electrically insulate small squares of tissue from one another. Conversely, Lashley, Kao Liang Chow and Josephine Semmes tried, without success, to short-circuit the electrical activity of the brain by placing strips of gold foil over the receiving areas. To accomplish a similar end I injected a minute amount of aluminum hydroxide cream at a number of points within a receiving area of an animal's cerebral cortex to produce electrical discharges resembling those seen in electroencephalograms during an epileptic seizure. Although these multiple discharging foci sharply retarded the animal's ability to learn a task of pattern discrimination, they did not interfere with recognition of these patterns when the multiple lesions were produced after learning.

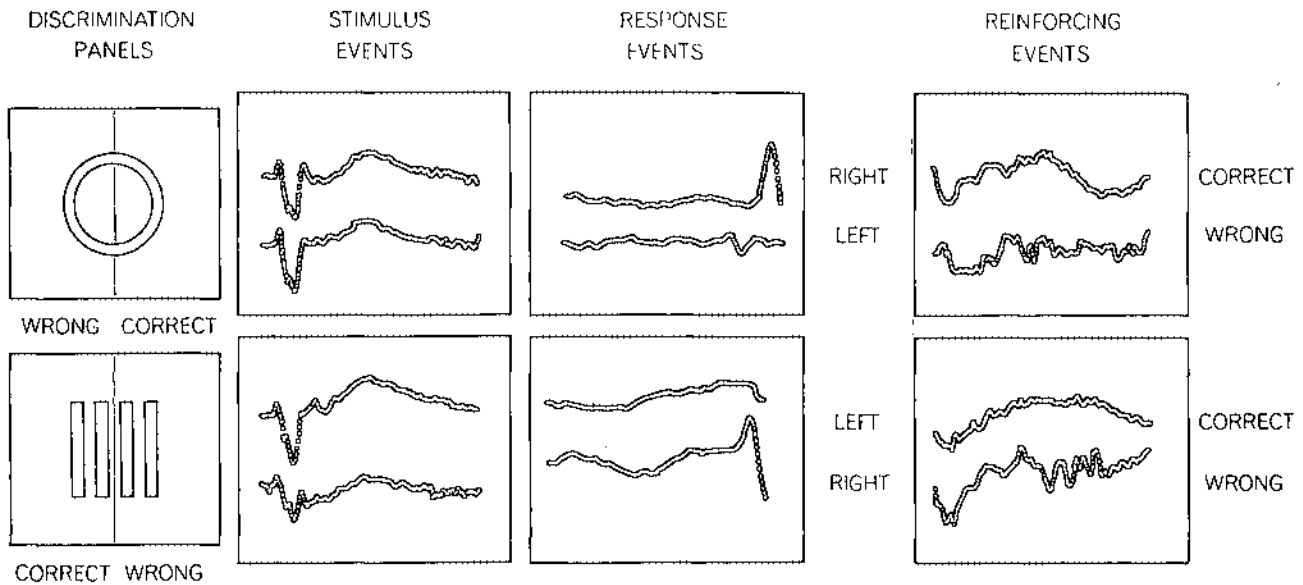
Such experiments have been interpreted as showing that each sensory system has considerable reserve capacity. Since it seems to make little difference in terms of performance which parts of the system are destroyed, it has been suggested that this reserve is distributed throughout the system, that the information needed to discriminate patterns is duplicated in many locations. According to this hypothesis, the discharging foci produced by injections of aluminum hydroxide cream interfere in some way with the reduplication that normally



LOCALIZATIONS OF FUNCTION in the cerebral cortex of monkeys have been known in general for many years. The evidence has been supplied in part by anatomical tracing of nerve pathways and more recently by electrical recording of wave forms, both through the intact skull and by use of implanted electrodes. Somesthesia refers to the sense of touch.



STRUCTURES IN CEREBRAL CORTEX AND BRAINSTEM mentioned in the text can be identified with the help of this illustration. Most of the cortical areas are labeled in adjectival form, the word "cortex" being omitted. The brainstem and its structures are shown in color. The corpus callosum is a bundle of nerve fibers that connects the two hemispheres of the brain. The lateral geniculate nucleus is the major relay station in the visual input system.



RESULTS OF VISUAL-DISCRIMINATION EXPERIMENT are shown in the wave forms recorded from the striate (visual) cortex of a monkey. The waves are those recorded after he has learned the task illustrated on page 4. The records under "Stimulus events" are wave forms that appear immediately after the monkey has been shown a circle or stripes. The records under "Response events" were generated just prior to the moment when the monkey actually responded by pressing either the left or the right half of the panel. The records under "Reinforcing events" were produced when the monkey was rewarded with a peanut if he was correct or not rewarded if he was wrong. The correct response was to press the right half of the panel on seeing a circle, the left half on seeing stripes.

A slight difference in the "stimulus" wave forms indicates whether the monkey has seen stripes or a circle. After he has learned his task well sharp differences appear in the response and reinforcing panels. The response wave forms, which are actually "intention" waves, show one pattern (the one with the sharp peak) whenever the monkey is about to press the right half of the panel, regardless of whether he has seen a circle or stripes. If he has actually seen stripes, of course, pressing the right half of the panel is the wrong response. Thus the wave forms reflect his intention to press a particular half of the panel. They could hardly reveal whether his response is going to be right or wrong because at this point he still "thinks" he is about to make the correct response.

takes place when information is being stored, but once storage is complete and the information is distributed all parts of the system are more or less "equipotent."

The correctness of this view has now been put to direct test. Over the past few years Nico Spinelli and I have shown that electrical activity recorded from widely distributed points in the striate, or visual, cortex of monkeys shows distinctive responses to different stimuli. Moreover, other widely distributed points within the cortex and brainstem give evidence that they have participated in storing information linked to the animal's response to particular stimuli. Let me describe the experiment more fully. (This is the one that took seven years to complete.)

Monkeys were placed in front of a translucent panel on which we could project either a circle or four vertical stripes [see illustration on page 4]. If, when the monkey saw the circle, he pressed the right half of the panel, he would be rewarded with a peanut. He would be similarly rewarded if he pressed the left side of the panel when the stripes appeared. Before the training begins we painlessly implant a num-

ber of tiny electrodes in the monkey's visual cortex. We then compare the electrical wave forms produced by the cortex during training with the wave forms produced after a high level of skill has been attained. We had expected that the wave forms would be different, and they were.

What we did not expect was that we would be able to tell from the waveform records whether the monkey saw a circle or vertical stripes, whether he responded correctly or made a mistake and, most surprising of all, whether he intended to press the right half or the left half of the panel once he was presented with the problem and before he initiated an overt response [see illustration above]. All these differing electrical responses arose in the visual cortex—the part of the brain that receives the visual input. We are forced to conclude that signals representing experience converge with and modify the input to the visual-input systems. We also found, however, that within the visual cortex different electrodes recorded different events.

Thus we now have direct evidence that signals become distributed within the input system. What we see (or at least what the monkey sees) is not a pure

and simple coding of the light patterns that are focused on the retina. Somewhere between the retina and the visual cortex the inflowing signals are modified to provide information that is already linked to a learned response, for example the monkey's intention to press one panel or another. Evidently what reaches the visual cortex is evoked by the external world but is hardly a direct or simple replica of it. Further, the information inherent in the input becomes distributed over wide regions of the visual cortex.

How might such a distribution of information occur? A possible clue to the puzzle came from an optical artifact, the hologram, which was then being made for the first time with the help of coherent laser light [see "Photography by Laser," by Emmett N. Leith and Juris Upatnieks; SCIENTIFIC AMERICAN, June, 1965]. The interference pattern of the hologram is created when a beam of coherent light is split so that a "reference" portion of the beam can interact with a portion reflected from a scene or an object. I reasoned (much as Lashley had) that neuronal events might interact in some way to produce complex patterns within the brain; the hologram now provided an explicit model.

Evidence for some such patterning of neuronal events, at least in the visual channels, has been provided by the work of R. W. Rodieck of the University of Sydney. He has shown that the initiating events in the visual channel that express the relations between the excitation of one receptor in the retina and the activity of neighboring points can be described mathematically through the use of "convolutional integrals," expressions somewhat similar to the familiar Fourier transformations. For example, the shape of the visual receptive field of a single retinal ganglion cell represents the convolution of a derivative of the shape of the retinal image produced at that point [see illustration on this page]. Convolutional integrals and Fourier transformations provide the mathematical basis on which holography was founded. Thus at least a first step has been taken to show that interference effects may operate in the central nervous system.

The question remains: How can interference effects be produced in the brain? One can imagine that when nerve impulses arrive at synapses (the junction between two nerve cells), they produce electrical events on the other side of the synapse that take the form of momentary standing wave fronts. Typically the junctions made by a nerve fiber number in the dozens, if not hundreds. The patterns set up by arriving nerve impulses presumably form a microstructure of wave forms that can interact with similar microstructures arising in overlapping junctional contacts. These other microstructures are derived from the spontaneous changes in electrical potential that ceaselessly occur in nerve tissue, and from other sources within the brain. Immediate cross-correlations result, and these can add in turn to produce new patterns of nerve impulses.

The hypothesis presented here is that the totality of this process has a more or less lasting effect on protein molecules and perhaps other macromolecules at the synaptic junctions and can serve as a neural hologram from which, given the appropriate input, an image can be reconstructed. The attractive feature of the hypothesis is that the information is distributed throughout the stored hologram and is thus resistant to insult. If even a small corner of a hologram is illuminated by the appropriate input, the entire original scene reappears. Moreover, holograms can be layered one on top of the other and yet be separately reconstructed.

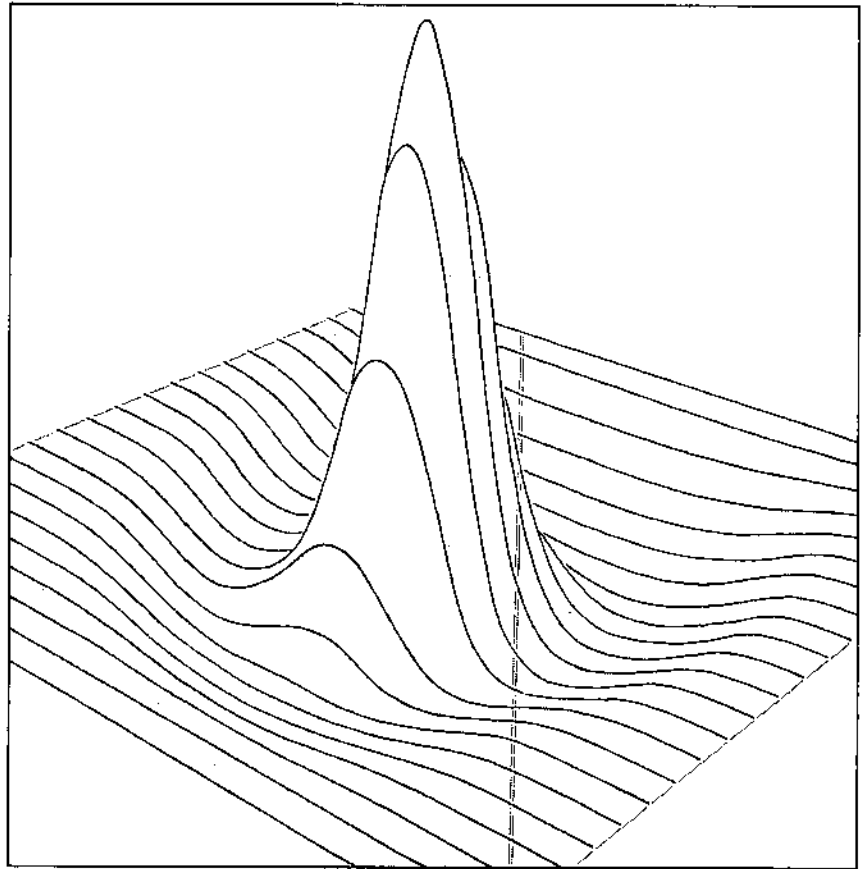
The holographic hypothesis imme-

diately raises many questions. Do the mathematical expressions that interpret the shape of visual receptive fields at the ganglion-cell layer of the retina yield equally useful interpretations at more central stations in the visual system? What kind of neural reference mechanism plays the role of the coherent light source needed to make and display holograms? Perhaps a kind of coherence results from the anatomical fact that the retina and the visual cortex are linked by many thousands of fibers arranged in parallel pathways. Or it could be that the nerve cells in the visual channel achieve coherence by rhythmic firing. Still another possibility is that coherence results from the operation of the variety of detectors that respond to such simple aspects of stimuli as the tilt of a line and movement that have recently received so much attention [see "The Visual Cor-


tex of the Brain," by David H. Hubel; SCIENTIFIC AMERICAN, November, 1963].

Other questions that flow from the holographic hypothesis are concerned with the storage of the memory trace. Two alternatives come to mind. The first involves a "tuning" of cell assemblies by changing synaptic characteristics so that a particular circuit will somehow resonate when it receives a familiar "note"; the second is some form of molecular storage, perhaps involving a change in structure at the synapses. Of course circuit-tuning may be secondary to just such structural changes, or the job may be done by a mechanism as yet unimagined. Such questions can be and are being investigated in the laboratory with techniques available today.


There is another line of investigation demonstrating that representations of



IDEALIZED MAP OF VISUAL RECEPTIVE FIELD represents recordings made from a single ganglion cell in the retina of the eye when a point source of light is presented in various parts of the visual field. The map contains smooth contour lines because the ganglion cell integrates the response of its neighbors, with which it is interconnected. The height of the contour at any point represents the number of times the individual nerve cell fires when the location of the point light source corresponds to that position on the map. Maximum firing occurs when the position of the light corresponds to that of the central peak. In mathematical terms, each contour line represents the "convolutional integral" of the first derivative of the shape of the stimulus figure. The interaction of many such convolutional integrals may produce hologram-like interference patterns within the visual system and elsewhere in the brain. Storage of such patterns could provide the basis of memory.




OLD VIEW OF VISUAL-RECOGNITION MECHANISM assumed that after visual information reached the striate cortex it was transferred to the prestriate cortex in two steps and from there to the inferior temporal cortex. Muscular response, according to the old view, then required that a message travel from the inferior temporal cortex to the precentral cortex (the motor cortex), which responded by sending signals down the brainstem to the muscles.



NEW VIEW OF VISUAL-RECOGNITION MECHANISM emphasizes the recent evidence that impulses from the inferior temporal cortex directly modify the visual input *before* it reaches the striate cortex (see illustration on page 10). This modification takes place sub-cortically through tracts leading to the visual colliculus and through interactions between that part of the brainstem and the lateral geniculate nucleus. There is also some evidence for an indirect pathway from the retina to the inferior temporal cortex. Visual information also seems to flow from the visual cortex to the visual colliculus. In the new view the body's muscle responses are relatively independent of the visual-recognition mechanism.

experience are distributed after entering the brain. The experiments I have described thus far demonstrated a distribution in space. There is also distribution in time; there are mechanisms in the brain for temporally distributing, or holding, events long enough so that they can be firmly registered. The evidence comes from an important group of experiments showing how animals (including man) gradually become habituated to a novel stimulus. Until recently habituation was thought to be due to a fatiguing of the nervous system. Eugene Sokolov of the University of Moscow showed, however, that when one is habituated, one can be dishabituated, that is, "oriented" anew, by a lowering in the intensity of the stimulus or even by complete silence when stimuli are expected. I like to call it the "Bowery-el phenomenon." For many years there was an elevated railway line (the "el") on Third Avenue in New York that made a fearful racket; when it was torn down, people who had been living in apartments along the line awakened periodically out of a sound sleep to call the police about some strange occurrence they could not properly define. Many such calls came at the times the trains had formerly rumbled past. The strange occurrences were of course the deafening silence that had replaced the expected noise.

In laboratory studies of this phenomenon the physiological concomitants of the orienting reaction are recorded and their reduction allows habituation to be investigated. The orienting reaction includes, among other things, changes in the conductivity of the skin (the galvanic skin response), changes in heart rate and respiratory rate, and changes in the electroencephalogram. Muriel H. Bagshaw and I found that we had to separate these physiological indicators of the orienting reaction into two classes. This was necessary because after we had surgically removed the frontal lobes of a monkey's brain, or the brainstem region known as the amygdala, the orienting stimulus no longer evoked the galvanic skin response or changes in heart rate and respiratory rate. (The responses themselves were not destroyed, because they could be evoked under other conditions.) On the other hand, surgery did not eliminate certain changes in the electroencephalogram and certain behavioral changes that also occur as a part of the orienting reaction. Surgery also interfered with habituation: a monkey lacking his frontal lobes or his amygdala continued much longer to show the behavioral and electroencephalographic orienting reactions. These results suggested that the loss of



INVESTIGATION OF VISUAL RECEPTIVE FIELDS is carried out by presenting a monkey with a small source of light that is systematically moved from point to point in a raster-like pattern. At each point the response of a single cell in the lateral geniculate nucleus is recorded by a microelectrode (a). During this mapping

a weak electrical stimulus can be delivered to other parts of the brain, such as the inferior temporal cortex (b) or the frontal cortex (c), to see if there is any effect. Some typical results are illustrated on the next page. The technique, which relies heavily on the computer, was developed by the author's colleague Nico Spinelli.

galvanic skin responses and heart and respiratory changes precluded habituation; when these indicators of orienting were not present, the stimulus, although perceived, failed to be registered in memory.

We have all had the experience of being preoccupied while a friend is recounting his experiences to us. Finally in exasperation he may say, "You aren't listening." Caught unaware, you may still be able quickly to repeat your friend's last sentence and from this even reconstruct what the "conversation" was about. If, however, your reverie is allowed to continue, much of what reached your ears will have been irretrievably lost; things just did not register. Thus there are two classes of indicators of orienting: the one concerned with just "sampling" the input, the other with its "registration," or storage.

E. D. Homskaya in the laboratory of A. R. Luria in Moscow and Mrs. Bagshaw in our laboratory at the Stanford University School of Medicine have also demonstrated that removal of the frontal lobes or the amygdala interferes with the indicators of registration when they appear in classical conditioning experiments. In normal animals the conditioning cue (such as a bell or a light) evokes changes in the galvanic skin response, in heart rate and in respiratory rate, as well as in the electroencephalogram. As the conditioning trials continue, these changes take place earlier and earlier until they actually precede the conditioning cue. It is as if the subject of the experiment were rehearsing the situation, anticipating what is coming next. After removal of the frontal lobes or the amyg-

dala, however, this rehearsal apparently ceases. Thus one can demonstrate that both anticipation and registration—a temporal distribution of mnemonic events—take place in a normal subject, and that these processes are impaired by surgery in certain parts of the brain. There is as yet little evidence to indicate how these parts of the brain bring about this temporal distribution of mnemonic events.

Given the fact that mnemonic events become distributed in the brain, what happens during remembering? Some kind of organizing process is clearly required. Experimental data make it likely that this process involves the "association cortex" of primates such as monkeys and man [see top illustration on page 5]. These regions are not to be confused with the "polysensory association cortex" that immediately surrounds the sensory projection areas and that has been studied so intensively in cats. The primate association areas consist of two general classes: the frontal and the posterior. The posterior association areas are located among the various primary sensory areas and consist of subareas that are specific for each of the senses.

In operations on several hundred monkeys my colleagues and I have made many kinds of lesion in this posterior system; the type, the size and the location of the lesions were based on a variety of anatomical and physiological criteria. These monkeys have been tested for their ability to learn and to retain discrimination tasks involving four senses. Vision is studied with a variety of patterns, colors and brightnesses; touch, with unseen objects of different shapes

and textures; taste, with samples differing in bitterness or sweetness; hearing, with different sound patterns. From the results of such experiments we are able to subdivide the posterior association cortex into areas, each serving a particular sense. These investigations show that the parieto-occipital area is concerned with touch, the anterior temporal cortex with taste, the middle temporal region with hearing. The inferior temporal cortex is important to visual discrimination [see illustrations on page 5].

These results present a number of questions. Why, following the removal of the inferior temporal cortex, do monkeys fail completely to accomplish visual discriminations while being perfectly able to accomplish discriminations in other senses and to perform more complex tasks, such as tasks involving delayed reactions and alternation of response? The problem is complicated by the following facts. Visual information passes from the retina to a relay point called the lateral geniculate nucleus and thence to the occipital (or striate) cortex. It had long been taught that the occipital cortex then sends information out to the surrounding areas, and that the information finally reaches the inferior temporal cortex. Since our monkeys fail in visual discrimination tasks after the removal of the inferior temporal cortex, the classical teaching would seem to be supported. Other considerations nonetheless argue that the classical view must be wrong.

First of all, the anatomical evidence shows that nerve impulses would have to be relayed by three synapses in traveling from the occipital cortex to the inferior temporal region. Three synapses,

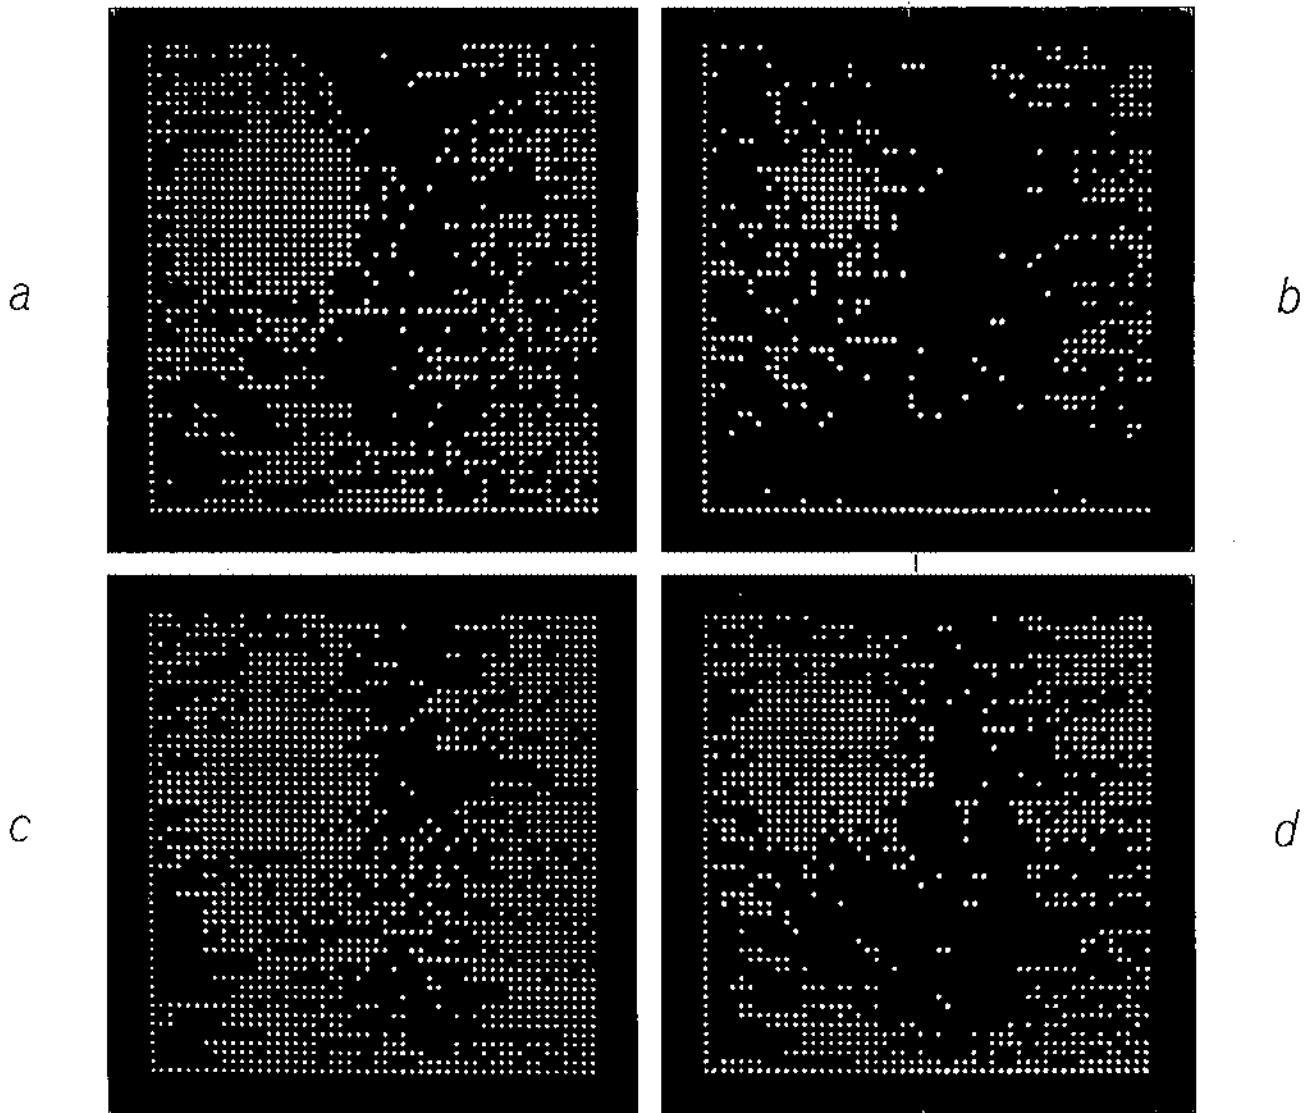
however, can get a signal from anywhere to almost anywhere else in the brain, so that this is hardly sufficient evidence for a mechanism that demands strict sensory specificity. Second, Chow, in a series of experiments confirmed by my own, removed all the tissue surrounding the occipital cortex in monkeys, so that the primary visual receiving area is totally isolated from the inferior temporal cortex. Such animals show no loss of visual performance in spite of the fact that a lesion in the inferior temporal cortex only a third or a quarter the size of the one made in the disconnection experiment will cause serious impairment on the same tasks. This makes it most unlikely that impulses reaching the inferior temporal cortex from the upper regions of the visual system account for the importance of this cortex in vision.

What, then, is the mechanism that

enables the inferior temporal cortex to play such a key role in the performance of visual tasks? Where does it get its information and where does it send it? The available evidence (much of which I have had to omit in this brief account) has led me to propose that the inferior temporal cortex exerts its control by organizing the traffic in the primary visual system. Recently the pathways from the inferior temporal cortex to the visual system have been traced. Applying the methods of electrophysiology, Spinelli and I have found, for example, that we can change the size and shape of visual receptive fields by stimulating the inferior temporal cortex [*see illustration below*]. These and other experiments demonstrate beyond doubt that the inferior temporal cortex is not the passive recipient of data relayed from the primary visual cortex, as was long believed,

but actively influences what enters the visual cortex. Similar results have been obtained in the auditory system by James H. Dewson in my laboratory.

An experiment that tells us a little about the meaning of this control over input is currently being completed by my associate Lauren Gerbrandt. A monkey sits in a chair inside a box that can be opened, so that he can see out, or closed. He can be stimulated through an electrode placed in the lateral geniculate nucleus (the relay station in the visual input system) while we record the level of activity in the visual cortex. When the box is closed, geniculate stimulation evokes only a small response in the cortex. When the box is open, the response is large. Gerbrandt found, however, that he could augment the strength of the cortical response when the box is closed (and only then) by stimulation of



VISUAL-RECEPTIVE-FIELD MAPS, made by the technique illustrated on the preceding page, show how information flowing through the primary visual pathway is altered by stimulation elsewhere in the brain. Map a is the normal response of a cell in the


geniculate nucleus when a light source is moved through a raster-like pattern. Map b shows how the field is contracted by stimulation of the inferior temporal cortex. Map c shows the expansion produced by stimulation of the frontal cortex. Map d is a final control.

the inferior temporal cortex. The response is then as strong as when the monkey is alertly looking about, examining the world around him. This suggests to us that electrical stimulation of the association cortex crudely reproduces the neural activity that goes on naturally when the animal is actively engaged in sampling and attending his visual environment.

A detailed and satisfactory mechanism for explaining these results remains to be worked out. A tentative hypothesis supported by considerable anatomical evidence, and very recently by limited electrophysiological evidence, might go something like this. There is evidently an input from the visual pathway, rather separate from the primary visual pathway, that leads to the inferior temporal cortex. This visual input to the inferior temporal cortex triggers a process that feeds back into the primary visual system and there exercises a control over the flow of visual impulses to the visual cortex [see bottom illustration on page 8]. This view is based on such evidence as our ability to change the size and shape of the visual fields in the optic nerve and lateral geniculate nucleus by stimulation of the inferior temporal cortex.

This, however, can be only a part of the story. A satisfactory hypothesis also has to explain the first experiment I described, in which recordings from the visual cortex foreshadowed the monkey's intention to press either the right or the left panel when he was presented with a circle or vertical stripes. Here we have evidence that the frontal cortex and the amygdala, which are involved in registration, also affect the visual mechanism, often in a direction just opposite to what is produced by stimulation of the inferior temporal cortex.

Pathways from the visual cortex to the superior colliculus of the brainstem are well known. Recently we have traced similar pathways from the inferior temporal cortex to this same superior colliculus, which is an important structure in the visual system. (In birds the collicular region plays a role comparable to the role of the cerebral cortex in primates.) One can now begin to see how surgically isolating the visual cortex from the inferior temporal cortex does not destroy an animal's capacity to perform visual tasks. Evidently the communication link between the visual cortex and the inferior temporal cortex (which is essential to the retention of visual discriminations) is buried deep within the brainstem. Just as the brainstem serves



INTERSECTION OF NERVE PATHWAYS in the visual-input channel can be depicted schematically in two dimensions by an array of dots, each representing a single nerve cell. The response of each cell in turn can be visualized as corresponding to the patterns shown in the visual-receptive-field maps presented on page 10. At a given instant a stimulus arising in a particular part of the visual field will cause a certain group of cells (*color*) to respond. Simultaneously a stimulus in another part of the field will excite a different group (*black*). Gray cells are inactive at this instant. As long as the scene in the visual field remains constant, these same groups of cells will "flash" off and on many times a second. The interference patterns resulting from the interacting fields of the flashing cells may provide the opportunity for the formation of holographic patterns. This diagram first appeared in "The Physiology of Imagination," by John C. Eccles; SCIENTIFIC AMERICAN, September, 1958.

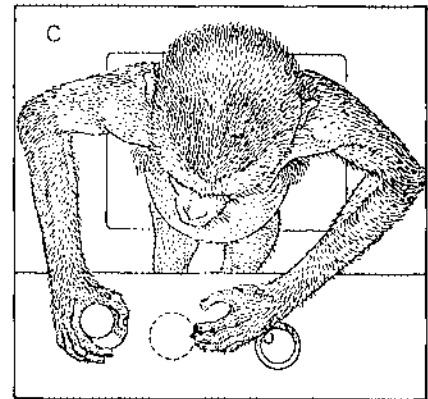


as a convergence station for the visual system, it serves (on the basis of Dewson's evidence) a similar function in hearing. The importance of such sub-cortical convergences, which in turn alter the input to the cortex, has been highlighted by these experiments.

Further evidence for this cortico-sub-cortical mode of operation of the brain (as opposed to a transcortical mode) comes from the same group of experiments in which our animals learned to distinguish between circles and stripes while a record of their brain waves was being made. The reason for doing these experiments in the first place was that I wanted to see how the wave forms recorded from various parts of the brain would be altered by making lesions in the inferior temporal cortex after the monkeys had learned their task. I fully expected that a lesion would selectively affect one of the wave forms and would leave others unchanged. Thus (I hoped) we would be able to identify the mechanism that accounted for the monkey's failure to perform satisfactorily after the lesion. We might conclude, for example, that the lesion had interfered with the monkey's capacity to differentiate between circles and stripes or that it had

interfered with some process linked to reinforcement or response. This is not what happened.

Instead of finding a selective change in one or another of these electrical waves, we found that the electrodes that provided the best differential recordings in advance of surgery subsequently showed no such differences; other electrodes whose wave forms had been undifferentiated now showed persistent and reliable differences. These differences turned out to be associated, for the most part, with responses, but in very peculiar ways that we have not as yet been able to decipher clearly.

It seems as if the frame of reference within which the brain activity had been working before the lesions were made was now shifted, and in fact was shifting from time to time. Judging by their behavior, the monkeys were as surprised by the effects of the surgery as we were. They approached their task confident of their ability to solve the problems, only to find they made errors (and hence received no peanut). This resulted in spurts, hesitations and variability in performance. It seemed as if they were completely baffled, not realizing, of course, that it was the inside of their



ALTERNATION TASK requires that a monkey remember which cup he lifted last in order to lift the correct one on his next trial and be rewarded with a peanut. Normally he is rewarded if he remembers to lift the cups in a simple alternating sequence: left, right, left and so on. After each trial a screen comes between him and the cups and remains there for periods that can be varied from

seconds to many minutes. In part c of this sequence the monkey has forgotten to alternate his response. Experiments demonstrate that certain brain lesions interfere with a monkey's ability to remember what he did even a few seconds earlier. By changing the task only slightly, however, the author found that brain-damaged monkeys were no longer perplexed (see illustration on page 13).

heads and not the situation that we had changed. What is the explanation?

Whatever the transformations of input (holographic or otherwise) that occur in the nervous system, such transformations are in effect coding operations. In order for a code to work, that is, to be decipherable, it must be framed within a context. This context must remain stable or the information conveyed by the code will be destroyed by successive transformations. Our reading of the recordings of the electrical brain activity of the monkeys who had their inferior temporal cortex removed is that the framework within which their discriminations had been made before surgery was now disrupted and shifting. The events observed by these monkeys no longer conveyed information because their brains had in a sense become unstable.

As a hypothesis this can be tested. We are about to investigate means of providing externally the stability that the brains of the brain-damaged monkeys evidently lack. Sandra Reitz, a student in my laboratory, recently suggested that this could be done simply by increasing the spatial redundancy of the visual cues (that is, the number of identical displays) that we present to our monkeys on a discrimination task. The expectation is, if our view is correct, that this change in the task will overcome the difficulty in discrimination experienced by the monkeys with lesions of the inferior temporal cortex.

Therefore a beginning has been made in specifying the structures that participate in the organization of memory inside the brain. The next task is to discover *how* these structures accomplish the physiological processes we call remembering, whether by holographic

representations or by some process even more subtle. In our concern with the storage mechanism, however, we should never overlook that aspect of memory which is of overriding importance to the process of effective remembering: the method of organizing or coding what is to be remembered.

In everyday life there are many homely examples to show that a given message is easier to remember in one form than in another. For example, rhymes are often employed in aphorisms ("A stitch in time saves nine"); many people cannot remember the number of days in the month without first recalling the jingle of their childhood. A more important example of the value of efficient coding is found in the 0-9 method of writing down numbers compared with the clumsy system of the Romans. By employing the concept of zero to indicate multiples of 10 our mathematical tasks are vastly simplified.

A coding mechanism need not necessarily be very complicated. Take, for example, the following "poem," which the neurophysiologist Warren McCulloch likes to intone with bishop-like solemnity: INMUDEELSARE/INCLAYNONEARE/INPINETARIS/INOAKNONEIS. When spaces are inserted where they belong, the message instantly becomes clear: IN MUD EELS ARE/IN CLAY NONE ARE/IN PINE TAR IS/IN OAK NONE IS. The passage has been decoded by the simple procedure of parsing, or what the psychologist George A. Miller, my sometime collaborator who is now at Rockefeller University, calls "chunking" [see "Information and Memory," by George A. Miller; SCIENTIFIC AMERICAN, August, 1956].

Many experiments with monkeys demonstrate that the frontal cortex—long regarded as the site of the "highest

mental faculties" in man and primates—plays an important role in short-term memory, whatever else it may be doing. When sufficiently complex tests, comparable to those used with monkeys, are given to lobotomized patients, they too show this memory disturbance. My experiments provide strong evidence that the primate frontal cortex performs its role by means of a coding operation that seems to resemble parsing, or chunking. When the frontal cortex of a monkey is damaged, the animal has difficulty performing tasks in which he has to remember what happened just a few seconds earlier.

Typical of such tasks is one in which the monkey faces two identical cups with lids that he must raise in a particular sequence to obtain a peanut [see illustration above]. In the simplest case he is rewarded with a peanut at each trial if he simply remembers to lift the lids alternately: left, right, left and so on. Then he must wait a specified interval, which can be varied from a few seconds to hours, between each trial, and while he is waiting an opaque screen is interposed between him and the cups. His task, then, is to remember which lid he lifted last so that he can lift the other one on the next trial. A monkey whose frontal cortex has been resected will fail at this simple task even when the interval between trials is reduced to three seconds.

It occurred to me that perhaps the task appears to these monkeys much as an unparsed passage does to us. I therefore changed the task so that the rewarded sequence became left-right (long interval), left-right (long interval) and so on. There was still a mandatory pause with the screen interposed of five seconds between each left-right trial, but


now a longer interval of 15 seconds was inserted between *pairs* of trials. Immediately the monkeys with frontal cortex damage performed as successfully as the control animals whose brains were intact [see illustration below]. That time-parsing was the key to the success of the brain-damaged monkeys was shown by other experiments in which the interval between trials was held constant but some other clue, such as a red light or a buzzer, was presented at every other trial. The clues were ignored; the monkeys with frontal lobe resections still failed at the task.

The experiment is important in several respects. First, it demonstrates at least one function of the frontal lobes, a func-

tion that may be basic to other functions. Second, it suggests that the difficulty the brain-damaged monkey has in recalling what he did last is not due simply to a premature fading of the memory trace; after all, he improved quickly when a longer interval was interposed, provided that the task was adequately structured. Third, this structuring, organizing or coding is in fact crucial to the process of recall.

Other studies show that the frontal cortex, like the posterior association cortex, exercises control over sensory information flowing into the cortical receiving areas. In many instances, as I have noted, electrical stimulation of the

frontal cortex produces effects that are opposite to those produced by posterior stimulation. Our studies are not advanced enough as yet to specify which pathways from the frontal lobes may be involved. Recent work done by Donald B. Lindsley and Carmine D. Clemente at the Brain Research Institute of the University of California at Los Angeles indicates that the pathway involved may be a large tract of fibers (running in the medial forebrain bundle) that carries inhibitory impulses to the reticular formation of the brainstem. I have on occasion attempted to spell out some possible relations between neural inhibitory processes and short-term memory but such efforts are at best tentative.



MODIFIED ALTERNATION TASK could be mastered as readily by monkeys with part of their frontal cortex removed (*colored curve*) as by normal monkeys (*black curve*). The brain-damaged monkeys had been unable to solve the standard left-right alternation task (described in the illustration on page 12) even when the interval between trials was only a few seconds. The task was then modified so that the interval between each left-and-right trial was

kept brief (five seconds) but a 15-second pause was inserted after every right-hand trial. When this change was made, brain-damaged monkeys performed about as well as normal monkeys, as shown here. Errors are the number made each day before a monkey achieved 40 successful trials. Bars indicate the range of errors made by different monkeys. Data for the 15th day show the result when all the trials were again separated by equal intervals of five seconds.

Coding and recoding are thus found to be essential operations in both memory storage and remembering. I have described evidence showing clearly that storage is distributed throughout a sensory system. I have also mentioned some evidence suggesting that the transformations (coding operations) that are performed within the input channels can be described in terms of convolutional integrals. The basic premise involved is that neighboring neural elements do not work independently of one another. By virtue of lateral interactions, neural elements spatially superpose the excitatory and inhibitory electrical potentials that arise among neighboring nerve cells. These transformations generate a microstructure of postsynaptic events, which

can be regarded as wave fronts that set up interference patterns with other (pre-existing or internally generated) wave fronts, producing in their totality something resembling a hologram. Given a mechanism capable of storing this hologram, an image could be evoked at some later time by the appropriate input. In order to be effective as codes, transformations must take place within some stable framework. To an extent this framework can be provided by the stored microstructure itself, by the parallel pathways of the input system, by the specific detector sensitivities of units in the system and by the very redundancy of the external environment. (We have no trouble recognizing automobiles because there are so many of them and

they are so much alike.)

For complex and novel events, however, a more powerful organizer must come into action. Experiments conducted in my laboratory and elsewhere suggest that this organizing mechanism critically involves the association areas of the cerebral cortex. The mechanism does not, however, seem to reside within these areas. Rather, the association areas exercise control on the input system by way of deeper structures in the brainstem. In short, the function of the association areas of the cortex turns out to be that of providing a major part of the organizing process necessary to remembering: the reconstruction of an image from distributed mnemonic events.

The Author

KARL H. PRIBRAM is research professor of psychology and psychiatry at Stanford University. He began his studies of the brain as a neurosurgeon. The advent of psychosurgery caused him to move into laboratory studies of monkeys; the premise, he writes, is that "one should know what one is about before mucking around with the human brain." The work has directed his career toward neuropsychology and experimental psychology, to the point where he was recently elected president of the Division of Physiological and Comparative Psychology of the American Psychological Association. Pribram's hobby is photography, and when holography became a practical reality, he perceived its relevance to problems of brain function that had been baffling for many years. As a result the image-forming properties of the brain have become a major focus of his current work.

Bibliography

- THE PROBLEM OF CEREBRAL ORGANIZATION IN VISION. K. S. Lashley in *Biological Symposia, a Series of Volumes Devoted to Current Symposia in the Field of Biology: Vol. VII, Visual Mechanisms*, edited by Jaques Cattell. The Jaques Cattell Press, 1942.
- IN SEARCH OF THE ENGRAM. K. S. Lashley in *Physiological Mechanisms in Animal Behavior: Symposia of the Society for Experimental Biology, No. 4*. Academic Press, 1950.
- TOWARD A SCIENCE OF NEUROPSYCHOLOGY (METHOD AND DATA). Karl H. Pribram in *Current Trends in Psychology and the Behavioral Sciences*. University of Pittsburgh Press, 1954.
- THE PHYSIOLOGY OF IMAGINATION. John C. Eccles in *Scientific American*, Vol. 199, No. 3, pages 135-146; September, 1958.
- SOME DIMENSIONS OF REMEMBERING: STEPS TOWARD A NEUROPSYCHOLOGICAL MODEL OF MEMORY. Karl H. Pribram in *Macromolecules and Behavior*, edited by John Gaito. Appleton-Century-Crofts, 1966.