

The Cognitive Neuroscience Approach

Kevin N. Ochsner
Stephen M. Kosslyn

Imagine that you are a physician trying to understand how digestion is accomplished; you want to know how the structure and function of internal organs cause food input to be converted into energy output. How might you begin and what types of information might you seek? You might begin by observing the problems experienced by people who have suffered injuries to particular internal organs, which could give you a general idea of what functions each helps to carry out in digestion. However, you might soon notice that damage to different organs results in similar, but not identical difficulties. Furthermore, you find that damage to a given organ does not always result in a single, consistent dysfunctional pattern, and the particular observed pattern seems to depend upon the integrity of related organs. Understanding digestion becomes an increasingly complex task as each new observation reveals a host of new questions. Moreover, each organ may participate in more than one aspect of digestion, and observations of deficits and locations of damage are not enough to identify the functions uniquely.

Happily, you can use information from other domains to help you interpret the data. Anatomical studies reveal that the organs in question, such as the stomach, intestines, liver and pancreas, are not connected haphazardly; in fact, they appear to be connected in a precise fashion, each organ connected to a specific set of other organs. These data help you piece together the notion of a “digestive tract” or system in which each organ carries out a particular process upon food as it passes through the system. This insight might help explain why the pattern of deficits fol-

lowing damage to one organ depends upon the integrity of others. You realize further that in order to understand the functional relationships among these organs you need to know more about how they communicate, what they do, and which aspects of their operation might be shared by all organs in the digestive tract. For example, you learn that a chemical called insulin is produced by the pancreas, is found in the liver, and has been linked to changes in blood sugar levels in rats. How does insulin relate to different aspects of digestion? How and when does the liver use it? Might insulin be influenced by, or itself influence, other processes? In answering these questions you may develop an understanding of digestion that draws on many levels of analysis, from the molecule to the system, ultimately providing a coherent description of the digestive process in general, and the role of specific components of the process, like insulin, in particular.

From this perspective, the digestive tract and the brain may have a lot in common. The goal of cognitive neuroscience is to understand how brain function gives rise to mental abilities such as memory, reasoning, vision, or movement, and to understand how such abilities interact with the systems underlying emotion. Our analogy illustrates one way in which multidisciplinary research on such a complex problem might proceed. Unfortunately, the task of understanding the relation of behavior to the structure and function of the brain is much more complex than that of understanding the relation of digestion to the structure and function of the liver, stomach, and other organs in the digestive tract. Not only is the brain structurally more complex than the digestive tract, it also carries out many more functions, and each function is both more complex and more difficult to describe operationally than are aspects of the digestive process.

Nevertheless, the same multidisciplinary approach that has proven successful in medical research has been incorporated in cognitive neuroscience. In cognitive neuroscience, we consider data collected by researchers studying behavior, cognition, neurophysiology, neuroanatomy, and computation, and each new finding provides additional fodder for theories of brain function. Theory building thus becomes a process of trying to fit together a wide variety of different types of information into a complex, but integrated whole. Thus the different types of information must be explained by the same theory; the theory is not simply “constrained” by different types of data, but rather it is an attempt to explain diverse phenomena with a single set of assumptions and principles.

Perhaps the most significant difference between cognitive neuroscience and cognitive science is that cognitive neuroscience aims to understand the neural implementation of mental abilities. Cognitive science focuses only on function, independently of the organ that gives rise to the function. From a cognitive science perspective, there are many ways that digestion could in principle take place, and the goal of research is seen as an attempt to specify them. But even though studying the function in its own right may implicate only a few candidate digestive processes, there is only one way that the body actually converts food input into energy output. If our goal is to understand how digestion works, viewing it as a functional

property of the physiology and anatomy of the digestive tract focuses and delimits the problem. Similarly, research in cognitive neuroscience characterizes function as a property of the brain itself, and in so doing necessarily integrates across physical and functional levels of analysis. Given that the human brain is currently the only system capable of producing the complex functions we call memory, emotion, and so forth, it makes sense to try to understand how its design is related to these abilities.

The purpose of this chapter is to illustrate the cognitive neuroscience approach in several problem domains. We selected domains that pertain to visual analyses of the world: attention, high-level vision, and visual memory. We focus on vision in large part because much of the research on the neural bases of cognition and behavior, especially using animals, has focused on vision. Our focus on vision also allows us to illustrate a key feature of the cognitive neuroscience approach: A major goal of cognitive neuroscience is the construction of integrated theories of cognition and behavior, and we wanted to illustrate the process of theory construction not simply within a single domain, such as attention, but across several related domains. Vision is thus the tie that binds several domains together.

Nevertheless, the reader should be aware that the cognitive neuroscientific approach outlined here is being profitably applied to the study of many topics, such as language (Caplan, 1993), movement (Georgopoulos, 1994), and emotion (J. E. LeDoux, 1994). Some of the topics we touch upon are reviewed in greater depth, from a different perspective, elsewhere in this volume (e.g., see chapters in this volume by Martin, 8, on language, LaBerge, 2, on attention, Gallistel, 1, on action).

With these considerations in mind, we discuss first key historical developments and general principles that have motivated research in cognitive neuroscience, and then consider briefly how research on attention, high-level vision, and visual memory has benefited from an application of the cognitive neuroscience perspective.

I. FOUNDATIONS OF COGNITIVE NEUROSCIENCE

A number of important advances in neurobiology, experimental psychology, and computer science laid the foundations for the emergence of cognitive neuroscience (see Kosslyn & Andersen, 1993).

In the late 1960s researchers began recording the electrical activity of cortical neurons in awake, behaving monkeys (Evarts, 1966). This technique allowed precise correlation of behavioral and physiological data, and researchers were able to characterize the functional organization of some parts of the brain. For example, Hubel and Wiesel (1968) discovered that the primary visual cortex is composed of a series of columns of cells, and the cells in each column are sensitive to the presence of bars or edges with particular orientations located in specific parts of the visual field. Although the parcellation of the brain into discrete functional components had been suggested by earlier work with brain-damaged patients and animals (e.g., Broca, 1863), direct measurement of neural activity allowed researchers sys-

tematically to map stimulus parameters onto the function of individual or groups of neurons.

A similar emphasis on understanding complex functions in terms of constituent processes also emerged in the cognitive psychology of the late 1960s and 1970s (e.g., see Neisser, 1967), but this approach had a very different origin: the computer metaphor. Researchers in psychology began to conceive of internal processing in humans by analogy to internal processing in a computer. For example, Sternberg (1969) developed a technique for isolating distinct information-processing stages that were characterized in terms of how information is stored, encoded, interpreted, or compared. Similarly, Posner developed tasks to tap simple component processes of complex abilities such as attention (e.g., Posner, Nissen, & Ogden, 1978). In addition, Shepard and his colleagues (e.g., Shepard & Feng, 1972; Shepard & Metzler, 1971) provided evidence that the brain can perform analog computations in some situations. Together, these findings led researchers to conceptualize behavior as arising when specific types of processing are performed upon specific internal representations.

However, cognitive psychologists recognized that descriptions of behavior are not enough to implicate one set of underlying mechanisms. Indeed, it was proved that any set of behavioral data could be explained by a number of theories (e.g., Anderson, 1978). Anderson argued that data from neurophysiology, such as that collected by Hubel and Wiesel, could provide critical insights into the nature of internal representations and the processes that operate upon them, thereby limiting our choices among theories to those that could accommodate these data.

Further key developments came from the rapidly developing field of artificial intelligence within computer science. Von Neumann (1958) and McCulloch and Pitts (1943) suggested that neural processes could be usefully conceptualized as computational processes. Early computational models demonstrated that neural activity could in fact be conceptualized as information processing. Combined with Hebb's (1949) associative model for learning in networks of neurons, and new findings of circumscribed learning deficits following focal brain damage (e.g., Scoville & Milner, 1957), there seemed good reason to believe that mental abilities could be viewed in terms of discrete processing stages operating upon internal representations.

Thus, links between mental abilities and sets of distinct processes were becoming apparent, and links between brain function and computation were becoming clear. But the whole was not greater than the sum of its parts until researchers saw how to combine the different sorts of information so that they mutually informed each other.

A. The Cognitive Neuroscience Approach

A comprehensive framework for understanding how the brain carries out computations was developed by David Marr (1982). Marr's work focused on vision, but his approach can be generalized to any type of biological information processing.

He posited that vision should be studied at three levels of analysis, which in turn must be integrated. These levels varied in abstraction: At the most abstract level, a theory specifies *what* is computed by a specific module; at an intermediate level, a theory specifies *how* a given computation is actually carried out (i.e., it specifies an algorithm); and at the most concrete level, a theory specifies how a set of processes is actually implemented in the brain. Marr argued that the three levels, particularly the more abstract ones, could be studied independently. This perspective is clearly compatible with the notion that the mind is like a computer program, which can be understood independently of the machine on which it runs. This view has recently been questioned, however, and many researchers are now impressed more by the close relationships among the levels than by their independence (e.g., Kosslyn & Maljkovic, 1990).

Indeed, the dominant paradigm in experimental psychology appears to be shifting because of two factors. First, many researchers have been impressed by the power of the connectionist (“neural network”) method of modeling cognitive abilities; such models conflate Marr’s levels of analysis (see Grossberg, 1980; Kosslyn & Koenig, 1992; J. L. McClelland & Rumelhart, 1986). Second, research in neurobiology has revealed a close relationship between the structure and function of the brain (for a review, see Kosslyn & Koenig, 1992). The brain is not a general purpose machine that can be programmed in any arbitrary way; rather, key aspects of the structure of the brain apparently have been tailored (via natural selection) for the specific types of computations that it performs. Cognitive neuroscience has emerged in part because researchers realized that facts about the evolution and biology of the brain could provide insight into the nature of cognition.

Research in neuroscience has led to several generalizations that in turn have guided much theorizing in the field. These generalizations can be summarized as follows.

1. *Limitations on “optimal” performance.* Brains have limited processing capacities because they are part of a biological system. Therefore, there can be no “optimal” or “logically correct” solution to a computational problem without reference to available hardware and resources; each computational step requires metabolic energy and must interact with the resource requirements of other processes. In addition, the brain was not engineered to perform optimally all computations; rather it is the product of hundreds of thousands of years of selection pressures that have added particular functions to those already present (Sherry & Schacter, 1987) if such functions enhanced the reproductive capability of the organism (but also sometimes even if they did not; see Gould & Lewontin, 1979). Thus any theory of the computation, algorithm, or implementation that does not take into account these limitations may make unfounded assumptions about what is possible, and therefore risks biological implausibility.

2. *Anatomical structure.* As will be discussed in more detail below, the brain is not a homogeneous “wonder net”; rather, different parts do different things. Moreover,

anatomical connectivity leads some sets of processes to take place in parallel, and other sets to take place in series (e.g., see DeYoe & Van Essen, 1988; Ungerleider & Mishkin, 1982). In addition, information typically does not flow in only one direction in the brain. In the vast majority of cases, every projection from a lower (i.e., closer to sensory input) area to a higher (i.e., further from sensory input) cortical area is accompanied by connections running in the opposite direction—and the two kinds of connections are of comparable size (Felleman & Van Essen, 1991). Furthermore, it appears that these reciprocal, feedback, connections have more diffuse target regions than the feed-forward connections; this anatomical fact may suggest that a given process can be modulated by many others. Thus it may not make sense to consider a single computation or algorithm in isolation because computations are carried out by systems of interacting subsystems (Kosslyn & Koenig, 1992; Posner & Petersen, 1990; Schacter, 1994). Connectionist (neural network) models are useful in part because they can be constructed to mimic the interactive nature of cortical processing; they consist of interacting layers of neuron-like nodes that can be designed to involve extensive bidirectional cross-talk between input and output levels.

3. *Physiological observations.* Basic facts about neural dynamics also shape the way the brain can process information. For example, the brain can carry out only about 10 serial steps to produce a response 250 ms after a stimulus has appeared (Churchland & Sejnowski, 1992; Feldman, 1985). Considering data from neuropsychology, neurophysiology, and other branches of neuroscience not only helps us understand existing data and evaluate theories of cognition, but also helps develop new theories and collect various types of additional data.

As illustrated in Figure 1, the cognitive neuroscience approach can be represented as an equilateral triangle with *abilities* at the apex, and *neuroscience* and *computation* at the two bottom corners (see also Kosslyn, 1994). Abilities is at the top because that is what one is trying, ultimately, to explain, and neuroscience and computation are at the bottom because the explanations rest on conceptions of how the brain computes. The equal length of the connections between each vertex reflects the fact that there is no privileged level of analysis or means of constraining or generating hypotheses. Explanations derived from multidisciplinary analyses necessarily turn on a confluence of facts about abilities (usually as manifested in observable behavior), the brain, and computation.

Many theories in cognitive neuroscience aim to specify the functional architecture for a specific type of processing. Such theories have two components: First, they may specify a set of processing subsystems, which either store or transform information in some way, and how information flows from one component of the system to another (e.g., see Kosslyn, 1994). Second, theories in cognitive neuroscience may specify the precise nature of processing within a single component subsystem. Such theories typically specify a type of neural network, which transforms input to a particular kind of output (e.g., Hasselmo, 1993; Hasselmo & Bower, 1993).

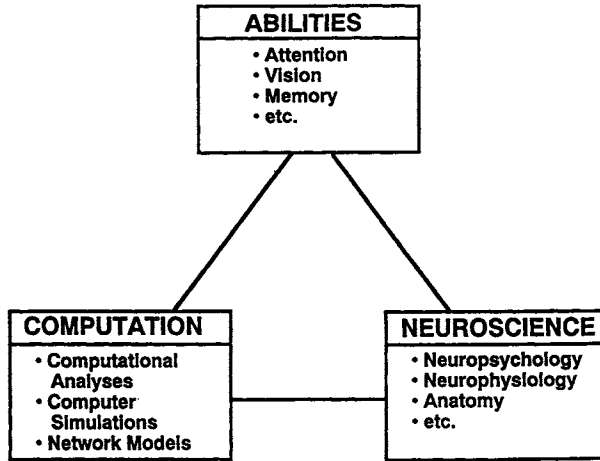


FIGURE 1 The cognitive neuroscience triangle.

B. The Isolable Systems Problem

Cognitive neuroscience thus is inherently multidisciplinary. Disciplines such as cognitive psychology, linguistics, and psychophysics are all concerned with delineating the nature of human abilities. The challenge here is to describe abilities in a way that can make contact with available data about brain function. A fundamental question researchers must answer in order to understand an ability has been termed the *isolable systems problem* (Posner, 1978): Does a given ability (or behavior that follows from it) come about through the function of a single system or many systems, and if many are involved, how do they coordinate their operations? This problem leads one to formulate and test a theory of a functional architecture. This problem is difficult to grapple with using behavioral data alone because any behavioral measure reflects the performance of the system as a whole (Posner & Carr, 1992). Even task-analytic procedures that rely on patterns of interaction between experimental variables (e.g., Sternberg, 1969) are subject to multiple interpretations (e.g., see Anderson, 1978; Townsend, 1974).

One approach to solving the isolable systems problem rests on the concept of convergent evidence. No one study provides conclusive evidence for a specific neurofunctional decomposition, but the results from a set of studies using diverse methods may point the way towards a single decomposition that explains all of them. Such studies have more power when the methods interlock, as occurs when one designs cognitive studies to investigate issues about the brain, and designs neuropsychological studies to investigate issues about cognition (Kosslyn, 1994). This is where neuroscientific data and techniques can be particularly important. For example, one can simultaneously record activity of neurons in different parts of the brain

while an animal performs a task that has been designed to require a specific type of processing; such findings may suggest which areas perform what computations and at what time (e.g., see Andersen, 1987). In humans, brain imaging techniques such as positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) allow us to determine which sets of brain areas are most active when a person performs a specific task (e.g., Posner, Petersen, Fox, & Raichle, 1988). If tasks are designed to tap specific types of processing, one can learn about the neural bases of such processing by comparing results from different tasks. The key is that

both brain imaging and neurophysiology add (on top of behavioral measures) the ability to break the system down into spatially defined parts in which the amount and timing of processing changes as a function of experimenter controlled manipulations of variables such as difficulty, type of task, decision rules, load, accuracy, etc.

(Posner & Carr, 1992, p. 8)

Given the truly staggering range of possibilities, it is no small thing to design tasks appropriately and to know where to look for specific effects in the brain. But these problems are more tractable if one has an hypothesis in hand. This hypothesis specifies one or more potentially isolable systems, and may or may not specify a candidate anatomical localization. Both parts of the hypothesis can be motivated in part by anatomical data that specify the connections between different areas. Such information not only can suggest separate processing components, as will be illustrated shortly, but also can be used to generate hypotheses about the flow of information in a system. These hypotheses can then be tested using brain imaging techniques, including time-sensitive measures such as event-related potentials (ERP) (e.g., Mangun, Hillyard, & Luck, 1992).

In addition to studying behavior and the brain, two types of computational data can help one to solve the isolable systems problem. First, *computational analyses* can lead one to formulate theories of how a given input can produce a given output. Such analyses rely on a careful consideration of the “problem” to be solved by an information-processing system, which often hinges on a consideration of the information that is available in the input (see Kosslyn & Koenig, 1992; Marr, 1982). Computational analyses typically result in hypotheses about the decomposition of a system into subsystems. These analyses must be informed by neurophysiology and neuroanatomy because we want to know how our cognitive system—not just any possible system—functions.

Second, one can construct *computational* models of hypothesized functional systems. These models are computer programs designed to mimic the operation of a dynamic system, and as such can help one to understand behavioral and neurobiological data in a number of ways. Models can lead one to discover unforeseen implications of a theory; by observing the behavior of an intact or “lesioned” model one can generate hypotheses about how the normal system functions, which can then be examined experimentally (e.g., predictions of Ambros-Ingerson, Granger, & Lynch, 1990, were tested by McCollum et al., 1991; see also S. Keele & Jennings,

1992). In addition, computational models can help address the isolable systems problem by specifying conditions under which it is more efficient to break a function down into component parts, each computed by a separate system, than to have the function carried out by a single system. For example, Rueckl, Cave, and Kosslyn (1989) demonstrated that some types of complex input–output mappings can be computed more efficiently by two networks rather than one, with each subnetwork specialized for carrying out different aspects of the mapping. By examining such models, one can experimentally determine when two mappings are “computationally incompatible,” and hence likely to interfere with each other if carried out within a single unified network.

In short, cognitive neuroscience can be characterized as having two general goals: First, it aims to carve the cognitive system at its functional and anatomical joints, along the way specifying the nature of, and interactions among, the component subsystems. Second, it aims to specify the ways specific neural networks operate to produce the requisite output when provided with an input. In both cases, the ultimate aim is to understand how computation in the brain confers specific abilities.

In the following three sections we consider how the cognitive neuroscience approach has begun to bear fruit in the study of some of our most fundamental mental abilities: selectively attending to objects, visual perception, and memory. There are several comprehensive cognitive neuroscience theories of attention (e.g., LaBerge, 1990; Posner & Petersen, 1990), visual perception (e.g., Hummel & Biederman, 1992; Kosslyn, 1994), and memory (e.g., Cohen & Eichenbaum, 1994; Schacter, 1990; L. R. Squire, 1987, 1992; see Schacter & Tulving, 1994a, 1994b, for summary of many recent theories). Space limitations preclude our considering each of these theories in detail. Our goal is not to review the literature exhaustively, but rather to convey the flavor of the cognitive neuroscience approach in action. Thus for each of the three content areas we will provide an overview of current theory, and illustrate the utility of a multidisciplinary approach and converging evidence.

II. ATTENTION

Attention is the selective aspect of information processing. This function allows us to focus on some information at the expense of other information. We typically are aware of what we attend to, and only specified pieces of information enter our conscious experience. Traditional conceptions of attention have posited either a limited “energy” resource or a structural bottleneck (Allport, 1992). Debate has focused on specifying exactly which types of processing do or do not require attention, which task variables play critical roles in demanding and directing attention, and exactly how far into the cognitive system information is processed before attention operates upon it (e.g., Broadbent, 1971; Shiffrin, 1988). However, as we have learned more about neural information processing, at least some of these questions have begun to appear ill posed (Allport, 1992). In particular, these questions are in large part predicated on the assumption that attention operates on information flowing

through the cognitive system in a precise, linear, increasingly abstract manner. However, the neuroanatomy suggests strongly that information processing in the brain is anything but simply linear and unidirectional (e.g., Felleman & Van Essen, 1991). Questions about capacity and the putative locus of attentional selection may be considered best with respect to particular types of information processed by particular components of neural systems (Allport, 1992; Posner & Petersen, 1990).

Attention has also been approached as a particular example of the isolable systems problem (Posner, 1978). Research in this mode begins with an analysis of the processing steps necessary for selective attention, which are then investigated by collecting a combination of behavioral and neurobiological data. Working within this paradigm, Posner and Petersen (1990) offer three general conclusions about attention: (a) the attention system is neurally distinct from, but interacts with, other processing systems of the brain; (b) this system consists of a network of different brain areas; and (c) each area carries out different computations that can be specified in cognitive terms. A series of seminal studies conducted by Posner and his colleagues (e.g., see Posner, Inhoff, Friedrich, & Cohen, 1987; Posner et al., 1988; Posner & Petersen, 1990) illustrates the utility of the cognitive neuroscience approach and will serve to flesh out our understanding of these three basic tenets.

A. Subsystems of Attention

At the computational level, attention can be viewed as involving the interaction of separable systems for (a) orienting to a stimulus; (b) detecting a stimulus; and (c) alerting and remaining vigilant for the appearance of a stimulus. The systems for orienting can be further distinguished as being used for overt (when the body, head, or eyes are moved) or covert (when no overt movement is made) shifts of attention. Furthermore, such shifts of attention appear to involve three processing subcomponents: in order to shift attention one must first disengage it from its current location, move it, and then engage attention at a newly specified location. The functional architecture of attention is illustrated in Figure 2.

1. Orienting to a Stimulus

An impressive amount about the mechanisms underlying spatial attention has been learned from a simple cuing task (e.g., Posner et al., 1978; Posner, Snyder, & Davidson, 1980). In this task, subjects first fixate on a cross and are cued to attend to a box that is either to the left or right of fixation. An asterisk then flashes in either the attended or the unattended box, and the subject simply presses a key as soon as he or she sees the asterisk. Subjects typically respond faster on validly cued trials, when the asterisk appears on the attended side, than on invalidly cued trials, when the asterisk appears in the box on the unattended side. The response time “cost” for invalidly cued trials has been interpreted as reflecting the time it takes to disengage

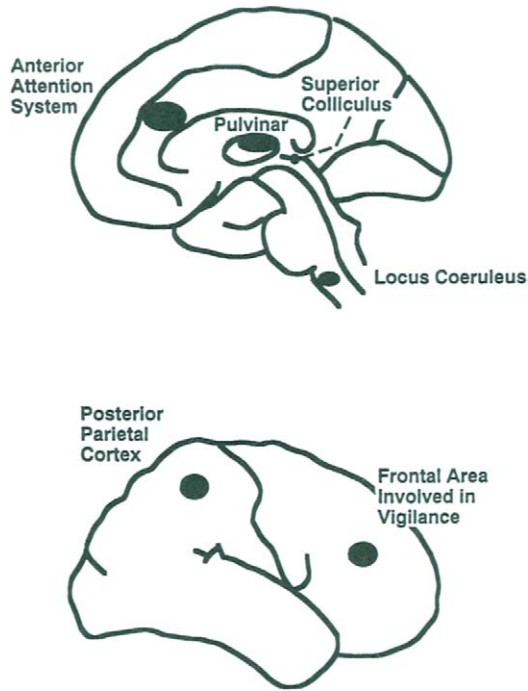


FIGURE 2 The functional architecture of attention shown superimposed on medial (top) and lateral (bottom) views of the right cerebral hemisphere. The neural locus of each subsystem is indicated by a solid black circle. These subsystems allow us to orient attention (posterior attention subsystem, which includes the posterior parietal cortex, pulvinar nucleus of thalamus, and superior colliculus), detect target stimuli (anterior attention subsystem), and maintain an alert, vigilant state (locus coeruleus and right hemisphere).

attention from the attended location in order to detect a target at the unattended location.

Neurophysiological studies have indicated that neurons in areas of the posterior parietal cortex (Wurtz, Goldberg, & Robinson, 1980), lateral pulvinar nucleus of the thalamus (D. L. Robinson & Petersen, 1992), and superior colliculus (Posner & Petersen, 1990) increase their firing rates when a monkey attends to a target stimulus to the exclusion of other distracting stimuli. These data suggest that these areas of the brain may be involved in attention, and in fact, patients with damage to any of these areas are impaired in the cuing task described above. Indeed, depending on the precise locus of the damage, patients are impaired at different aspects of the cuing task. First, patients with damaged parietal lobes have difficulty on invalidly cued trials, in which attention is initially focused in the incorrect location (Posner et al., 1987; Posner, Walker, Friedrich, & Rafal, 1984). These patients appear to have

particular difficulty moving attention away from an ipsilesional (same-side) cue to detect a target in their neglected visual field. Subsequent brain-imaging data have confirmed that the parietal lobes are activated when subjects shift their attention (Corbetta, Miezin, Shulman, & Petersen, 1993). These findings are consistent with the fact that damage to the parietal lobes often causes a deficit known as visual neglect (Bisiach & Luzzatti, 1978; Bisiach, Luzzatti, & Perani, 1979). For example, patients with right parietal damage appear unaware of or may not be able to respond to stimuli on the left side of space (Bisiach et al., 1979). Such patients typically exhibit *extinction*, or the loss of awareness of a stimulus appearing on the side contralateral to their lesion when it appears simultaneously with a stimulus on the same side as the lesion.

This deficit in disengaging attention differs markedly from the impairments exhibited by patients with lesions of the pulvinar nucleus of the thalamus (see Figure 2). Such patients are slow to respond to targets at cued locations (D. L. Robinson & Petersen, 1992), and similar results have been found with monkeys. Furthermore, PET scanning has revealed that the pulvinar becomes more active when subjects must attend to one aspect of a display to the exclusion of others (LaBerge & Buchsbaum, 1990). Thus the pulvinar seems to play a special role in engaging attention at a target location.

Finally, the ability to shift attention may be selectively disrupted by damage to the midbrain. In progressive supranuclear palsy, damage to the superior colliculus (see Figure 2) results in a slowing of responses to targets that appear at cued and uncued locations; the advantage for targets at cued locations appears only if the subjects are given a long time to focus on the cue before the asterisk appears; these patients apparently need extra time to shift attention to the cue (Posner, Choate, Rafal, & Vaughan, 1985).

2. Detecting a Stimulus

Researchers have also begun to understand the neural mechanisms that allow primates to detect behaviorally significant target events. However, theories of these mechanisms rest primarily on post hoc explanations of data; computational analyses that can motivate theory-driven research on the processes that underlie target detection are only now beginning to take shape. One reason for this is that studies have only recently revealed an area of the brain, anterior cingulate cortex, that appears to play a special role in target detection (see Figure 2). By examining the connectivity between the anterior cingulate and other parts of the brain, researchers can formulate hypotheses about interactions among specific subsystems (e.g., Ochsner & Baker, 1994). This is an example of how neuroscientific data can motivate a theory, which in turn prompts researchers to design behavioral experiments that bear on this new hypothesis.

Our knowledge about the function of the anterior cingulate in attention rests in large part on results from PET studies. For example, Petersen, Fox, Posner,

Mintun, and Raichle (1988) found anterior cingulate activation when subjects generated verbs that describe functions of nouns (e.g., when given “hammer” they might say “pound”), decided whether an animal was dangerous, or passively listened to words read aloud. Hypothesizing that this area might be involved in detecting targets, or selecting stimuli relevant to task demands, these researchers predicted, and found, that activity increased when greater numbers of targets were presented in the “dangerous animal detection” task. Anterior cingulate cortex is also active during performance of the Stroop task (Pardo, Pardo, Janer, & Raichle, 1990), during diffuse attention, as opposed to focal attention (Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1991), when painful stimuli are applied to the forearm (Talbot et al., 1991), when subjects generate attention-based visual mental images (Kosslyn, Alpert et al., 1993), and when subjects generate finger movements from memory (Dieber et al., 1991). In contrast, anterior cingulate activity decreases when subjects become more practiced in the verb generation task (Raichle et al., 1993), or when they mentally manipulate forms (Haier et al., 1988, cited in LaBerge, 1990); it also decreases when subjects are in a vigilant state, waiting to detect an infrequently presented target tone (see Posner & Rothbart, 1992).

The almost ubiquitous changes of activation in the anterior cingulate during task performance may suggest that it is something of a general-purpose attentional area, which is recruited whenever relevant stimuli cannot be detected on the basis of simple stimulus features or automatized routines (cf. Corbetta et al., 1991; LaBerge, 1990; Petersen & Fiez, 1993; Posner & Petersen, 1990; Posner & Rothbart, 1992). Posner and colleagues have used dual-task methods to test the generality of the anterior cingulate’s role in attention. When subjects shadowed speech while they also performed the cued probe-detection task described earlier, the response time difference between valid and invalid trials disappeared when the cues were presented to the left hemisphere. Passive listening to spoken words has been shown to activate the left anterior cingulate by Petersen et al., 1988, thus the anterior cingulate may have been engaged by the shadowing task, and hence was not able to confer an advantage for valid trials. Similarly, a concurrent auditory task can slow engagement of attention in parietal patients (Posner et al., 1987). Hence, there is evidence that language and visual spatial attention may share some common attentional mechanisms, although the nature of the shared computation(s) remains unclear.

Insights into the possible computations carried out by the anterior cingulate can also be garnered by examining the pattern of behavioral deficits that occurs when it is lesioned. Psychiatric patients for whom other interventions have failed sometimes receive bilateral stereotactic lesions in the rostral portions of anterior cingulate cortex, just above the genu of the corpus callosum. This operation is thought to alleviate anxiety (Ballantine, Cassidy, Brodeur, & Giriunas, 1972). Janer and Pardo (1991) examined the performance of one such patient on three tasks found in PET studies to activate anterior cingulate cortex: verb generation, identifying dangerous animals, and the Stroop task. Compared to her preoperative level of performance, the cingulotomy patient had deficits on all three tasks 2 weeks after the operation.

However, the deficits on the Stroop and identifying dangerous animals tasks disappeared 6 weeks later, which suggests that the attentional system can compensate (somehow) for small lesions. It is clear, however, that whatever the anterior cingulate does, that computation or computations is normally drawn upon when one performs these tasks.

Additional information about the role of the anterior cingulate comes from research on discriminative aversive conditioning in rats. Although it is always dangerous to generalize across species, basic sensory and motor processes (and at least some forms of attention may be included in these categories) are often similar among different mammals; in any case, findings about rat brains are a good source of plausible hypotheses about processing in the human brain. Such research has shown different patterns of firing in neurons in anterior cingulate cortex to a conditioned stimulus than to a stimulus that was not conditioned, and lesions of this area impair acquisition—but not expression—of discriminative avoidance behavior (Gabriel, 1990). These results suggest that the anterior cingulate cortex helps to identify behaviorally significant stimuli. This function is consistent with the fact that this area has major connections to the amygdala—which plays a critical role in emotion (Amaral et al., 1992).

We might expect that a system playing a general role in detecting target events would enjoy widespread connections with cortical areas involved in attention, memory, and motor control. And in fact, Goldman-Rakic (1988) has documented the connections between the anterior cingulate and some of the other areas known to be involved in attention, specifically the posterior parietal cortex and possibly the pulvinar nucleus; the anterior cingulate is also connected to the dorsolateral prefrontal cortex and parahippocampal cortex, which are involved in short-term visual-spatial and long-term object memory, respectively. In addition, within the cingulate sulcus, the anterior cingulate has reciprocal connections with primary and supplementary motor cortices (Barbas & Pandya, 1981; V. B. Brooks, 1986; Vogt & Miller, 1983), and the neurons there are sensitive to errors made during motor skill learning (V. B. Brooks, 1986).

Thus the anatomy and neurophysiology of the anterior cingulate suggest that it has a general role in attention. Posner and Petersen (1990) have termed anterior cingulate cortex “the anterior attention system,” distinct from the “posterior attention system,” which consists of posterior parietal cortex, pulvinar thalamus, and the superior colliculus (as summarized earlier). Posner and Petersen conceptualize the anterior system as a general purpose target detector, which gates various components of the posterior system as well as mediates attention to other functions such as language. Recent research suggests that the cingulate may be specialized not just for detecting targets, but for monitoring the relationship of stimuli to the goals of the individual (for discussion see Ochsner & Feldman-Barrett, in press). This is suggested by the finding that painful stimulation (e.g., Rainville et al., 1997) and attention to one’s current emotional state activates the cingulate (Lane et al., 1998), and by event-related potential (Gehring et al., 1993) and functional magnetic resonance imaging studies (Carter et al., 1998) that show cingulate activity when participants

make an error in simple reaction time tasks (cf. Brooks, 1986). It is possible that different areas of the anterior cingulate subserved slightly different, but related functions (Posner & DiGirolamo, 1998) and future work will serve to differentiate them.

3. Maintaining Vigilance

The brain stem and right hemisphere apparently play key roles in alerting and maintaining a vigilant, aroused, state (see Figure 2). Norepinephrine (NE) released by the locus coeruleus (a structure in the brain stem) apparently modulates the alert state (Aston-Jones, Foote, & Bloom, 1984), and right-hemisphere lesions lead to depletion of NE in *both* hemispheres (R. G. Robinson, 1985). Furthermore, NE strongly innervates the thalamus and parietal cortex (Morrison & Foote, 1986), and NE agonists (which facilitate the uptake of NE by receptors) may enhance processing in the parietal cortex, speeding the disengage operation (Clark et al., 1989) and thereby speeding the process of attentional selection (Posner & Petersen, 1990). Given these data, we would predict that patients with damage to the right—but not patients with damage to the left—hemisphere would have a deficit in alerting (Coslett, Bowers, & Heilman, 1987). As expected, Posner et al. (1987) found that patients with damage to the right parietal lobe had an increasingly smaller benefit from validly cued targets as the delay between cue and target increased: the patients were unable to keep attention engaged at the cued location over a short delay. Furthermore, PET studies have shown that regions of the right frontal lobe are activated during maintenance of a vigilant state (Corbetta et al., 1991, 1993).

B. Summary

Attention can be divided into three major systems, and at least one of these systems can in turn be divided into two subsystems (for a caveat, however, see Farah, 1994). The emerging theories have been built on a convergence of findings from different patient populations, brain-imaging techniques, and behavioral results from normal subjects. Advances thus made have the effect of systematizing and concretizing our notions of attention and “attentional resources” while at the same time providing a testable framework that makes contact with research in other domains. Such a framework provides a starting point for examining the roles of other brain areas in attention. For example, recent work indicates that the basal ganglia (Alexander, Crutcher, & DeLong, 1990; Clark et al., 1989; Jackson & Houghton, 1995) may modulate interactions between the anterior and posterior attention systems. As we shall see in the following section, attention plays an important role in perceiving the visual world.

III. HIGH-LEVEL VISION

A hallmark of the human visual system is the ability to recognize and identify objects presented in various orientations, from different perspectives, and in many differ-

ent viewing conditions (such as poor lighting or partial occlusions; see Kosslyn, 1994, for a taxonomy of these abilities). It is useful to distinguish between low-level and high-level vision. Low-level visual processing is bottom-up, driven solely by properties of the perceptual input that strikes the retina. It is concerned with specifying information such as edges, regions of homogeneous color or texture, and depth. In contrast, high-level visual processing makes use of stored information to help one identify an object or use stored knowledge to guide reaching and navigation. We focus here on the mechanisms that underlie high-level vision, which are of most interest to cognitive scientists.

A. Subsystems of High-Level Vision

Kosslyn (1994; see also Kosslyn & Koenig, 1992) has argued that the system subserving high-level vision can be broken down into a set of major subsystems, each of which is instantiated in a discrete cortical area. These subsystems are illustrated in Figure 3. We briefly describe each subsystem below.

1. Visual Buffer

When viewing an object, information from the eyes is passed through the brain stem and thalamus to the primary visual cortex. From the primary visual cortex, information is distributed to over a dozen distinct visual areas in the occipital lobe (see Felleman & Van Essen, 1991). These areas are “retinotopically organized”: their




FIGURE 3 The functional architecture of high-level vision shown superimposed on a lateral view of the left cerebral hemisphere. A set of subsystems (described in text) allow one to recognize and identify objects. The putative location of each subsystem is indicated by a solid black circle. Arrows indicate possible directions of information flow between subsystems. Information flowing in a bottom-up fashion from lower to higher level areas follows paths marked with solid arrowheads. Information can also flow between higher level areas or in a top-down fashion from higher to lower level areas, following paths marked with open arrowheads.

spatial structure corresponds (approximately) to that of the retina itself; however, these maps typically are distorted so that there is a disproportionately large area devoted to the high-resolution central portion of the retina, and not all of the remainder of the visual field is represented equally well. A particularly vivid demonstration of the existence of such areas was reported by Tootell, Silverman, Switkes, and DeValois (1982), who had monkeys view a flashing circular spoked pattern after injection of radioactive sugar. The more a neuron fired while the animal watched the pattern, the more sugar the neuron took up, and hence the more radioactivity was taken up. The monkeys were then sacrificed and their cortices “developed” so that cells that had taken up the tracer were visible. Tootell et al. showed that in the primary visual cortex (in addition to other areas) there was a physical pattern of active cells laid out on the surface of the cortex in roughly the same shape as the spoked pattern; the map was distorted so that parts of the pattern that fell on or near the fovea received larger representation.

Kosslyn (1994) groups into a single functional structure the set of retinotopically mapped areas that work together to segregate figure from ground. This structure is called the “visual buffer.” It is clear that this component can be decomposed into more specialized components; indeed, in the monkey some of the constituent areas include a preponderance of neurons that are sensitive to wavelength (area V4), others to motion (e.g., area MT), and so on. Nevertheless, patterns of activity in the set of areas that comprise the visual buffer preserve key features of the local geometry of images that strike the retina. Data from patients with occipital lobe damage who cannot see in particular regions of the visual field (e.g., Holmes, 1918), and more recent data from PET studies (e.g., Fox et al., 1986) confirm that this conclusion can be extended to the human brain.

2. Attention Window

There is much more information in the visual field than can be processed at any one time; hence some of this information must be selected over others. The mechanisms that orient attention (discussed in the previous section) not only shift one’s body, head, and eyes so that a specific stimulus is fixed, but also can shift the locus of attention covertly. An internal “attention window” selects patterns in the visual buffer for further processing (for a review of supporting evidence, see Kosslyn, 1994). We are led to infer the existence of such a mechanism by the fact that subjects can covertly shift attention over an ionic image (Sperling, 1960) or display (e.g., Posner et al., 1980) to search for a particular item. Furthermore, A. M. Treisman and Gelade (1980) have shown that in some circumstances such covert attention is necessary to bind together the location and form of an object.

The position of the attention window gates the information that is passed along for further processing. For example, consider the results from an experiment reported by Moran and Desimone (1985). They first located neurons in monkeys that responded selectively to a certain stimulus (e.g., a vertical green bar). They then

mapped out the receptive fields of these neurons; a receptive field is the area of space where a stimulus will drive the neuron. Moran and Desimone then rewarded the monkeys for responding to stimuli that appeared only in one quadrant of the receptive field of a cell. After such training the cell fired vigorously to stimuli in the reinforced quadrant and would still show some response to stimuli appearing in other quadrants—but responses to stimuli in nonreinforced quadrants were quickly squelched. In this case at least, it seems clear that the “engage” component of attention is operating via inhibition: stimuli in the unselected regions begin to evoke increased neural activity, but this activity is soon suppressed.

3. Ventral and Dorsal Systems

Information selected by the attention window is sent along two parallel cortical pathways, one specialized for processing the “object properties” of a stimulus, such as its shape and color, and the other specialized for processing the “spatial properties” of a stimulus, such as its location and orientation. Ungerleider and Mishkin (1982) term these the “what” and “where” pathways, or ventral and dorsal systems because they are located in the inferior temporal and posterior parietal lobes, respectively. This distinction between the ventral and dorsal systems is motivated by a number of different findings. In the monkey, removing the inferior temporal lobes devastates the ability to recognize shapes of objects but not the ability to recognize location; in contrast, removing the parietal lobes devastates the ability to recognize spatial locations but not the ability to recognize shape. For example, Ungerleider and Mishkin (1982) trained monkeys to select food hidden under one of two lids; if the monkeys had to select a lid with a particular pattern in order to get the food, inferior temporal lobe lesions impaired performance, whereas if they had to select the lid closest to a visual landmark, parietal lobe lesions impaired performance. Consistent with these findings in monkeys, damage to the posterior inferior temporal lobes of humans may impair perception of the visual form of objects, whereas parietal damage impairs orientation in space (e.g., Farah, 1990; Kosslyn, 1994; Levine, 1982).

In addition, single-cell recording studies in monkeys have found neurons in inferior temporal cortex that are sensitive to shape and color (Desimone, Albright, Gross, & Bruce, 1984; Gross, Desimone, Albright, & Schwartz, 1984; Maunsell & Newsome, 1987; Perrett et al., 1985). These neurons typically have very large receptive fields, and are relatively insensitive to an object’s location (Gross & Mishkin, 1977); such cells may underlie our ability to recognize objects regardless of their spatial location (see also Kosslyn, 1994). In contrast, cells in posterior parietal cortex fire in response to the location, size, and motion of an object (Andersen, 1987; Andersen, Essick, & Siegel, 1985; Hyvarinen, 1982; Maunsell & Newsome, 1987).

In addition, PET studies of face comparison (Haxby et al., 1993), face recognition (Sergent, Ohta, & MacDonald, 1992), and object recognition (Kosslyn, Alpert et al., 1994) have documented activation in inferior portions of the temporal lobes. In contrast, PET studies that require encoding spatial relations have shown activa-

tion of the posterior parietal lobes (typically the inferior portion; e.g., Corbetta et al., 1993). Further support for this distinction comes from psychophysical studies in humans. These studies have shown that information about location and shape can independently influence perception (Sagi & Julesz, 1985; A. Treisman & Gormican, 1988).

The division of higher-level visual processing into two major processing streams makes sense from a computational point of view. As noted earlier, Rueckl et al.'s (1989) computational models showed that a single network that identified both an object's form and spatial location is substantially less efficient than two subnetworks, one for each computation (provided that enough resources were allocated to the subnetworks). Just as in humans, the single-network model needed to ignore location to recognize the shape in different locations, but needed to encode location to specify it in the output. Encoding object identity and spatial relations apparently were "computationally incompatible" processes, and hence were difficult to compute in the same system.

4. Associative Memory

Processing in the ventral system can allow one to *recognize* an object; recognition occurs when the shape matches the stored representation of another shape. But the ventral system is modality-specific: it only encodes visual input. One knows that an object is familiar after it has been recognized, but knows nothing else about it. In order to *identify* an object, one needs to access representations of its categories, its name, and various other kinds of nonvisual information. Identification can occur even if recognition is not very good, provided that the object has strong spatial cues (e.g., such as occurs when one encodes the size of an ant). Thus, information from the ventral (what) and dorsal (where) pathways must make contact with information stored in a long-term "associative memory" (which may or may not be further divisible into an "episodic" and "semantic" memory; for our purposes, we need not take a position on this issue). This memory system stores relations among object and spatial properties, as well as other attributes such as names and categories to which an object belongs. The same information in associative memory can be accessed when an object is recognized in any modality, as would occur if one heard a cat meow, felt it caress one's shin, or saw it walking towards one.

The literature is vague with respect to the locus of associative memory. However, the object and spatial properties systems are known to converge on the dorsolateral prefrontal cortex as well as regions of the parietal-temporal junction (Goldman-Rakic, 1988). The dorsolateral prefrontal regions appear to store information temporarily, as part of "working memory" (e.g., Goldman-Rakic, 1988), and hence this region is not a good candidate for the site of a long-term associative memory structure. In contrast, lesions in the region of the parietal-occipital junction can result in deficits in linguistic and semantic processing (e.g., see Geschwind, 1965).

5. Information Look-up

Under ideal conditions, one can identify an object via the route described thus far. However, if the input image of an object is degraded, or the object projects a novel shape (perhaps because it is contorted or seen from an unusual viewpoint), it may not correspond well to a representation in the ventral system. In such circumstances, the bottom-up input may lead only to a tentative hypothesis about an object's identity. In such straits, one can look up information in memory that would support this hypothesis, but has not yet been detected. This information can then guide one to search "top down" for this decisive part or characteristic (e.g., such as a particular dimple on the bottom of a Delicious apple; see Gregory, 1970; Kosslyn, 1994).

The frontal lobes are the likely locus of systems used to look up possibly diagnostic information in memory. PET studies have found activation of regions in the frontal lobe when subjects are retrieving from memory information about objects (Petersen et al., 1988; Tulving et al., 1994). In addition, retrieving stored information and holding it temporarily on-line are important for formulating and testing hypotheses, and a substantial literature indicates that lesions to the frontal lobe impair this ability. For example, the Wisconsin Card Sort test requires subjects to infer a rule that relates patterns on successively presented cards. This rule changes periodically, and patients with damaged frontal lobes perseverate, or get stuck, using one rule even when they realize that the rule has changed (e.g., Milner, 1964).

6. Attention Shifting

In the previous section we discussed the mechanisms underlying attention shifting. In addition to those mechanisms, the frontal lobes play a role in using information accessed from memory to shift attention. The frontal eye fields (also known as Area 8) direct voluntary eye movements, and frontal lesions disrupt systematic visual search and visual working memory (Alexander et al., 1990; Luria, 1980; D. L. Robinson & Petersen, 1986). Frontal lesions can also cause a form of unilateral neglect (Heilman & Valenstein, 1985), which might be expected given their rich connections with the anterior and posterior attention systems (Posner & Petersen, 1990).

Once one has shifted one's attention to the location where a diagnostic part or property should be located, that pattern is recognized and identified. If the expected part or property is present, one may have enough information to identify the object. If not, additional information may need to be encoded.

B. Summary and Extensions of the Theory

In summary, information striking the retina sets up a pattern of activation in a set of retinotopically mapped regions of cortex, which we call the visual buffer. Some of this information is selected by an attention window for further processing, and this information is passed to the object-properties and spatial-properties encoding

systems, which operate in parallel. The outputs from these systems are sent to a long-term associative memory structure. If the set of information reaching associative memory is not consistent with the properties of a single object, the best matching description in associative memory is treated as an hypothesis. This hypothesis in turn guides a top-down search for a distinctive part or characteristic, which will either confirm or reject the hypothesis. The frontal lobes play a key role in this top-down search process; mechanisms implemented there retrieve information from memory, that guides attention to select disambiguating information. This process is repeated until the object is identified.

Each of the major component processes just described can in turn be divided further. For example, the dorsal (spatial-properties encoding) system consists of at least three distinct subsystems, which have different functions. The “spatiotopic mapping” subsystem converts the retinotopic coordinates of the visual buffer, which depend on where one’s eye is positioned, to spatiotopic coordinates, which are anchored in external space. The “categorical spatial relations encoding subsystem” encodes spatial relations such as above/below, left/right, and on/off; this subsystem operates more effectively in the left than in the right cerebral hemisphere. And the “coordinate spatial relations encoding subsystem” encodes metric spatial relations, and it operates more effectively in the right cerebral hemisphere (see Hellige & Michimata, 1989; Kosslyn, 1987; Kosslyn et al., 1989; Laeng, 1994; but see also Sergeant, 1991, versus Kosslyn, Chabris, Marsolek, & Koenig, 1992, and Cook, Fruh, & Landis, 1995, versus Kosslyn, Chabris, Jacobs, Marsolek, & Koenig, 1995). Similarly, the ventral (object-properties encoding) system can be divided into more specialized subsystems that extract distinctive features, that match such features to stored memories, and so on.

The theory of high-level vision is more detailed than the theory of attention, in large part because of the enormous volume of research on vision. What do such detailed theories buy us? For one, they allow us to interpret a large body of data, which addresses computational, neural, and behavioral properties. Because such a theory must accommodate a wide range of different types of findings, it is likely to have more general and powerful principles than a theory that is restricted to only one type of data. In addition, we have seen that a multicomponent theory of attention can help us understand the deficits exhibited by patients and make predictions about the roles components of the system should play in different tasks; the same is true for the theory of high-level vision (e.g., Kosslyn & Koenig, 1992). The following two examples illustrate the utility of such a theory.

Warrington and her colleagues have found that patients with posterior cortical lesions have difficulty recognizing objects that are seen from unusual (noncanonical) points of view, but not objects seen from a usual (canonical) perspective (e.g., Warrington & James, 1991; Warrington & Taylor, 1973, 1978). This finding makes sense within the framework just developed because posterior lesions may disrupt the spatial properties encoding system; damage to this system may impair one’s ability to recognize objects shown in unusual views because their three-dimensional

structure cannot be reconstructed. However, Warrington also found that patients with frontal lobe lesions were not impaired when asked to recognize objects seen from unusual views, and this finding is not as predicted by the theory: presumably objects shown from unusual views are difficult to recognize initially, and hence one would typically identify them only after engaging in top-down search—and this process is putatively guided by the frontal lobes.

Kosslyn et al. (1994) suggested that the reason Warrington and her colleagues did not find deficits in frontal lobe patients in this task was because they failed to record response times: One can locate a distinctive part or characteristic by random search, but this method would take longer than when top-down search can be employed (and hence one can use knowledge to search immediately for distinctive parts or characteristics). To test the hypothesis that the frontal lobes play a role in top-down search when objects are viewed from unusual perspectives, Kosslyn et al. used PET to compare the brain areas that were active when subjects identified objects that were seen from typical points of view with the areas that were active when they identified objects seen from unusual points of view. In one condition the subjects decided whether objects shown from a typical perspective matched an object name; in another condition the subjects performed the same task with objects seen from unusual perspectives. To isolate the brain areas that were activated selectively when the subjects identified objects seen from unusual points of view, which the theory predicts should reveal evidence of the role of the frontal lobes in top-down search, Kosslyn et al. subtracted cerebral blood flow recorded in the typical-view condition from that recorded in the unusual-view condition.

As predicted, dorsolateral prefrontal cortex in both hemispheres was more active when subjects identified objects seen from unusual points of view. This is good evidence that this region plays an important role in looking up information in memory to test hypotheses. The specific locus of activation was similar to that reported by Petersen et al. (1988) when subjects accessed information about uses and functions of objects. Furthermore, as was also expected, the set of brain areas predicted to be involved in object identification was also activated: the occipital cortex corresponding to the visual buffer was active, as were areas of the parietal lobe associated with shifting attention, and areas of the parietal lobe associated with encoding spatial properties (part of the dorsal system), and the inferior and middle temporal lobes (part of the ventral system). Moreover, an area at the occipital-temporal-parietal junction was activated, which may play a critical role in implementing associative memory.

A second example illustrates how knowledge of the systems involved in high-level vision can inform cross-domain hypothesis testing (Schacter, 1992). This study used the theory of visual perception to illuminate the nature of the neural mechanisms involved in visual mental imagery. Historically, much debate has surrounded the status of mental images; in recent years, much interest has focused on questions about the nature of the representation underlying imagery and the relation of imagery to perception. Marshaling evidence from various disciplines, Kosslyn (1980,

1994) argued that visual mental images are depictive (i.e., that they use space to represent space, thereby preserving geometric properties of imaged objects), and furthermore that such images correspond to patterns of activation in the visual buffer. In fact, according to this theory, imagery relies on many of the same neural mechanisms as high-level visual perception. Specifically, frontal lobe mechanisms access stored information from associative memory, which is used to activate visual information stored in the ventral system; this information in turn engenders an image proper by causing a pattern of activation in the visual buffer—this inverse mapping procedure apparently is necessary because visual memories are not stored as topographic representations, but rather as “population codes” (e.g., Fujita, Tanaka, Ito, & Cheng, 1992). Additional parts can be added to an imaged object by shifting the attention window over it, and activating stored representations of parts or properties so that they are positioned in the correct relative locations (see Kosslyn, 1994, for a detailed theory of how such processing may occur).

According to this theory, once the geometric properties of an object have been reconstructed in the visual buffer, the object properties and spatial properties of the imaged object can be reinspected. For example, once one has formed an image of a German shepherd dog, one can “see” the shape of its ears (an object property) and also determine which is longer, its tail or rear leg (a spatial property).

Some researchers have challenged this theory. Not only have some (e.g., Pylyshyn, 1973, 1981) suggested that image representations are language-like propositions (and the depictive properties evident to introspection are epiphenomenal, like the heat of a lightbulb while one is reading), but others have questioned the commonality of the neural systems underlying imagery and perception (e.g., Roland & Gulyas, 1994, versus Kosslyn & Ochsner, 1994). Kosslyn, Alpert et al. (1993) tested these claims using several PET studies of imagery. In one, the subjects closed their eyes and visualized letters at either small or large sizes. Not only was the topographically organized visual cortex activated during this task, but the locus of activation depended on the size of the imaged letters; indeed, the precise coordinates of the activated regions were close to what one would predict if subjects were actually viewing objects at the corresponding sizes.

Other studies in this series were designed to study image generation, the process of building up an image from stored information. The theory predicts that the same areas used to encode objects when top-down hypothesis testing is used should be activated when an image is built up from parts. In this case, instead of searching for a distinctive part or property at a particular location, one searches for the location in order to add another part or property to the image. Subjects viewed a 4×5 grid with a lowercase cursive letter printed underneath. An X mark occupied one cell of the grid. Subjects either simply saw the stimuli and responded (in a baseline condition), or visualized the corresponding uppercase letter in the grid and decided whether it would cover the X if it were actually in the grid. As predicted, a very similar set of brain areas was activated when visual mental images are formed and when top-down search is used in visual perception: when blood flow in the base-

line condition was subtracted from blood flow in the imagery condition, very much the same areas were identified as were identified when blood flow in the typical-point-of-view condition was subtracted from blood flow in the unusual-point-of-view conditions in the object-identification task described above (the same subjects participated in both sets of tasks; see Kosslyn, Thompson, & Alpert, 1997). The fact that such similar patterns of activity were observed in such seemingly different tasks (evaluating names of pictures versus visualizing letters in grids) is strong evidence that the theory is on the right track.

In the following section we will consider how research on memory reveals additional properties of some systems used in both vision and imagery.

IV. MEMORY

Memory allows us to use knowledge gained from previous experience to guide current and future actions, and is the cornerstone of many cognitive processes. Indeed, memory is crucial for identifying and recognizing objects that our attentional systems have selected for further processing. Memory, like visual perception and attention, is accomplished by a set of subsystems working together. One not only can store and recall the meanings of and associations among words, images, and concepts, but also can recognize objects and encode relationships among particular stimuli and visceral or motor responses. Each of these abilities is accomplished primarily by a distinct system or set of systems. In this section we consider more fully the memory encoding and storage systems that play critical roles in visual perception and attention, and we also consider memory systems involved in encoding and storing other types of information. The view presented here draws on and is consistent with aspects of many contemporary theories of memory (e.g., Schacter, 1990; Squire, 1992), but is derived primarily from the analysis offered by Kosslyn and Koenig (1992). The functional architecture of memory is illustrated in Figure 4.

A. Perceptual Encoding Subsystems

In order to recognize an object we must have previously stored a representation of its object properties. These representations are stored in perceptual encoding subsystems that store the structural and feature properties of modality-specific inputs (Kosslyn & Koenig, 1992; Schacter, 1990). Examples are the object-properties-encoding ventral system and spatial-properties-encoding dorsal systems discussed above, although every sensory modality has its own perceptual encoding subsystems.

1. Object Properties Encoding Subsystem

After initial processing by the visual buffer in the occipital lobe, information is passed along to the object properties encoding subsystem in the inferior temporal




FIGURE 4 The functional architecture of memory shown superimposed on (A) a medial view of the right hemisphere, (B) a lateral view of the left cerebral hemisphere, and (C) a transparent view of the left hemisphere. A set of subsystems (described in text) encode and/or store different types of information. The putative neural locus of each subsystem is indicated.

lobe (see Figure 4B). Research in monkeys has revealed some of the basic characteristics of representations in this subsystem. Neurons in this area are sensitive to the form, color, and shape, but not orientation or size of objects (Desimone et al., 1984; Gross et al., 1984; Gross & Mishkin, 1977; Maunsell & Newsome, 1987; Perrett et al., 1985), and lesions impair memory for the form of an object but not memory for its relationship to other objects in space (Levine, 1982; Pohl, 1972; Ungerleider & Mishkin, 1982).

Much has been learned about the nature of object representations in normal

human subjects from studies of priming. In a typical priming task, subjects first are shown a set of objects or words and asked to make some simple decision about their perceptual or semantic attributes; this task requires them to look at each stimulus and produces an “incidental” memory representation. Later, degraded versions of these stimuli are presented along with degraded versions of new objects, and the subjects are asked to identify, read, or make some decision about them; priming is assessed by measuring the gain in performance for the previously seen stimuli compared to the new ones. The notion is that when a stimulus is encoded initially, one or more representations are activated in memory; this activation decays rather slowly, and hence subjects can subsequently encode the stimulus more easily if it appears soon enough after it was shown initially. Similar to the neurophysiological findings in monkeys, studies of priming for familiar objects (e.g., a shoe) have shown that primed identification of pictures is long-lasting (Cave & Squire, 1992; Mitchell & Brown, 1988) and is unaffected by study-to-test changes in object size or reflection in both normal subjects (Biederman & Cooper, 1992) and amnesics (Cave & Squire, 1992). Although we will discuss amnesia in more detail below, this latter finding is important because it indicates that perceptual representations can guide performance even when they are not accessed consciously.

Schacter, Cooper, and their colleagues (e.g., Schacter, Cooper, & Delaney, 1990) have used an object decision priming task to study the nature of the representations stored in memory. In their task, subjects decide whether drawings depict structurally possible or impossible three-dimensional objects. These objects are novel, and hence a new representation must be encoded for all of them during the initial exposure phase; priming is measured by comparing the errors when previously shown and new objects are subsequently presented very briefly, and the subjects are asked to determine whether the object is structurally possible or impossible. Priming in this task is found only for possible objects (Schacter, Cooper, & Delaney, 1990), is preserved in amnesics (Schacter, Cooper, Tharan, & Rubens, 1991), and depends upon encoding the global three-dimensional structure of the object when it was first shown (Schacter, Cooper, Delaney, Peterson, & Tharan, 1991). Thus it appears that representations in the object-properties encoding subsystem incorporate regularities that characterize actual objects. As one would expect given the properties of interior temporal lobe neurons in monkeys, priming in this task is not affected by changes in the size of an object or changes in the direction it faces; however, priming is reduced by changes in orientation on the picture plane (Cooper, Schacter, Ballesteros, & Moore, 1992; Schacter, Cooper, & Treadwell, 1993; for a review and interpretation of these and similar findings, see chapter 5 of Kosslyn, 1994).

The fact that the object-properties encoding subsystem cannot represent impossible objects easily does not imply that it can only represent well-formed objects. Rather, it appears to store perceptual representations of objects and parts of objects. For example, when a subject is shown a picture of an object that has had many of its recognizable features eliminated, and the global structure of the object itself is very difficult to recover based on this picture, subjects show greater subsequent

priming for that picture of the fragmented object than for a picture of the whole, undegraded object—even though the picture of the whole object is in some sense “less degraded” (Srinivas, 1993).

We earlier distinguished between modality-specific perceptual representations that underlie recognition and amodal representations in “associative memory.” If this distinction is correct, then we might predict that brain damage can disrupt one structure while leaving the other intact. And in fact, when the cortical areas that implement the object-properties encoding subsystems are damaged, subjects may have a “visual object agnosia” (see Farah, 1990; Kosslyn & Koenig, 1992): They cannot identify the object visually, but can identify it via other sensory modalities (e.g., by touching it). This disorder was originally characterized as “mind blindness.” Such damage impairs recognition of objects, but not access to their semantic attributes. In contrast, selective brain damage may produce the opposite pattern of deficits: one may lose the ability to access semantic, but not perceptual characteristics of objects. The neuropsychological literature includes many reports of patients with cortical lesions who perform normally when asked to match or copy pictures, to decide whether a design represents a real object or is nonsensical, and similar visual tasks, but are impaired when asked to display knowledge of the semantic attributes of pictured objects, such as naming or describing an object’s function (e.g., Farah, 1990; Riddoch & Humphreys, 1987; Warrington & Taylor, 1978; see Kohn & Friedman, 1986, for analogous deficits in audition).

2. A Word Form System?

Recognizing letters or words is similar to recognizing objects, but is not exactly the same: Compared to objects, words are defined solely by patterns of lines whose meaning has been arbitrarily assigned; recognizing them does not require computation of global, three-dimensional structure, and for many adults words are more familiar than are most objects. Such observations have led some to argue that word forms are stored in a distinct visual memory (e.g., J. L. McClelland & Rumelhart, 1981; Petersen & Fiez, 1993; Petersen, Fox, Synder, & Raichle, 1990). It is possible that frequent exposure to words biases the object-properties encoding subsystem to dedicate part of its structure to encoding words; if so, we might expect word recognition to involve a brain area distinct from those used in object recognition. Consistent with this view, Petersen et al. (1990) found that both real words and non-words that could be words (according to the rules of English) activated an area of left medial extrastriate cortex; this area is distinct from areas of the temporal lobe that are activated when one recognizes objects or faces (Kosslyn, Alpert et al., 1994; Sergent et al., 1992). In addition, some brain-damaged patients have difficulty accessing word meanings but can recognize word forms and identify objects (Warrington & Shallice, 1980). Thus, part of the object properties pathway may be specialized for representing highly familiar words; we do not yet know, however, whether this pathway deals with all highly familiar stimuli, or words per se. Various

researchers have posited a *word form subsystem* (see Figure 4A) that represents the perceptual or orthographic properties of words (e.g., Schacter, 1990; Warrington & Shallice, 1980).

Priming studies have revealed properties of the representations of word shapes in memory. These tasks often require the subjects first to view a set of words, and later to complete three-letter “word stems” or fragments with the first words that come to mind. Alternatively, a subject might be asked to identify a briefly presented word. The increased probability of completing the fragments to form one of the initially seen words or of identifying the briefly presented stimulus is the measure of priming. By and large, the findings with these tasks dovetail nicely with the findings for objects, though there are some important differences. As has been reported with object priming, word priming is long lasting (MacAndrews, Glisky, & Schacter, 1987; Sloman, Hayman, Ohta, Law, & Tulving, 1988), and semantic encoding during the initial exposure phase (e.g., having the subjects judge the number of meanings of each word on the list) enhances recall, but has little or no effect upon priming (e.g., Graf, Squire, & Mandler, 1984; Jacoby & Dallas, 1981; H. L. I. Roediger, Weldon, Stadler, & Riegler, 1992). Moreover, priming is substantially reduced when the presentation modality is changed (e.g., auditory to visual; Jacoby & Dallas, 1981; H. L. Roediger & Blaxton, 1987). Like objects, studies of word priming have shown that the object properties subsystem encodes highly specific features of perceptual input. For example, changes in typefont or letter case can reduce priming (e.g., Hayman & Tulving, 1989; Jacoby & Hayman, 1987; H. L. Roediger & Blaxton, 1987), but usually only when the incidental encoding task (administered when subjects are given a list of words at the outset of the experiment) focuses the subjects on the perceptual characteristics of the words (e.g., counting the number of T-junctions in the letters of a word; Graf & Ryan, 1990).

Additional findings suggest that the enhanced priming when the identical form is presented during the initial exposure phase and the test phase arises from a particular type of word form system, which is localized in the right cerebral hemisphere. Marsolek, Kosslyn, and Squire (1992) found that changes in typefont had no effect upon word-stem completion priming when word stems were presented to the left hemisphere at test (this is done by having the subject stare straight ahead and flashing the stem to the left or right—which causes the input to be encoded initially by the right or left hemisphere, respectively). In contrast, preserving the typefont enhanced the amount of priming when word stems were presented to the right hemisphere at test. This led Marsolek et al. to infer that a right-hemisphere system stores form-specific representations, and a left-hemisphere system stores more abstract visual form representations. Furthermore, PET investigations of word-stem completion priming have found a decrease in activation of right extrastriate occipital cortex when typefont is unchanged between study and test, which may reflect that priming has facilitated processing (Squire et al., 1992). Similar studies of object priming have yet to be reported.

B. Associative Memory

As discussed earlier, all perceptual encoding systems send input to associative memory (see Figures 3 and 4B; see Kosslyn & Koenig, 1992, for extended discussion). Three characteristics of associative memory are of interest in the context of memory per se. First, although it encodes relations among perceptual representations, associative memory represents this information in an abstract or propositional format. These propositions can specify complex relations such as “is a,” “has a,” and so on. These relations are qualitatively distinct and often abstract, and so could not be implemented by simple direct connections between perceptual representations.

Second, relations in associative memory appear to involve pointers back to representations in the perceptual subsystems, and in that sense have “meaning.” These pointers are bidirectional, allowing perceptual input to activate associative memory, and vice-versa. Thus associative memory is distinct from the systems that provide it input, and as one would expect, dissociations between impaired access to semantics and intact access to perceptual features, and vice versa, can be found in different sensory modalities (e.g., Schacter, Cooper, Delaney, Peterson, & Tharan, 1991; Warrington & Taylor, 1978). In addition, patients with category-specific associative memory deficits have been described (e.g. Hart, Berndt, & Caramazza, 1985; McCarthy & Warrington, 1986); for example, a patient might be unable to identify pictures of living things, but has no trouble identifying nonliving things. Although such findings may sometimes reflect damage to associative memory per se, in many cases the deficits may reflect disruptions of the pointers from associative memory to perceptual memories. Careful analysis of these deficits and results from neural network models has revealed that such deficits may arise from damage to modality-specific representational systems, rather than damage to a special “living things” memory system (Farah & McClelland, 1991).

Third, it is not clear where in the brain associative memory is implemented. Although the occipital-temporal-parietal area appears to play a critical role in associative memory (e.g., Geschwind, 1965; Kosslyn & Koenig, 1992), deficits in associative memory have been described after lesions to many different brain areas (e.g., Hart et al., 1985; Tulving, Hayman, & MacDonald, 1991). A problem in localizing this subsystem is that activation of areas associated with semantic processing may reflect either the memory structure itself or the processes that access it.

C. Information Look-up Subsystem

When encoding new information into memory or looking up information to help identify an object, generate a mental image, or answer a question, one can use the look-up subsystem to access associative memory. As noted earlier, the frontal lobes play a key role in implementing this subsystem (see Figures 3 and 4B). PET investigations have shown that various tasks that involve accessing semantic information in memory activate the left frontal lobe; such tasks include verb generation (Petersen

et al., 1988), verbal fluency (Frith, Friston, Liddle, & Frackowiak, 1991), completion of nonstudied word stems (Buckner et al., 1996), image generation (Kosslyn, Alpert et al., 1993), and identifying objects seen from unusual views (Kosslyn et al., 1994). Left frontal lesions may also impair short-term semantic priming that depends on the strength of association between word pairs (Milberg & Blumstein, 1981). In addition, accessing semantic information in memory often activates the left anterior cingulate cortex, part of the anterior attention system discussed earlier (e.g., Frith et al., 1991; Kapur et al., 1994; Kosslyn, Daly et al., 1993; Petersen et al., 1988).

In some situations, however, the right-frontal lobe—not the left—is activated when people access information in memory. Retrieval of episode-specific memories for auditory sentences (Tulving, Kapur, Markowitsch et al., 1994), words from three-letter cues (Buckner et al., in press; Squire et al., 1992), faces (Haxby et al., 1993), or scents (Jones-Gotman, Zaforre, Evans, & Meyer, 1993) activate the right frontal lobe. Kosslyn (1994) suggests that different look-up subsystems are implemented in the left and right frontal lobes, which access categorical information (such as that specified by words) and specific information (such as specific events), respectively.

D. Memory Formation Subsystem

When the term *memory* is used in common parlance, it usually refers to memory for specific events (a person who has difficulty encoding or retrieving such memories is often referred to as having a “bad memory”). We can encode various types of new information: new perceptual representations, new associations between items in associative memory, new associations between items in associative memory and perceptual representations, and all of this information is often embedded in a particular spatio-temporal context. How do we flexibly and quickly encode these new memories?

The ability to store new facts in memory depends upon the integrity of the diencephalon, and medial temporal lobe structures that include the mammillary bodies, the hippocampus, and the surrounding perirhinal, entorhinal, and parahippocampal cortices (although findings in monkeys suggest that the contributions to memory of these latter four areas may differ in interesting ways; for discussion, see Gaffan & Murray, 1992; Squire, 1992; Zola-Morgan, Squire, Amaral, & Suzuki, 1989). The hippocampal region and medial dorsal thalamic nucleus, part of the diencephalon, are indicated in Figure 4A. The hippocampal formation receives inputs from a variety of cortical areas in the frontal, parietal, and temporal lobes, and seems ideally positioned to encode relations among cues and context (Squire, 1992). Indeed, hippocampal lesions impair a rat’s ability to learn to navigate in a water maze or eight-arm radial maze, to learn to discriminate among locations containing food, and to acquire conditioned responses that are specific to a spatial context (Jarrard, 1993). Similar lesions impair a monkey’s ability to recall a rewarded object across delays lasting from minutes to days (Squire, 1992).

Results from studies of brain-damaged humans converge with the findings from nonhuman animals. In a series of classic studies of patient H. M., Milner and her colleagues (e.g. Milner, Corkin, & Teuber, 1968; Scoville & Milner, 1957) discovered that bilateral medial temporal lobe removal caused permanent anterograde amnesia (i.e., disrupted memory for new events) and slight retrograde amnesia (i.e., disrupted memory for past events), but spared short-term memory and prior semantic knowledge. After his operation, H. M. could not remember events beyond a few minutes; for example, after more than a few minutes, he would consistently forget ever having met the experimenter. Squire and his colleagues later showed that damage confined exclusively to the CA1 region of the hippocampus causes marked memory impairment, but not as severe as that of H. M. (who had complete removal of the medial temporal area; see Squire, 1992; Squire, Amaral, & Press, 1990; Zola-Morgan, Squire, & Amaral, 1986). Memory deficits have been found following CA1 damage in rats (Auer, Jensen, & Whishaw, 1989) and monkeys (Zola-Morgan & Squire, 1990a).

The memory formation subsystem is needed for normal encoding of information into associative memory: amnesics typically cannot acquire new semantic information (e.g., Rozin, 1976; Squire & Shimamura, 1986), such as word meanings (Gabrieli, Cohen, & Corkin, 1988) or paired associates (D. N. Brooks & Baddeley, 1976). If they do acquire semantic information, it is usually tied to specific aspects of the learning environment, and occurs only after a slow and laborious training (Glisky & Schacter, 1988; MacAndrews et al., 1987; Schacter, Harbluk, & McLachlan, 1984; Shimamura & Squire, 1987, 1988; Tulving et al., 1991). In contrast, damage to the medial temporal lobe does not substantially impair encoding of representations by the perceptual encoding subsystems: amnesic patients show normal priming on both visual object (e.g., Cave & Squire, 1992; Schacter, Cooper, Tharan, & Rubens, 1991) and visual word and nonword priming tasks (word identification, Cermak, Talbot, Chandler, & Wolbarst, 1985; word stem completion, Graf et al., 1984; word fragment completion, Tulving et al., 1991; Warrington & Wieskrantz, 1974; nonwords, Cermak, Verfaellie, Milberg, Letourneau, & Blackford, 1991; Haist, Musen, & Squire, 1991; see, however Cermak et al., 1985; Smith & Oscar-Berman, 1990).

One important aspect of the memory encoding subsystem is that it takes time to complete the encoding process. Damage to the hippocampus causes a temporally graded retrograde amnesia: there is a large loss of memory when the lesions occur soon after learning, but memory loss tapers off as the time between learning and lesion increases (Kim & Fanselow, 1992; MacKinnon & Squire, 1989; Zola-Morgan & Squire, 1990b). Computational models of the hippocampus have provided insights into why such a delay might be necessary for memory encoding. J. L. McClelland, McNaughton and O'Reilly (1994; see also Gluck & Meyers, 1993) argue that typically we do not want the influence of any one learning event to have a large effect on representations in associative memory. Rather, it would be more useful to have the connections among items in an associative memory network

change slowly and gradually as a function of events that recur in time. But we still want to be able to take a quick “snapshot” of the environment, so that relations among cues can be encoded if necessary. The memory encoding subsystem takes this “snapshot,” which allows rapid encoding and orthogonalization (i.e., creation of distinct representations) of memories in a sparse, compact code. This code specifies spatial, temporal, and other contextual variables and can later be used to “train” cortical areas to store a structural representation of the information. This training is slow in order to reduce interference among cortically based memory representations (cf. McCloskey & Cohen, 1989).

This model accounts for findings of retrograde amnesia in humans and animals following damage to the hippocampus: these memory problems arise because recently acquired memory traces have yet to be stored as structural representations in the cortex. This model also suggests why it is so difficult for amnesics to acquire new semantic information: the connections in associative memory change only very slowly in the face of perceptual input without training by the memory formation subsystem. It is also possible that loss of an orthogonalization process underlies some aspects of human amnesia.

E. Stimulus–Motor Response Connection Subsystem

We not only can store representations of facts, but also can store relatively direct connections between stimuli and motor responses. Stimulus–response (SR) learning is akin to behaviorist notions of a direct link between eliciting stimuli and consequent actions without any mediating internal representations (Skinner, 1957). Typically these associations are built over the course of many repeated pairings of a stimulus and a response. A particularly clear example of such learning was provided by Mishkin and his colleagues (Mishkin & Appenzeller, 1987). These researchers presented monkeys with an object discrimination task in which pairs of items were presented once per day for a period of about 4 weeks. Upon presentation of each pair the monkey had to choose the item that was consistently paired with the reward. Not only could monkeys with lesioned hippocampi acquire the correct response, eventually choosing the rewarded object in each pair, but they did so at the same rate as normal monkeys. The memory formation subsystem is not necessary for acquiring this sort of information.

This object discrimination task involves a consistent mapping of a single stimulus onto a response and depends upon integrity of a set of subcortical structures known as the neostriatum (often simply called the striatum). The striatum has two parts, the caudate and the putamen (both of which are parts of the basal ganglia), which have connections to the perceptual encoding systems and motor output systems (see Figure 4C). Lesions of the striatum impair learning of brightness discriminations, avoidance learning, reversal learning, and alternation (McDonald & White, 1993).

It is important to note that tasks that tap the stimulus–motor response subsystem

can be disrupted independently of tasks that tap the memory formation subsystems, and vice versa (this type of pattern of results is called a “double dissociation”; Teuber, 1955). Researchers have contrasted performance in two types of tasks, which superficially may appear similar: In the win–stay task, a rat must learn to return to a single arm of an eight–arm radial maze in order to receive food; in the win–shift task, food is available in any of the eight arms, and a rat must learn to visit each arm only once, noting the spatial location of each arm visited so as not to visit it a second time. Lesions of the stratum impair learning in the win–stay task, but not the win–shift task, and fornix lesions (disrupting input to the hippocampus) have the opposite effect (Packard, Hirsh, & White, 1989; Packard & McGaugh, 1992). Thus, it appears that the striatum is necessary for coding consistent S–R mappings built up over time.

The stimulus–motor response subsystem may also be involved in the control of sequences of successive S–R mappings. Striatal lesions can disrupt production of sequences of rat instinctual grooming behaviors, although the individual constituent movements can still be elicited (Berridge & Whishaw, 1992). Furthermore, neurophysiological recordings of activity in the striatum have revealed that such neurons are active only when the animal produces grooming sequences (Aldridge, Berridge, Herman, & Zimmer, 1993).

Similarly, humans with Parkinson’s disease (which reflects impaired functioning of the striatum, due to a depletion of the neurotransmitter dopamine) also have deficits in sequential processing, such as the timing of vocal utterances and syntactic comprehension (Lieberman et al., 1992). In addition, patients with Huntington’s disease (caused by degeneration of the striatum) may show deficits in a variety of tasks that require either acquisition of simple S–R mappings or acquisition of sequences of such mappings. In contrast, amnesics do not show such impairments. For example, in the serial reaction time (SRT) task subjects press keys in response to visual cues that appear in one of four locations. Cues may appear at these locations in either a random order or in a repeating sequence (usually 10 items long). Learning is indexed by a decrease in reaction time as more trials are completed, and acquisition of the sequence is shown by greater improvement in repeating blocks as compared to random blocks of trials. Amnesics acquire the sequence normally, but patients with Huntington’s disease fail to show this learning (e.g., Knopman & Nissen, 1991). A similar dissociation is revealed by a weight judgment task, in which prior exposure to a set of weights biases subsequent judgments of them: amnesics show normal biases, whereas patients with Huntington’s disease do not (Heindel, Salmon & Butters, 1991). Moreover, patients with Huntington’s disease are also impaired on the pursuit rotor task in which subjects must hold a stylus on a disk located near the edge of a rapidly rotating platter (Heindel, Salmon, Shults, Walicke, & Butters, 1989). Abnormal metabolic activity in the striatum has been linked to psychiatric syndromes involving repetitive thoughts or actions, such as obsessive–compulsive disorder (Rauch et al., 1997) and Tourette’s syndrome (Witelson, 1994).

F. Stimulus-Visceral Response Subsystem

In some situations it is necessary to form an association between an external, neutral stimulus and an internal state evoked by an event with negative consequences. This type of learning is different from that involved in S-R learning because the association is formed between a stimulus and a physiological state resulting from a stimulus, rather than between a stimulus and an overt motor response. For example, in the fear-conditioning paradigm a rat is shocked following presentation of a light or tone (J. LeDoux & Hirst, 1986). Over time, the rat comes to exhibit fearful behavior to the light alone, as evidenced by changes in sympathetic and parasympathetic nervous system activity when the light is presented. Lesions of the striatum do not impair conditioning in this paradigm, whereas lesions to the central and lateral nuclei of the amygdala do (see Figure 4A for location of this subsystem). These two amygdaloid nuclei have proven crucial for learning to occur in a variety of tasks that involve association of an aversive stimulus and a neutral stimulus, such as fear-potentiated startle (Davis, 1992), passive avoidance (Cahill & McGaugh, 1990), and conditioning of autonomic responses such as heart rate or blood pressure (Kapp, Whalen, Supple, & Pascoe, 1992). What is common across all these tasks, even though the motor responses may differ, is the link between a stimulus and a visceral, internal state.

This subsystem is also important for acquiring associations between stimuli and appetitive events, though different amygdaloid nuclei may be involved (basolateral and lateral; Everitt et al., 1992; McDonald & White, 1993). Thus, damage to the amygdala may impair acquisition of a variety of conditioned reward tasks in which a neutral stimulus is paired with reward. McDonald and White (1993) showed that rats with lesions of the amygdala, but not the striatum or fornix, were impaired in learning a task in which rats were allowed to feed in different, although perceptually similar, locations whenever a light was present. Other rats were fed only in dark areas. Learning was assessed by the amount of time spent in the lighted or darkened area that had been associated with food. In this task, the only memory that could underlie the animals' preference was the association of the cue and the internal state generated by food. The amygdala lesions did not impair learning of the win-stay or win-shift tasks, although performance in these tasks is impaired following striatal and hippocampal damage, respectively.

G. Summary and Extensions

Like other complex mental abilities, memory is subserved by a host of specialized subsystems. Perceptual encoding subsystems represent modality-specific inputs at a presemantic level; associative memory stores relational, identifying, and classifying information in a propositional format; information look-up subsystems access information in associative memory; the memory formation subsystem enables flexible, rapid encoding of episodes and events; and the stimulus-motor response and

stimulus-visceral response subsystems encode pairings of stimuli and behaviors or stimuli and physiological states.

This conceptualization of memory not only allows one to account for a wide range of experimental findings, but also leads to predictions that follow from specific interactions among different subsystems. For example, consider a task that is impaired following hippocampal damage, but may involve the stimulus-motor response subsystem as well. In the negative patterning paradigm, an animal is rewarded if it presses a bar when a tone is presented or when a light is on. However, the animal is not rewarded if it presses when the tone and light appear simultaneously. Normal animals learn not to respond when both cues are present. Hippocampal lesioned animals, however, cannot withhold responses to the tone-light pairing, although they respond normally to each of the stimuli in isolation (McDonald and White, 1993). Presumably, the intact stimulus-motor response system mediates responses to the individual stimuli, but the hippocampus is necessary to encode the association between the two simultaneous cues and the lack of reward. If this account is correct, then animals with striatal lesions should also show impairments on this task, even when only a single stimulus is present. Future work may address this prediction.

Similar analyses may inform and motivate research with human subjects. Most of the work on the learning of S-R sequences in the SRT in humans has employed sequences in which responses predict each other with unequal probability. For example, key 1 might be followed by key 2 with .67 probability, by key 3 with .33 probability, and by key 4 with .00 probability. Having some responses predict the occurrence of others may reduce the number of S-R mappings that must be acquired. These probabilistically *unbalanced* sequences are the type that patients with striatal damage have been shown to be unable to acquire (Knopman & Nissen, 1991), which has led to the conclusion that the striatum alone may participate in sequence learning and performance.

However, some researchers suggest that the memory formation subsystem may participate in some forms of sequence learning that cannot be learned on the basis of predictive S-R chains but require hierarchical grouping of response sequence clusters (Keele & Curran, 1995; but see also Keele et al., 1998). It is possible that a stimulus-motor response subsystem may allow expression of only a few simple S-R chains, and when more complex mappings are required the memory formation subsystem is recruited as necessary (cf. Squire & Frambach, 1990).

Finally, consider an example in which a cognitive neuroscientific view of memory can lead one to infer properties of previously unstudied subsystems, which in turn may lead to novel results that can be explained with reference back to the systems that generated the initial hypotheses. This process is being played out in investigations of the auditory perceptual encoding subsystem. Given that there are modality-specific, cortically based systems that represent the structure and form of objects and words at a "presemantic" level (i.e., the level of recognition, as opposed to identification), we expect similar systems to exist in various sensory modalities.

Indeed, neuropsychological and PET research has shown that encoding phonological information involves the posterior superior temporal lobe (Ellis & Young, 1988; Petersen et al., 1988).

Such findings led Schacter and Church (1992) to infer the existence of an auditory word form subsystem that is dedicated to representing the acoustic, but not the semantic, properties of spoken words. Support for this claim comes from studies of priming on tests in which subjects identify perceptually degraded spoken words that have had their low frequencies removed, and in tests of auditory stem completion, in which subjects complete an auditory stem to form the first word that comes to mind. Consistent with the notion that auditory priming is modality-specific and presemantic, word identification priming is reduced when the presentation modality is changed from the initial exposure to the test phases of the experiment (A. G. R. McClelland & Pring, 1991); moreover, such priming is not affected by semantic encoding tasks that enhance explicit memory (Schacter & Church, 1992). Such priming is also specific to the acoustic properties of the input. Church and Schacter (1994) found that changing the emotional tone, gender, or fundamental frequency of a speaker's voice from exposure to test phases reduces priming. Furthermore, auditory priming is preserved in amnesia (Schacter, Church, & Treadwell, 1994) as well as in patients who suffer from word meaning deafness (Schacter, McGlynn, Milberg, & Church, 1993).

V. CONCLUSIONS

In this chapter we have tried to illustrate the ways in which cognitive neuroscientific analyses make use of multiple, converging streams of evidence to inform theory construction. There are five general points about cognitive functions revealed by this analysis:

1. Many of the subsystems that confer a specific ability such as attention, vision, or memory, can interact in multiple ways.
2. These systems consist of networks that are implemented in distinct brain areas.
3. Each area carries out computations that can be characterized specifically enough to be implemented in a simulation model.
4. Each system processes information both serially and in parallel.
5. Processing is highly interactive, with higher-level areas sending feedback to lower level areas.

At the present stage of research, different types of data carry more or less weight in different domains. For instance, in the study of visual perception, theorizing rests in large part on the results of studies of monkeys. Only recently have neuroimaging studies begun to confirm and extend some of the basic findings from the animal literature; theories have also gained leverage by attempting to explain the effects of focal brain damage in humans. In the case of attention, recent neuroimaging

findings are beginning to gather force as well, complementing early work with brain-damaged patients. And in the study of memory, the multiple systems account offered here is guided primarily by work with human subjects, using data from animals to help address specific questions.

There are perhaps two major reasons for the differences in approaches among the fields. First, the study of each topic is strongly influenced by the first disciplines to make significant contributions to theory in that area. Second, and this is especially true for study of various forms of attention and memory, we currently do not have the techniques to study some human abilities, such as priming, in animal populations. But the day is still young. The term *cognitive neuroscience* was only coined in 1970 (Kosslyn & Andersen, 1993) and the Cognitive Neuroscience Society had its inaugural meeting in 1994. Despite its relative infancy, the field is making steady progress in many areas, and the rapid development and increased availability of new imaging techniques will help to address the functional anatomy of abilities.

This chapter has surveyed theoretical advances in only a handful of areas that currently are being explored from a cognitive neuroscience perspective. Emotion, language, categorization and reasoning, movement, and audition are but a few of the topic domains in which theories of this sort are now being advanced (Gazzaniga, 1995; Kosslyn & Koenig, 1992; LeDoux & Hirst, 1986; Ochsner & Schacter, in press; Weingartner & Lister, 1991). Current work is extending the cognitive neuroscience approach to problems of interest to social psychologists, such as attitude change (Lieberman, in press; Lieberman et al., 1999; Ochsner & Lieberman, 1999). We anticipate that future work will only broaden the horizons of these exciting research programs.

Acknowledgments

Completion of this article was supported by a National Science Foundation Graduate Research Fellowship awarded to K. N. Ochsner, and National Institutes of Health grant NINDS 2 PO1 17778-09 and AFOSR grant F49620-99-1-0114 awarded to S. M. Kosslyn. We thank Daniel L. Schacter and Tim Curran for helpful discussion of relevant issues.

References

- Aldridge, J. W., Berridge, K. C., Herman, M., & Zimmer, L. (1993). Neuronal coding of serial order: Syntax of grooming in the neostriatum. *Psychological Science*, *4*, 391–395.
- Alexander, G. E., Crutcher, M. D., & DeLong, M. A. (1990). Basal ganglia-thalamocortical circuits: Parallel substrates for motor, oculomotor, “prefrontal” and “limbic” functions. In H. B. M. Uylings, C. G. Van Eden, J. P. C. De Bruin, M. A. Corner, & M. G. P. Feenstra (Eds.), *Progress in brain research* (pp. 119–146). Amsterdam: Elsevier.
- Allport, A. (1992). Attention and control: Have we been asking the wrong questions? A critical review of twenty five years. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV* (pp. 187–213). Cambridge, MA: MIT Press.
- Amaral, D. G., Price, J. L., Pitkanen, A., & Carmichael, S. T. (1992). Anatomical organization of the pri-

- mate amygdaloid complex. In J. P. Aggleton (Ed.), *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 1–66). New York: Willey Liss.
- Ambros-Ingerson, J., Granger, R., & Lynch, G. (1990). Simulation of paleocortex performs hierarchical clustering. *Science*, *247*, 1344–1348.
- Andersen, R. A. (1987). Inferior parietal lobule function in spatial perception and visuomotor integration. In F. Plum (Ed.), *Handbook of physiology* (Sect. 1, Vol. 5, pp. 483–518). Bethesda, MD: American Physiological Society.
- Andersen, R. A., Essick, G. K., & Siegel, R. M. (1985). Encoding of spatial location by posterior parietal neurons. *Science*, *230*, 456–458.
- Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review*, *85*, 249–277.
- Aston-Jones, G., Foote, S. L., & Bloom, F. E. (1984). Anatomy and physiology of locus coeruleus neurons: Functional implications. In M. G. Ziegler (Ed.), *Frontiers in clinical neuroscience*. Baltimore, MD: Williams & Watkins.
- Auer, R. N., Jensen, M. L., & Whishaw, I. Q. (1989). Neurobehavioral deficit due to ischemic brain damage limited to half of the CA1 section of the hippocampus. *Journal of Neuroscience*, *9*, 1641–1647.
- Ballantine, H. T., Cassidy, W. C., Brodeur, J., & Giriunas, I. B. (1972). Frontal cingulotomy for mood disturbance. In E. Hitchcock, L. Laitinen, & K. Vearnat (Eds.), *Psychosurgery* (pp. 221–228). Springfield, IL: Charles C. Thomas.
- Barbas, H., & Pandya, D. N. (1981). Frontal lobe afferent input to area 6 in the rhesus monkey. *Society for Neuroscience Abstracts*, *7*, 414.
- Berridge, K. C., & Whishaw, I. Q. (1992). Cortex, striatum and cerebellum: Control of serial order in a grooming sequence. *Experimental Brain Research*, *90*, 275–290.
- Biederman, I., & Cooper, E. E. (1992). Size invariance in visual object priming. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 121–133.
- Bisiach, E., & Luzzatti, C. (1978). Unilateral neglect of representational space. *Cortex*, *14*, 129–133.
- Bisiach, E., Luzzatti, C., & Perani, D. (1979). Unilateral neglect, representational scheme, and consciousness. *Brain*, *102*, 609–618.
- Broadbent, D. E. (1971). *Decision and stress*. London: Academic Press.
- Broca, P. (1863). Localisations des fonctions cérébrales. Siège du langage articulé. *Bulletins de la Société d'Anthropologie*, *4*, 200–204.
- Brooks, D. N., & Baddeley, A. D. (1976). What can amnesic patients learn? *Neuropsychologia*, *14*, 111–122.
- Brooks, V. B. (1986). How does the limbic system assist motor learning? A limbic comparator hypothesis. *Brain, Behaviour and Evolution*, *29*, 29–53.
- Buckner, R. L., Petersen, S. E., Ojemann, J. G., Miezin, F. M., Squire, L. R., & Raichle, M. E. (1996). Functional anatomical studies of explicit and implicit memory retrieval tasks. *Journal of Neuroscience*, *16*, 6219–6235.
- Cahill, L., & McGaugh, J. L. (1990). Amygdaloid complex lesions differentially affect retention of tasks using appetitive and aversive reinforcement. *Behavioral Neuroscience*, *104*(4), 532–543.
- Caplan, D. (1993). *Language*. Cambridge, MA: MIT Press.
- Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, *280*, 747–749.
- Cave, C. B., & Squire, L. R. (1992). Intact and long-lasting repetition priming in amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 509–520.
- Cermak, L. S., Talbot, N., Chandler, K., & Wolbarst, L. R. (1985). The perceptual priming phenomenon in amnesia. *Neuropsychologia*, *23*, 615–622.
- Cermak, L. S., Verfaellie, M., Milberg, W., Letourneau, L., & Blackford, S. (1991). A further analysis of perceptual identification priming in alcoholic Korsakoff patients. *Neuropsychologia*, *29*(8), 725–736.
- Church, B. A., & Schacter, D. L. (1994). Perceptual specificity of auditory priming: implicit memory for voice intonation and fundamental frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 521–533.

- Churchland, P. S., & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, MA: MIT Press.
- Clark, C. R., Geffen, G. M., Geffen, L. B. (1989). Catecholamines and the covert orientation of attention in humans. *Neuropsychologia*, 27(2), 131–139.
- Cohen, N. J. (1984). Preserved learning in amnesia: Evidence for multiple memory systems. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of memory* (pp. 83–103). New York: Guilford Press.
- Cohen, N. J., & Eichenbaum, H. (1994). *Memory, amnesia and the hippocampal system*. Cambridge, MA: MIT Press.
- Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern analyzing skill in amnesics: Dissociation of knowing how and knowing that. *Science*, 210, 207–210.
- Cook, N. D., Frueth, H., & Landis, T. (1995). The cerebral hemispheres and network simulations: Design considerations. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 410–422.
- Cooper, L. A., & Schacter, D. L. (1992). Dissociations between structural and episodic representations of visual objects. *Current Directions in Psychological Science*, 1, 141–146.
- Cooper, L. A., Schacter, D. L., Ballesteros, S., & Moore, C. (1992). Priming and recognition of transformed three-dimensional objects: Effects of size and reflection. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 43–57.
- Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: Functional anatomy by positron emission tomography. *Journal of Neuroscience*, 11(8), 2383–2402.
- Corbetta, M., Miezen, F. M., Schulman, G. L., & Petersen, S. E. (1993). A PET Study of visuospatial attention. *Journal of Neuroscience*, 13, 1202–1226.
- Coslett, H. B., Bowers, D., & Heilman, K. M. (1987). Reduction in cerebral activation after right hemisphere stroke. *Neurology*, 37, 957–962.
- Davis, M. (1992). The role of the amygdala in conditioned fear. In J. P. Aggleton (Ed.), *The amygdala: Neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 255–306). New York: Wiley-Liss.
- Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. J. (1984). Stimulus selective properties of inferior temporal neurons in the macaque. *Journal of Neuroscience*, 4, 2051–2062.
- DeYoe, E. A., & Van Essen, D. C. (1988). Concurrent processing streams in monkey visual cortex. *Trends in Neurosciences*, 11, 219–226.
- Dieber, M. P., Passingham, R. E., Colebatch, J. G., Friston, K. J., Nixon, P. D., & Frackowiak, R. S. J. (1991). Cortical areas and the selection of movement: A study with PET. *Experimental Brain Research*, 84, 393–402.
- Ellis, A. W., & Young, A. W. (1988). *Human cognitive neuropsychology*. Hove, UK: Erlbaum.
- Evarts, E. V. (1966). Methods for recording activity of individual neurons in moving animals. In R. F. Ruhmer (Ed.), *Methods in medical research* (pp. 241–250). Chicago: Chicago Year Book.
- Everitt, B. J., & Robbins, T. W. (1992). Amygdala-ventral striatal interactions and reward-related processes. In J. P. Aggleton (Ed.), *The amygdala: Neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 401–429). New York: Wiley-Liss.
- Farah, M. J. (1990). *Visual agnosia: Disorders of object recognition and what they tell us about normal vision*. Cambridge, MA: MIT Press.
- Farah, M. J. (1994). Neuropsychological inference with an interactive brain: A critique of the “locality” assumption. *Behavioral and Brain Sciences*, 17, 43–104.
- Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality-specificity and emergent category-specificity. *Journal of Experimental Psychology: General*, 120, 339–357.
- Feldman, J. A. (1985). Four frames suffice: A provisional model of vision and space. *Behavioral and Brain Sciences*, 8, 265–289.
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.
- Fox, P. T., Mintun, M. A., Raichle, M. E., Miezen, F. M., Allman, J. M., & Van Essen, D. C. (1986). Mapping human visual cortex with positron emission tomography. *Nature (London)*, 323, 806–809.

- Frith, C. D., Friston, K. J., Liddle, P. F., & Frackowiak, R. S. J. (1991). A PET study of word finding. *Neuropsychologia*, *29*, 1137–1148.
- Fujita, I., Tanaka, K., Ito, M., & Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. *Nature (London)*, *360*, 343–346.
- Gabriel, M. (1990). Functions of anterior and posterior cingulate cortex during avoidance learning in rabbits. In H. B. M. Uylings, J. P. C. Van Eden, M. A. De Bruin, & M. G. P. Feenstra (Eds.), *Progress in brain research* (pp. 467–483). Amsterdam: Elsevier.
- Gabrieli, J. D. E., Cohen, N. J., & Corkin, S. (1988). The impaired learning of semantic information following bilateral temporal lobe resection. *Brain and Cognition*, *7*, 157–177.
- Gaffan, D., & Murray, E. A. (1992). Monkeys (*Macaca fascicularis*) with rhinal cortical ablations succeed in object discrimination learning despite 24-hr intertrial intervals and fail at matching to sample despite double presentations. *Behavioral Neuroscience*, *106*, 30–38.
- Gazzaniga, M. S. (Ed.). (1995). *The cognitive neurosciences*. Cambridge, MA: MIT Press.
- Gehring, W. J., Goss, B., Coles, M. G., Meyer, D. E., & Donchin, E. (1993). A neural system for error detection and compensation. *Psychological Science*, *4*(6), 385–390.
- Georgopoulos, A. P. (1994). Motor cortex and cognitive processing. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 507–517). Cambridge, MA: MIT Press.
- Geschwind, N. (1965). Disconnexion syndromes in animals and man. *Brain*, *88*, 237–294.
- Glisky, E. L., & Schacter, D. L. (1988). Long-term retention of computer learning by patients with memory disorders. *Neuropsychologia*, *26*, 173–178.
- Gluck, M. A., & Meyers, C. E. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus*, *3*, 491–516.
- Goldman-Rakic, P. S. (1988). Topography of cognition: Parallel distributed networks in primate association cortex. *Annual Review of Neuroscience*, *11*, 137–156.
- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist paradigm. *Proceedings of the Royal Society of London, Series B*, *205*, 581–598.
- Graf, P., & Ryan, L. (1990). Transfer-appropriate processing for implicit and explicit memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 978–992.
- Graf, P., Squire, L. R., & Mandler, G. (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 164–178.
- Gregory, R. L. (1970). *The intelligent eye*. London: Weidenfeld & Nicholson.
- Gross, C. G., Desimone, R., Albright, T. D., & Schwartz, E. L. (1984). Inferior temporal cortex as a visual integration area. In F. Reinoso-Suarez & C. Ajmone-Marsan (Eds.), *Cortical integration*. New York: Raven Press.
- Gross, C. G., & Mishkin, M. (1977). The neural basis of stimulus equivalence across retinal translation. In S. Harnad, R. Doty, J. Jaynes, L. Goldstein, & G. Krauthamer (Eds.), *Lateralization in the nervous system*. New York: Academic Press.
- Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, *87*, 1–51.
- Haier, R. J., Siegel, B. V., Neuchterlein, K. H., Hazlett, E., Wu, J. C., Paek, J., Browning, H. L., & Buchsbaum, M. S. (1988). Cortical glucose metabolic rate correlates of abstract reasoning and attention studied with positron emission tomography. *Intelligence*, *12*, 199–217.
- Haist, F., Musen, G., & Squire, L. R. (1991). Intact priming of words and nonwords in amnesia. *Psychobiology*, *19*, 275–285.
- Hart, J., Berndt, R. S., & Caramazza, A. (1985). Category-specific naming deficit following cerebral infarction. *Nature (London)*, *316*, 439–440.
- Hasselmo, M. E. (1993). Acetylcholine and learning in a cortical associative memory. *Neural Computation*, *5*, 32–44.
- Hasselmo, M. E., & Bower, J. M. (1993). Acetylcholine and memory. *Trends in Neurosciences*, *16*, 218–222.
- Haxby, J. V., Horwitz, B., Maisog, J. M., Ungerleider, L. G., Mishkin, M., Schapiro, M. B., Rapoport, S. I., & Grady, C. L. (1993). *Journal of Cerebral Blood Flow and Metabolism*, *13*(Suppl. 1), S499.

- Hayman, C. A. G., & Tulving, E. (1989). Is priming in fragment completion based on "traceless" memory system? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 941–956.
- Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.
- Heilman, K., & Valenstein, K. M. (1985). *Clinical neuropsychology*. New York: Oxford University Press.
- Heindel, W. C., Salmon, D. P., & Butters, N. (1991). The biasing of weight judgments in Alzheimer's and Huntington's disease: A priming or programming phenomenon? *Journal of Clinical and Experimental Neuropsychology*, *13*, 189–203.
- Heindel, W. C., Salmon, D. P., Shuits, C. W., Walicke, P. A., & Butters, N. (1989). Neuropsychological evidence for multiple implicit memory systems: A comparison of Alzheimer's, Huntington's, and Parkinson's disease patients. *Journal of Neuroscience*, *9*, 582–587.
- Hellige, J. B., & Michimata, C. (1989). Categorization versus distance: Hemispheric differences for processing spatial information. *Memory & Cognition*, *17*, 770–776.
- Holmes, G. (1919). Disturbances of visual space perception. *British Medical Journal*, *2*, 449–468, 506–615.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London)*, *195*, 215–243.
- Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, *99*(3), 480–517.
- Hyvarinen, J. (1982). Posterior parietal lobe of the primate brain. *Physiological Review*, *62*, 1060–1129.
- Jackson, S., & Houghton, G. (1995). Sensorimotor selection and the basal ganglia: A neural network model. In J. Howe (Eds.), *Models of information processing in the basal ganglia* (pp. 336–367). Cambridge, MA: MIT Press.
- Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General*, *110*, 306–340.
- Jacoby, L. L., & Hayman, C. A. G. (1987). Specific visual transfer in word identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 456–463.
- Janer, K. W., & Pardo, J. V. (1991). Deficits in selective attention following bilateral anterior cingulotomy. *Journal of Cognitive Neuroscience*, *3*(2), 231–241.
- Jarrard, L. E. (1993). On the role of the hippocampus in learning and memory in the rat. *Behavioral and Neural Biology*, *60*, 9–26.
- Jones-Gotman, M., Zatorre, R. J., Evans, A. C., & Meyer, E. (1993). *Society for Neuroscience Abstracts*, *19*, 1002.
- Kapp, B. S., Whalen, P. J., Supple, W. F., & Pascoe, J. P. (1992). Amygdaloid contributions to conditioned arousal and sensory information processing. In J. P. Aggleton (Eds.), *The amygdala: Neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 229–254). New York: Wiley-Liss.
- Keele, S. W., & Jennings, P. (1992). Attention in the representation of a sequence: Experiment and theory. *Human Movement Science*, *11*, 125–138.
- Keele, S. W., & Curran, T. (1995). On the modularity of sequence learning systems in humans. In E. Covey (Ed.), *Neural representation of temporal patterns* (pp. 197–225). New York: Plenum Press.
- Keele, S. W., Ivry, R. B., Hazeltine, E., Mayr, U., & Heuer, H. (1998). *The cognitive and neural architecture of sequence representation*. (Technical Report 98-03 REVISED). University of Oregon: Institute of Cognitive and Decision Sciences.
- Kim, J. J., & Fanselow, M. S. (1992). Modality-specific retrograde amnesia of fear. *Science*, *256*, 675–677.
- Knopman, D., & Nissen, M. J. (1991). Procedural learning is impaired in Huntington's disease: Evidence from the serial reaction time task. *Neuropsychologia*, *29*, 245–254.
- Kohn, S. E., & Friedman, R. B. (1986). Word-meaning deafness: A phonological-semantic dissociation. *Cognitive Neuropsychology*, *3*, 291–308.
- Kosslyn, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- Kosslyn, S. M. (1987). Seeing and imagining in the cerebral hemispheres: A computational approach. *Psychological Review*, *94*, 148–175.
- Kosslyn, S. M., Koenig, O., Barrett, A., Cave, C. B., Tang, J., & Gabrieli, J. D. E. (1989). Evidence for

- two types of spatial representations: Hemispheric specialization for categorical and coordinate relations. *Journal of Experimental Psychology: Human Perception & Performance*, *15*, 723–735.
- Kosslyn, S. M. (1994). *Image and brain*. Cambridge: MIT Press.
- Kosslyn, S. M., Alpert, N. A., Thompson, W. L., Chabris, C. F., Rauch, S. L., & Anderson, A. K. (1994). Identifying objects seen from canonical and noncanonical viewpoints: A PET investigation. *Brain*, *117*, 1055–1071.
- Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Maljkovic, V., Weise, S. B., Chabris, C. F., Hamilton, S. E., Rauch, S. L., & Buonanno, F. S. (1993). Visual mental imagery activates primary visual cortex: PET investigations. *Journal of Cognitive Neuroscience*, *5*, 263–287.
- Kosslyn, S. M., & Andersen, R. A. (1993). General introduction. In S. M. Kosslyn & R. A. Andersen (Eds.), *Frontiers in cognitive neuroscience* (pp. xv–xxix). Cambridge, MA: MIT Press.
- Kosslyn, S. M., Chabris, C. F., Jacobs, R., Marsolek, C., & Koenig, O. (1995). [reply to Cook, Fruh & Landis]. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 423–431.
- Kosslyn, S. M., Chabris, C. F., Marsolek, C. J., & Koenig, O. (1992). Categorical versus coordinate spatial representations: Computational analyses and computer simulations. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 562–577.
- Kosslyn, S. M., Daly, P. F., McPeck, R. M., Alpert, N. M., Kennedy, D. N., & Caviness, V. S. J. (1993). Using locations to store shape: An indirect effect of a lesion. *Cerebral Cortex*, *3*, 567–582.
- Kosslyn, S. M., & Koenig, O. (1992). *Wet mind*. New York: Free Press.
- Kosslyn, S. M., & Maljkovic, V. M. (1990). Marr's metatheory revisited. *Concepts in Neuroscience*, *1*, 239–251.
- Kosslyn, S. M., & Ochsner, K. N. (1994). In search of occipital activation during visual mental imagery. *Trends in Neurosciences*, *17*, 288–290.
- Kosslyn, S. M., Thompson, W. L., & Alpert, N. M. (1997). Neural systems shared by visual imagery and visual perception: A positron emission tomography study. *NeuroImage*, *6*, 320–334.
- LaBerge, D. (1990). Thalamic and cortical mechanisms of attention suggested by recent positron emission tomographic experiments. *Journal of Cognitive Neuroscience*, *2*(4), 358–372.
- LaBerge, D., & Buchsbaum, M. S. (1990). Positron emission tomography measurements of pulvinar activity during an attention task. *Journal of Neuroscience*, *10*, 613–619.
- Laeng, B. (1994). Localization of categorical and coordinate spatial functions: A study of unilateral stroke patients. *Journal of Cognitive Neuroscience*, *6*, 189–203.
- Lane, R. D., Reiman, E. M., Axelrod, B., Yun, L.-S., Holmes, A., & Schwartz, G. E. (1998). Neural correlates of levels of emotional awareness: Evidence of an interaction between emotion and attention in the anterior cingulate cortex. *Journal of Cognitive Neuroscience*, *10*, 525–535.
- LeDoux, J. E. (1994). In search of an emotional system in the brain: Leaping from fear to emotion and consciousness. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 1049–1061). Cambridge, MA: MIT Press.
- LeDoux, J. E., & Hirst, P. (Eds.). (1986). *Mind and brain: Dialogues in cognitive neuroscience*. New York: Cambridge University Press.
- Levine, D. N. (1982). Visual agnosia in monkey and man. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 629–670). Cambridge, MA: MIT Press.
- Lieberman, P., Kako, E., Friedman, J., Tajchman, G., Feldman, L. S., & Jimenez, E. B. (1992). Speech production, syntax comprehension, and cognitive deficits in Parkinson's disease. *Brain and Language*, *43*, 169–189.
- Lieberman, M. D. (in press). Intuition: A social cognitive neuroscience approach. *Psychological Bulletin*.
- Lieberman, M. D., Ochsner, K. N., Gilbert, D. T., & Schacter, D. L. (1999). *Attitude change: A social cognitive neuroscience approach*. Cambridge, MA: Harvard University Manuscript.
- Luria, A. R. (1980). *Higher cortical functions in man*. New York: Basic Books.
- MacAndrews, M. P., Glisky, E. L., & Schacter, D. L. (1987). When priming persists: Long-lasting implicit memory for a single episode in amnesic patients. *Neuropsychologia*, *25*, 497–506.
- MacKinnon, D., & Squire, L. R. (1989). Autobiographical memory in amnesia. *Psychobiology*, *17*, 247–256.

- Mangun, G. A., Hilyard, S. A., & Luck, S. J. (1992). Electrocortical substrates of visual selective attention. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 219–244). Cambridge, MA: MIT Press.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Freeman.
- Marsolek, C. J., Kosslyn, S. M., & Squire, L. R. (1992). Form-specific visual priming in the right cerebral hemisphere. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 492–508.
- Marsolek, C. J. (1995). Abstract visual-form representations in the left cerebral hemisphere. *Journal of Experimental Psychology: Human Perception & Performance*, *21*, 375–386.
- Maunsell, J. H. R., & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience*, *10*, 363–401.
- McCarthy, R. A., & Warrington, E. K. (1986). Visual associative agnosia: A clinico-anatomical study of a single case. *Journal of Neurology, Neurosurgery, & Psychiatry*, *49*, 1233–1240.
- McClelland, A. G. R., & Pring, L. (1991). An investigation of cross-modality effects in implicit and explicit memory. *Quarterly Journal of Experimental Psychology*, *43A*, 19–33.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1994). *Why we have complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory* (Tech. Rep. No. PDP.CNS.94.1). Pittsburgh, PA; Carnegie-Mellon University.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, *88*, 375–407.
- McClelland, J. L., & Rumelhart, D. E. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2. Psychological and biological models*. Cambridge, MA: MIT Press.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (Ed.), *The psychology of learning and motivation*. Orlando, FL: Academic Press.
- McCollum, J., Larson, J., Otto, T., Schottler, Granger, R., & Lynch, G. (1991). Single unit processing in olfactory cortex. *Journal of Cognitive Neuroscience*, *3*, 293–299.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas imminent in nervous activity. *Bulletin of Mathematical Biophysics*, *5*, 115–133.
- McDonald, R. J., & White, N. M. (1993). A triple dissociation of memory systems: Hippocampus, amygdala, and dorsal striatum. *Behavioral Neuroscience*, *107*, 3–22.
- Milberg, W., & Blumstein, S. E. (1981). Lexical decision and aphasia: Evidence for semantic processing. *Brain and Language*, *14*, 371–385.
- Milner, B. (1964). Some effects of frontal lobectomy in man. In J. M. Warren & K. Akert (Eds.), *The frontal granular cortex and behavior*. New York: McGraw-Hill.
- Milner, B., Corkin, S., & Teuber, H. L. (1968). Further analysis of the hippocampal amnesic syndrome: Fourteen year follow-up study of H.M. *Neuropsychologia*, *6*, 215–234.
- Mishkin, M., & Appenzeller, T. (1987). The anatomy of memory. *Scientific American*, *256*, 80–89.
- Mitchell, D. B., & Brown, A. S. (1988). Persistent repetition priming in picture naming and its dissociation from recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 213–222.
- Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, *229*, 782–784.
- Morrison, J. H., & Foote, S. L. (1986). Noradrenergic and serotonergic innervation of cortical, thalamic and tectal structures in old world and new world monkeys. *Journal of Comparative Neurology*, *243*, 117–128.
- Neisser, U. (1967). *Cognitive psychology*. New York: Appleton-Century-Crofts.
- Nordahl, T. E., Benkelfat, C., Semple, W. E., & Gross, M. (1989). Cerebral glucose metabolic rates in obsessive compulsive disorder. *Neuropsychopharmacology*, *2*, 23–28.
- Ochsner, K. N., & Baker, D. P. (1994). *An integrative neural systems approach to motivation and emotion*. Unpublished manuscript, Harvard University, Cambridge, MA.

- Ochsner, K. N., & Feldman-Barrett, L. (in press). The neuroscience of emotion. In T. Mayne & G. Bonnano (Eds.), *Emotion: Current issues and future directions*. Oxford: Oxford University Press.
- Ochsner, K. N., & Lieberman, M. D. (1999). *The social cognitive neuroscience approach*. Cambridge, MA: Harvard University Manuscript.
- Ochsner, K. N., & Schacter, D. L. (in press). Constructing the emotional past: A social cognitive neuroscience approach to emotion and memory. In J. Borod (Ed.), *The neuropsychology of emotion*. Oxford: Oxford University Press.
- Packard, M. G., Hirsh, R., & White, N. M. (1989). Differential effects of fornix and caudate lesions on two radial maze tasks: Evidence for multiple memory systems. *Journal of Neuroscience*, *9*, 1465–1472.
- Packard, M. G., & McGaugh, J. L. (1992). Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience*, *106*, 439–446.
- Pardo, J. V., Pardo, P. J., Janer, K. W., & Raichle, M. E. (1990). The anterior cingulate cortex mediates processing selection in the Stroop attentional conflict paradigm. *Proceedings of the National Academy of Sciences of the U.S.A.*, *87*, 256–259.
- Perrett, D. I., Smith, P. A. J., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A. D., & Jeeves, M. A. (1985). Visual cells in the temporal cortex sensitive to face view and gaze direction. *Proceedings of the Royal Society of London, Series B*, *223*, 293–317.
- Petersen, S. E., & Fiez, J. A. (1993). The processing of single words studied with positron emission tomography. *Annual Review of Neuroscience*, *16*, 509–530.
- Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M., & Raichle, M. E. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. *Nature (London)* *331*, 585–589.
- Petersen, S. E., Fox, P. T., Snyder, A. Z., & Raichle, M. E. (1990). Activation of extrastriate and frontal cortical areas by visual words and word-like stimuli. *Science*, *249*, 1041–1044.
- Petersen, S. E., Robinson, D. L., & Morris, J. D. (1987). *Neuropsychologia*, *25*, 97–105.
- Posner, M. I. (1978). *Chronometric explorations of mind*. Hillsdale, NJ: Erlbaum.
- Posner, M. I., & Carr, T. H. (1992). Lexical processing and the brain: Anatomical constraints on cognitive models of word recognition. *American Psychologist*, *105*, 1–27.
- Posner, M. I., Choate, L., Rafal, R. D., & Vaughan, J. (1985). Inhibition of return: Neural mechanisms and function. *Cognitive Neuropsychology*, *2*, 211–228.
- Posner, M. I., & DiGirolamo, G. J. (1998). Executive attention: Conflict, target detection, and cognitive control. In R. Parasuraman (Ed.), *The attentive brain* (pp. 401–423). Cambridge, MA: The MIT Press.
- Posner, M. I., Inhoff, A. W., Friedrich, F. J., & Cohen, A. (1987). Isolating attentional systems: A cognitive-anatomical analysis. *Psychobiology*, *15*, 107–121.
- Posner, M. I., Nissen, M. J., & Ogden, W. C. (1978). Attended and unattended processing modes: The role of set for spatial location. In H. L. Pick & I. J. Saltzman (Eds.), *Modes of perceiving and processing information* (pp. 137–158). Hillsdale, NJ: Erlbaum.
- Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, *13*, 25–42.
- Posner, M. I., Petersen, S. E., Fox, P. T., & Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science*, *240*, 1627–1631.
- Posner, M. I., & Rothbart, M. K. (1992). Attentional mechanisms and conscious experience. In M. Rugg & A. D. Milner (Eds.), *Neuropsychology of consciousness* (pp. 91–111). San Diego, CA: Academic Press.
- Posner, M. I., Snyder, C. R. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, *109*(2), 160–174.
- Posner, M. I., Walker, J. A., Friedrich, F. J., & Rafal, R. D. (1984). Effects of parietal lobe injury on covert orienting of visual attention. *Journal of Neuroscience*, *4*, 1863–1874.
- Pylyshyn, Z. W. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin*, *80*, 1–24.
- Pylyshyn, Z. W. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, *87*, 16–45.

- Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A. K., Pardo, J. V., Fox, P. T., & Petersen, S. E. (1993). Practice-related changes in human brain functional anatomy. *Society for Neuroscience Abstracts*, 17, 21.
- Rainville, P., Duncan, G. H., Price, D. D., Carrier, B., & Bushnell, M. C. (1997). Pain affect encoded in human anterior cingulate but not somatosensory cortex. *Science*, 277, 968–971.
- Rauch, S. L., Savage, C. R., Alpert, N. M., Dougherty, D., Kendrick, A., Curran, T., Brown, H. D., Manzo, P., Fischman, A. J., & Jenike, M. A. (1997). Probing striatal function in obsessive-compulsive disorder. A PET study of implicit sequence learning. *Journal of Neuropsychiatry & Clinical Neurosciences*, 9, 568–573.
- Riddoch, M. J., & Humphreys, G. W. (1987). Visual object processing in optic aphasia: A case of semantic access agnosia. *Cognitive Neuropsychology*, 4, 131–186.
- Robinson, D. L., & Petersen, S. E. (1986). The neurobiology of attention. In J. LeDoux & P. Hirst (Eds.), *Mind and brain: Dialogues in cognitive neuroscience*. New York: Cambridge University Press.
- Robinson, D. L., & Petersen, S. E. (1992). The pulvinar and visual salience. *Trends in Neurosciences*, 15, 127–132.
- Robinson, R. G. (1985). Lateralized behavioral and neurochemical consequences of unilateral brain injury in rats. In S. G. Glick (Ed.), *Cerebral lateralization in nonhuman primates*. Orlando, FL: Academic Press.
- Roediger, H. L. III., & Blaxton, T. A. (1987). Effects of varying modality, surface features, and retention interval on priming in word fragment completion. *Memory & Cognition*, 15, 379–388.
- Roediger, H. L. III., Weldon, M. S., Stadler, M. L., & Riegler, G. L. (1992). Direct comparison of two implicit memory tests: Word fragment and word stem completion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1251–1269.
- Roland, P. E., & Gulyas, B. (1994). Visual imagery and visual representation. *Trends in Neurosciences*, 17, 281–287.
- Rozin, P. (1976). The psychobiological approach to human memory. In M. R. Rosenzweig & E. L. Bennett (Eds.), *Neural mechanisms of learning and memory*. Cambridge, MA: MIT Press.
- Rueckl, J. G., Cave, K. R., & Kosslyn, S. M. (1989). Why are “what” and “where” processed by separate cortical systems? A computational investigation. *Journal of Cognitive Neuroscience*, 1, 171–186.
- Sagi, D., & Julesz, B. (1985). “Where” and “what” in vision. *Science*, 228, 1217–1219.
- Schacter, D. L. (1990). Perceptual representation systems and implicit memory: Toward a resolution of the multiple memory systems debate. *Annals of the New York Academy of Sciences*, 608, 543–571.
- Schacter, D. L. (1992). Understanding implicit memory: A cognitive neuroscience approach. *American Psychologist*, 47, 559–569.
- Schacter, D. L. (1994). Priming and multiple memory systems: Perceptual mechanisms of implicit memory. In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994*. Cambridge, MA: MIT Press.
- Schacter, D. L., & Church, B. (1992). Auditory priming: Implicit and explicit memory for words and voices. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 915–930.
- Schacter, D. L., Church, B. A., & Treadwell, J. (1994). Implicit memory in amnesic patients: Evidence for spared auditory priming. *Psychological Science*, 5, 20–25.
- Schacter, D. L., Cooper, L. A., & Delaney, S. M. (1990). Implicit memory for unfamiliar objects depends on access to structural descriptions. *Journal of Experimental Psychology: General*, 119, 5–24.
- Schacter, D. L., Cooper, L. A., Delaney, S. M., Peterson, M. A., & Tharan, M. (1991). Implicit memory for possible and impossible objects: Constraints on the construction of structural descriptions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 3–19.
- Schacter, D. L., Cooper, L. A., Tharan, M., & Rubens, A. B. (1991). Preserved priming of novel objects in patients with memory disorders. *Journal of Cognitive Neuroscience*, 3, 118–131.
- Schacter, D. L., Cooper, L. A., & Treadwell, J. (1993). Preserved priming of novel objects across size transformation in amnesic patients. *Psychological Science*, 4, 331–335.
- Schacter, D. L., Harbluk, J. L., & McLachlan, D. R. (1984). Retrieval without recollection: An experimental analysis of source amnesia. *Journal of Verbal Learning and Verbal Behavior*, 23, 593–611.

- Schacter, D. L., McGlynn, S. M., Milberg, W. P., & Church, B. A. (1993). Spared priming despite impaired comprehension: Implicit memory in a case of word meaning deafness. *Neuropsychology*, *7*, 107–118.
- Schacter, D. L., Rapcsak, S. Z., Rubens, A. B., Tharan, M., & Laguna, J. M. (1990). Priming effects in a letter-by-letter reader depend on access to the word form system. *Neuropsychologia*, *28*, 1079–1094.
- Schacter, D. L., & Tulving, E. (Eds.). (1994a). *Memory systems 1994*. Cambridge, MA: MIT Press.
- Schacter, D. L., & Tulving, E. (1994b). What are the memory systems of 1994? In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 1–38). Cambridge, MA: MIT Press.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology and Neurosurgery and Psychiatry*, *20*, 11–21.
- Sergent, J. (1991). Judgments of relative position and distance on representations of spatial relations. *Journal of Experimental Psychology: Human Perception and Performance*, *91*, 762–780.
- Sergent, J., Ohta, S., & MacDonald, B. (1992). Functional neuroanatomy of face and object processing: A positron emission tomography study. *Brain*, *115*, 15–36.
- Shepard, R. N., & Feng, C. (1972). A chronometric study of mental paper folding. *Cognitive Psychology*, *3*, 228–243.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, *171*, 701–703.
- Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review*, *94*, 439–454.
- Shiffrin, R. M. (1988). Attention. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Steven's handbook of experimental psychology* (pp. 739–811). New York: Wiley.
- Shimamura, A. P., & Squire, L. R. (1987). A neuropsychological study of fact memory and source amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 464–473.
- Shimamura, A. P., & Squire, L. S. (1988). Long-term memory in amnesia: Cued recall, recognition memory, and confidence ratings. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *14*, 763–770.
- Skinner, B. F. (1957). *Verbal behavior*. New York: Appleton-Century-Crofts.
- Slooman, S. A., Hayman, C. A. G., Ohta, N., Law, J., & Tulving, E. (1988). Forgetting in primed fragment completion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 223–239.
- Smith, M. E., & Oscar-Berman, M. (1990). Repetition priming of words and pseudowords in divided attention and amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 1033–1042.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, *74* (11, Whole No. 498).
- Squire, L. R. (1987). *Memory and brain*. New York: Oxford University Press.
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, *99*, 195–231.
- Squire, L. R., Amaral, D. G., & Press, G. A. (1990). Magnetic resonance measurement of hippocampal formation and mammillary nuclei distinguish medial temporal lobe and diencephalic amnesia. *Journal of Neuroscience*, *10*, 3106–3117.
- Squire, L. R., & Frambach, M. (1990). Cognitive skill learning in amnesia. *Psychobiology*, *18*, 109–117.
- Squire, L. R., Ojemann, J. G., Miezins, F. M., Petersen, S. E., Videen, T. O., & Raichle, M. E. (1992). Activation of the hippocampus in normal humans: A functional anatomical study of memory. *Proceedings of the National Academy of Sciences of the U.S.A.*, *89*, 1837–1841.
- Squire, L. R., & Shimamura, A. P. (1986). Characterizing amnesic patients for neurobehavioral study. *Behavioral Neuroscience*, *100*, 866–877.
- Srinivas, K. (1993). Perceptual specificity in nonverbal priming. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *19*, 582–602.
- Sternberg, S. (1969). Memory scanning: Mental processes revealed by reaction-time experiments. *American Scientist*, *57*, 421–457.

- Talbot, J. D., Marret, S., Evans, A. C., Meyer, E., Buswell, M. C., & Duncan, D. H. (1991). Multiple representations of pain in human cerebral cortex. *Science*, *251*, 1355–1358.
- Teuber, H. L. (1955). Physiological psychology. *Annual Review of Psychology*, *6*, 267–296.
- Tootell, R. B. H., Silverman, M. S., Switkes, E., & DeValois, R. L. (1982). Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science*, *218*, 902–904.
- Townsend, J. T. (1974). Issues and models concerning the processing of a finite number of inputs. In B. H. Kantrowitz (Ed.), *Human information processing: Tutorials in performance and cognition* (pp. 133–185). New York: Wiley.
- Treisman, A. M., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, *12*, 97–136.
- Treisman, A. M., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, *95*, 15–48.
- Tulving, E., Hayman, C. A. G., & MacDonald, C. (1991). Long-lasting perceptual priming and semantic learning in amnesia: A case experiment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 595–617.
- Tulving, E., Kapur, S., Craik, F. I. M., Moscovitch, M., & Houle, S. (1994). Hemispheric encoding/retrieval asymmetry in episodic memory: Positron emission tomography findings. *Proceedings of the National Academy of Science*, *91*, 2016–2020.
- Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. *Science*, *247*, 301–306.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549–586). Cambridge, MA: MIT Press.
- Vogt, B. A., & Miller, M. W. (1983). Cortical connections between rat cingulate cortex and visual, motor, and postsubicular cortices. *Journal of Comparative Neurology*, *216*, 192–210.
- von Neumann, J. (1958). *The computer and the brain*. New Haven, CT: Yale University Press.
- Warrington, E. K., & James, M. (1991). A new test of object decision: 2D silhouettes featuring a minimal view. *Cortex*, *27*, 377–383.
- Warrington, E. K., & Shallice, T. (1980). Word-form dyslexia. *Brain*, *30*, 99–112.
- Warrington, E. K., & Taylor, A. M. (1973). The contribution of the right parietal lobe to object recognition. *Cortex*, *9*, 152–164.
- Warrington, E. K., & Taylor, A. M. (1978). Two categorical stages of object recognition. *Perception*, *7*, 695–705.
- Warrington, E. K., & Weiskrantz, L. (1974). The effect of prior learning on subsequent retention in amnesic patients. *Neuropsychologia*, *12*, 419–428.
- Weingartner, H., & Lister, R. (1991). *Perspectives on cognitive neuroscience*. New York: Oxford University Press.
- Witelson, S. F. (1994). Clinical neurology as data for basic neuroscience: Tourette's syndrome and human motor systems. *Neurology*, *43*, 859–861.
- Wurtz, R. H., Goldberg, M. E., & Robinson, D. L. (1980). Behavioral modulation of visual responses in the monkey: Stimulus selection for attention and movement. In J. M. Sprague & A. N. Epstein (Eds.), *Progress in psychobiology and physiological psychology* (pp. 48–83). New York: Academic Press.
- Zola-Morgan, S., & Squire, L. R. (1990a). The neuropsychology of memory: Parallel findings in humans and nonhuman primates. *Annals of the New York Academy of Sciences*, *608*, 434–456.
- Zola-Morgan, S., & Squire, L. R. (1990b). The primate hippocampal formation: Evidence for a time-limited role in memory storage. *Science*, *250*, 288–290.
- Zola-Morgan, S., Squire, L. R., & Amaral, D. G. (1986). Human amnesia and the medial temporal lobe region: Enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus. *Journal of Neuroscience*, *6*, 2950–2967.
- Zola-Morgan, S., Squire, L. R., Amaral, D. G., & Suzuki, W. A. (1989). Lesions of perirhinal and parahippocampal cortex that spare the amygdala and hippocampal formation produce severe memory impairment. *Journal of Neuroscience*, *9*, 4355–4370.