In M.S. Gazzaniga, (Ed.),
Perspectives in Memory Research.
Cambridge, MA: The MIT Press (1988).

III

Present and Future of Memory Research and Its Applications

William Hirst and
Michael S. Gazzaniga

Anyone with even a passing professional interest in memory has been faced with the smiling face of a relative probing, "So, you study memory, do you? Well, I have a terrible memory. Can't remember names. How can you help me?" Most students of memory return the smile, mumble something about things being complex, and scan the room for a more pleasant conversationalist, for the truth is that despite many years of study little is known about how to improve memory. Students of memory might caution that improvement and application is not their goal; rather, they want to understand how memory works. We often reply to our inquisitive relatives that their memory is quite fine, at least as good as one can expect. This assurance might be true for many, but for the elderly sinking into dementia, for students unable to keep up with the rest of the class, for employees who have trouble mastering their jobs in a timely fashion, memory is clearly not working as well as expected. For these people application, not understanding, is important.

In this part we summarize the lessons learned from the previous chapters. The book was motivated by the question, Does present and forseeable biological and psychological knowledge about the structure and function of memory hold hope for people with memory problems? Each contributor was asked to push his or her current understanding about either the neuroscience or psychology of memory beyond the current state of knowledge. They were asked to see if any insights into the functions or structure of memory could be applied to a variety of educational, medical, and vocational settings. Thus much of the book is speculative. Black et al. offer a model of neuronal change with long-term persistence; Lynch and Baudry infer psychological function of brain structures from lesion work and the route of connections; Shepherd investigates commonalities among neural circuitries; Sejnowski and Rosenberg explore the implication that recent advances in neuronal modeling have for memory; Kutas pushes event-related potential (ERP) data toward questions of psychological relevance; Hirst examines how psychological principles of memory might be applied to the job environment; Kosslyn develops a model of imagery based on neurological and psychological data.

We began our review of these chapters with a prejudice about what is meant by the term "memory." At some level, even a spring can be said to have a memory, in that it bounces back to its original position every time it is squeezed. Clearly, if we are interested in human memory, as we must be if our concern is application, then a broad-ranging definition of memory will not do. Accompanying any human memory is the belief that the information was acquired in the rememberer's personal past. People may not remember the incident in which a remembered event or fact occurred, but nevertheless they must believe that they either experienced the remembered event or learned the remembered fact. This concept of belief is a bit woolly, and clearly with a better understanding of memory a less subjective term may be possible. But something along these lines is necessary if memory "images" are ever to be differentiated from images as products of pure imagination. From this perspective cells cannot be said to have "memories." They may show persistent structural changes, but this structural change is quite removed from the human experience of memorizing and remembering. This persistence may indeed serve as a basis for memory, in the sense that we will use it, but work must be done to establish the connection. It cannot be taken for granted.

The Present State of Knowledge

None of the authors here is ready to offer a strong solution for people's memory difficulties. Hirst argued that curriculum and the to-be-remembered material itself should be designed with the demands and principles of memory in mind. He actually shows some examples of how this designing might be done. But even his chapter is programmatic in nature because it points in the right direction rather than maps out an easily applicable algorithm. The real question, however, is not whether the chapters offer solutions for memory problems but whether they suggest ways in which research might progress so that solutions are eventually found.

We discuss first the work on the neuroscience of memory and then the chapters written from a psychological perspective.

Neuroscience

The neuroscience of memory has made much progress since Hebb (1949) first sketched a theory of memory and learning, but Hebb's thoughts on the matter still guide much of the work. Hebb argued

that any experience leads to the transmission of nervous impulses through a neuronal circuit. Repeated exposure facilitates transmission. The resulting cell assembly, with the rapid transmission of impulses, constitutes a biological trace of the experience.

Hebb supplied a framework for understanding memory in biological terms. He left a host of questions, however. How are nerve impulses transmitted? What changes occur in the neuron to facilitate transmission? How long is this change in effect? How are the circuits structured? Are there similarities among circuits? And so on. In other words, one needed to fill in the Hebbian framework with details about anatomy, biochemistry, physiology, and computational power.

The contributors to the neuroscience part of this book summarize recent attempts to fill in these details. It is probably worthwhile to pause to see how far each of them goes in developing a biological theory of memory.

The Biochemistry of Memory In their chapter, Black et al. are interested in the biochemistry of "information storage in the nervous system." As they note, the transmission of nerve impulses across a synapse depends crucially on the neurotransmitters present. On average, a neuron uses three different neurotransmitters, and, if we accept Black et al.'s assumption that these transmitters come in five different concentrations, then each neuron can be in 1 of 244 different neurotransmitter-defined states. When it is considered that there are 100 billion neurons, the potential that neurotransmitters have for defining unique brain states is transparent. For Black et al. these unique brain states can be said to represent information, and, as a consequence, neurotransmitters can be treated as "communicative symbols."

If these communicative symbols are to have anything to do with the storage of information, as opposed to the expression of information, then the time course of any change is important. If environmental stimulation causes a change in the production and concentration of neurotransmitters, then the change must be maintained for a long time if it is to represent the storage of information. The effort for Black et al. then, is to find changes in neurotransmitter concentrations that have a long-term presence.

Black et al. use as their model the sympathetic nervous system, which is involved in fight-or-flight responses, and examine the role of norepinephrine (NE) and substance P. Consider the discussion of NE. The biosynthesis of NE depends on the conversion of circulating tyrosine to L-DOPA. The key enzyme for this conversion is tyrosine hydroxylase (TH), and the thrust of Black et al.'s argument is that TH is a communicative symbol important to the storage of information. Black and his co-workers show that stressful stimuli resulting in sympathetic activation elicit a two- to threefold elevation of TH in sympathetic neurons within two days, and that the enzyme remains elevated for at least three days. Moreover, the elevated level of TH will depend on the extent of the initial stimulation and the number of stimulations, reflecting well-known properties of memory. TH clearly acts as if it plays a role in information storage.

One, of course, must face the question of whether this mechanism, defined on the periphery, has anything to do with memory, which after all involves the central nervous system. Black et al. raise this question and argue that the same properties of TH found in the sympathetic nervous system can also be found in the nucleus locus ceruleus. Indeed, enzyme activity following stimulation is significantly elevated after twelve days and in the frontal cortex can last up to three weeks. But this argument does not directly confront the relevance of the TH model to a theory of memory. Therein lies the limitations of Black et al.'s theory.

Black and his co-workers' energies are focused on the actions of neurotransmitters and related enzymes. They present evidence for various neuronal changes with stimulation that have long-lasting effects. How these long-lasting effects influence memory have yet to be determined. So far, the duration of any observable change is no more than three or four weeks; yet some memories last for years. Clearly the mechanisms sketched by Black et al. cannot be the complete story.

A Neural Model of Memory Function Lynch and Baudry try to map some of the functions of memory onto neurophysiological structures and processes. Olfactory memory is considered for two main reasons: first, because it is simple (the sensory epithelium is only two synapses away from the cortex) and, second, because there is little interspecies variation. Lynch and Baudry trace the flow of information from the olfactory bulb through the pyriform cortex to the dorsomedial thalamic nucleus or the hippocampus, and from the thalamic structure to the frontal cortex. Each of these structures is linked to an aspect of olfactory memory on the basis of lesion work. The pyriform cortex is associated with representational memory because of the quasi-random nature of the connections between the lateral olfactory tract and the pyriform cortex. The distribution ensures that any combination of cell firing in the tract causes excitation in a small area of the pyriform cortex. The dorsomedial thalamus is associated with either the mapping of appropriate response patterns to particular

olfactory cues or the categorization of odors. As support for this claim, Lynch and Baudry point to the highly convergent nature of the projections from the pyriform cortex, where individual odors are represented, to the thalamus. Finally, the hippocampus is associated with consolidation of representational but not procedural memory by means of encoding of contextual information, in part because the hippocampus receives equal-size inputs of olfactory and nonolfactory (contextual) information. Lynch and Baudry also note that the hippocampus is capable of long-term potentiation (LTP), that is, a persistent potentiation of hippocampal synapses following brief periods of high-frequency synaptic stimulation. This LTP may be the mechanism by which memories are consolidated in the hippocampus, a proposal that Lynch and Baudry explore in detail.

It is interesting that Lynch and Baudry propose a context theory of memory on the basis of the functional neuroanatomy of the olfactory memory system and the hippocampus. That is, they assign to the hippocampus the role of integrating olfactory and contextual or nonolfactory information rather than the traditional role of memory consolidation per se. This thesis conforms with some work on temporal lobe amnesia in humans. Several investigators have suggested that damage to the hippocampus disrupts the encoding of information about the context in which the target occurred, not the encoding of the target itself (Hirst 1982; Hirst and Volpe 1984a; Mayes et al. 1985). The resulting mnemonic representation is impoverished, with targets connected semantically but not contextually. This proposal is supported by several lines of evidence. First, amnesiacs can show retention of the to-be-remembered material if properly cued. Thus amnesiacs' responsiveness to a recognition probe (Hirst et al. 1986) and strong semantic cues is better than one might expect from their poor free recall (Hirst et al. 1987; Warrington and Weiskrantz 1971). Second, although there is some controversy over the contribution of the frontal lobe (Schacter 1987b), amnesiacs' memory for the spatial location of to-be-remembered objects (Hirst and Volpe 1984b; Smith and Milner 1981), the temporal order of events (Hirst and Volpe 1982; Mayes et al. 1985), and the source of the to-be-remembered material (Schacter et al. 1984) is much worse than one would expect from their target recognition scores. Third, work with normal subjects indicates that the encoding of contextual information is qualitatively different from the encoding of target information. Finally, again in normal subjects, the ability to discriminate to-be-remembered events on the basis of contextual information is more important for successful recall than it is for successful recognition.

Thus Lynch and Baudry's functional analysis, their review of the lesion work with rats, the recent work on amnesia in humans, and the studies of contextual encoding in normal subjects converge to underline the importance of contextual information for successful encoding of information. But this progress should not obscure some of the problems inherent in Lynch and Baudry's model. First, to make a point similar to one raised in the discussion of Black et al.'s work, there is at this point no direct evidence that LTP is related to psychological memory. LTP is essentially a neuronal change that can last for weeks. The causal link between this change and memory must still be established. Second, the assignment of psychological functioning to various structures on the basis of their connections to other structures is at best conjectural. Again, the necessary behavioral work has not yet been done. Finally, it is unclear how the role assigned to the hippocampus for odors will generalize to other stimuli, despite Lynch and Baudry's assurances.

The Circuitry of Memory Shepherd is concerned with microcircuitry, the structure that links an input to a neuron with its output. Shepherd shows that nature is conservative and the structure present in the more primitive paleocortex holds, with variations, for the archicortex and even the neocortex. Consequently the genetic mechanism governing cortical growth may be much simpler than might be expected, inasmuch as the genetic mechanism does not have to be sensitive to the cortical area in which the growth is taking place.

Shepherd organizes his discussion of the extant research around the synaptic triad: principal output neuron, input fiber, and interneurons. For the paleocortex, the hippocampus, the submammalian general cortex, and the mammalian neocortex, Shepherd identifies the three elements and characterizes their organization in terms of input processing and output control. Some of the similarities that he finds between the basic circuitry of these cortical structures are as follows. First, each cortex has as its principal output neuron a pyramidal type of cell, with apical and basal dendritic compartments. Second, specific inputs to the three types of simple cortex are made onto the apical dendrites of the pyramidal neurons. For all cortices the input fibers make excitatory synapses onto spines of the distal apical branches. Feed-forward inhibitory pathways are also found in several of the cortical structures.

Moreover, the input processing in each cortex is remarkably similar. The process of integrating different inputs begins at the individual spines of the apical dendrite and continues as spines and branches interact. Integrative interactions are governed by several factors, including spatial distribution of excitatory and inhibitory inputs and the distribution of voltage-gated membrane channels. Output control is mediated in each cortex by inhibitory interneurons, which can either feed forward or feed back. Reexcitation within a circuit is mediated by axon collaterals.

The challenge of Shepherd's research is to go from the detailed microcircuitry to a discussion of memory. Shepherd is claiming that the structures are similar across the cortex, but memory functioning is not uniformly distributed across cortices. Nevertheless, Shepherd does suggest that the reexcitation provided by axon collaterals can account for short-term memory. He further suggests that, whatever the mechanism of long-term memory, it probably rests at the connection between input fibers and the spines on the apical dendrites.

It is unfair to Shepherd's aims to demand a model of memory from his work. His chapter is more a discussion of constraints rather than a framework for understanding memory and learning. If memories are to be represented in the weighted connections of a network, as many researchers have claimed since Hebb—the contribution of Sejnowski and Rosenberg in this book fits into this school—then Shepherd places firm constraints on the architecture of these networks. It is often difficult to see what the implications of these constraints are for psychological functioning. The virtue of building them into a network model such as Sejnowski and Rosenberg's is that one can then compare the functioning of the model with and without the constraints. Such a comparison should provide some hint about the implications of the constraints mentioned by Shepherd. We return to this point when discussing Sejnowski and Rosenberg's chapter.

Summary Most biological models of memory assume that memories are laid down through changes in neuronal connections, either through neuronal growth or through changes at the level of the synapse, or, to follow Shepherd, at dendritic spines. Most of the models presented here involve changes that occur rather quickly following stimulation but last for at most four weeks. Thus they can nicely account for the rapid acquisition of memories, especially short-term memories. Their difficulties arise when trying to describe changes that would account for memories that last years if not decades. Moreover, even if any of the accounts could be extended to cover lasting memories, the necessary behavioral research clearly establishing a link between the posited mechanism and memory needs to be pursued.

Thus in the chapters here we see several proposals for a biological theory of memory. They would be more accurately described as sketches of a theory. Their chief virtue is that they raise many empirically addressable questions; their chief vice is that they are more data impoverished, especially in terms of behavioral results, than a theory should be.

Cognitive Science

In considering the psychology of memory, one is tempted to forgo any theoretical development and just use one's intuitions. Ancient Greeks did just that, with some success, but this intuitive method has its limits. The Greeks, for instance, disagreed over the effectiveness of mnemonics like the method of loci for memorizing material such as prose. This issue can clearly be resolved empirically using the methodology of psychology. The same point can be made for several other deeper issues. To what extent can training improve memory? To what extent does mnemonic capacity differ across individuals? What are the biological constraints on memory, and are they susceptible to psychological manipulation? These questions are raised in several of the chapters, and empirical research, not intuitions, form the basis of the tentative answers offered by the authors. All of the research described in part II on the psychology of memory involves humans as research subjects. We consider each chapter separately.

Using Biological Measures to Constrain Theories of Mind Kutas reviews recent evidence from the event-related potential literature that bears on cognitive psychological modeling. In doing so, she underlines the close relation between mind and brain. Her discussion focuses on three issues. First, she argues that one should examine the function of latency of P300 plotted against memory set size when considering the Sternberg search experiment. In particular, she argues that the slope of this function is a more accurate measure of speed of serial comparison than the standard reaction time (RT) slope because the RT measure includes the time it takes to encode a stimulus and the time it takes to respond, whereas the P300 measure directly assesses serial comparison speed. The need for alternative interpretations of P3 latencies and RT is dramatically illustrated in work on memory scanning in the aged. The slope of the standard function plotting RT against set size increases with age, whereas the slope of the function plotting P3 against set size does not. If P3 is a more accurate measure of speed of serial comparison, then memory scanning may slow in the elderly not because their speed of serial comparison slows but because the speed with which stimuli are encoded or responses are made decreases.

Second, Kutas reviews research on P3 in verbal learning experiments. She adopts the viewpoint that large P3's can be found

whenever an item in working memory must be updated. To the extent that this updating will lead to a strong memory trace, a large P3 in encoding should predict subsequent recall. In a variety of experiments this prediction seems to be verified. There are some interesting twists. For instance, Karis et al. (1984) examined the von Restorff effect and presented subjects with lists of words in which all of the words but one were typed in small letters. The "isolate" was typed in capital letters. Karis et al. found that P3 measured during encoding best predicted recall for the isolate for those subjects who seemed most sensitive to the von Restorff effect. Thus some subjects recalled the list quite well but did not remember the isolate any better than the other words in the list. These subjects presumably used associational and other semantic strategies when memorizing. The updating of the isolate in working memory had little bearing on the outcome of these strategies. Consequently P3 was not very predictive of recall for this group. Other subjects recalled the list poorly but remembered the isolates quite well. Clearly the oddness of the word was important to their memorization strategy. For this group P3 was predictive of subsequent recall. In this experiment P3 appears to be a sensitive measure of cognitive strategies.

Third, Kutas reviews work investigating whether ERPs can serve as a measure of depth of processing, a measure that has eluded traditional cognitive psychologists. Unfortunately, at present there does not seem to be any difference in the ERPs found when subjects process a word superficially or meaningfully. However, the nature of the response—for instance, a "yes" versus a "no" response—does elicit differentiable ERPs.

Much research has built on this observation. The N400 wave is relevant here. It is produced by semantic incongruity, for instance, an unexpected ending to a sentence. Semantically incongruous words, that is, words with strong N400's, are not remembered well. Interestingly Kutas is able to predict that incongruous words related to a possible congruous word (for example, "The game was called when it started to umbrella") would be remembered better than incongruous words unrelated to a possible congruous word (for example, "George was fired but he could not tell his fog") purely on the basis of the N400's that she observed with these sentences. Other work had shown that words with small N400's were better remembered than words with large N400's and that related incongruous words produced smaller N400's than the unrelated incongruous words. Thus Kutas accurately predicts that the related incongruous words would be better remembered than the unrelated incongruous words. N400 and P300 appear to probe mental processing sensitively enough that detailed predictions about the consequences of this processing can be made.

Theories of Imagery and Implications about Biological Constraints Kosslyn is mainly concerned with imagery, not memory, but the principles that he expounds are central to any cognitive theory of memory. In his chapter he notes that cognitive scientists attempt to characterize the mind in terms of separate subsystems. Imagery, for instance, is not a single undifferentiated event or ability; rather, it can be divided into subsystems that carry out four separate functions: the generation of images from information stored in long-term memory, the retention of an image, the inspection of an image, and the transformation of an image, such as rotation. A chief claim of Kosslyn's is that the different subsystems are neurologically real; that is, various neuroscientific probes will alter behavior in ways consistent with the disruption of normal functioning of a subsystem. Kosslyn has developed a model of imagery detailed enough to permit successful computer simulation. The computer program's subsystems include a normalizer, a shape categorizer, a position calculator, associative memory, and so on. Kosslyn maps these subsystems onto two well-known visual tracts. This effort is important because the tracts serve two quite different perceptual functions. The ventral system is involved with the perception of shape, whereas the dorsal system is involved with the perception of location. Thus subsystems important to the analysis of shape, such as the shape categorizer and the normalizer, are treated as parts of the ventral system, whereas subsystems important in placing an image in a mental coordinate plan, such as the categorical relations encoder and the coordinate relations encoder, are treated as parts of the dorsal system.

This modularity is important when considering training for imaging or other cognitive tasks. As Kosslyn notes, people do not differ in their ability to image; they differ in their ability to transform, create, or maintain an image. That is, differences exist in the subsystems and not the overall process. Consequently, when devising a training curriculum, one must concentrate on improving the individual subsystems. It may help to teach someone to image, but it would be better to teach him or her to normalize an image or to categorize a shape. For Kosslyn a fuller understanding of the components of imaging is a necessary prerequisite for any training program because these components indicate what skills must be taught. The same point can be made for memory training. Of course, Kosslyn is not as concerned about memory as he is with imaging, but a componential analysis of memory may guide the development of memory training just as it should guide the development of imaging training.

At the end of his chapter Kosslyn asserts that "training probably will improve only the subsystems used in a task and only in the precise way they are being used. If the option is available, it is better to select a person with the right skills for a task than to try to train a person for the task." Two claims are embedded in this proposal. First, training is usually quite task specific and often does not transfer from one situation to another, even if the same subsystem is used. Second, basic abilities, representing what Kosslyn claims are neurologically real subsystems, may be individual traits not easily altered with training.

Application Most of the chapters in this book are concerned with development of theory. In his chapter Hirst tries to apply what theory exists to practical problems. He offers a brief review of the literature on improving memory and emphasizes that memory is to a great extent what people do with stimuli, not what their intentions or motivations are. He articulates several well-known principles to a better memory. People, for instance, can organize material and search for its meaning, see connections between the material itself (organizing), establish relations between the material and other knowledge (elaborating), and draw distinctions between different to-be-remembered material and general world knowledge (discriminating). They can also image material or build retrieval cues in order to aid memorizing and remembering. Almost all courses on memory aids and all the well-known mnemonic techniques taught in them involve one or more of these general strategies. Hirst shows how these strategies could be applied to practical settings, illustrating his discussion with an example in which trainees have to learn the function of meters and levers on a control panel. Hirst argues that, although these strategies could be applied intentionally by memorizers in order to facilitate their memory, it is probably more worthwhile for the material to be structured and tasks be given so that learners must organize, elaborate, discriminate, image, and build retrieval cues regardless of their intentions or motivations. Thus cover tasks could force individuals to memorize and remember in effective ways, even though the learner may not conceive of the task as involving memory. Hirst argues that the careful design of such cover tasks should lead to as effective memorization as any intentional memorization would yield, but without the need for motivation or even good strategic thinking.

Hirst confronts many of the issues raised by Kosslyn, but he is not as pessimistic about the possibilities of training as Kosslyn is. First, consider the problem of task specificity. Hirst readily admits that, whatever the nature of any training, it must be designed for the task at hand. For Kosslyn this constraint creates severe limitations, in that Kosslyn probably in the end wants training to make a more intelligent person or a better imager or a better memorizer. Hirst would agree with Kosslyn that the prognosis for such drastic changes is not encouraging [but see Herrnstein et al. (1986)]. But if you are concerned with designing a curriculum to teach someone how to use a piece of equipment on the job or the attributes of a set of boats, you do not have to make a more intelligent person or a better memorizer. You just have to teach the material at hand. From this perspective both Kosslyn and Hirst would agree that training can be designed to be more effective.

Hirst outlines several things that should be considered when planning a training procedure or developing equipment. They are a direct outcome of theories of memory that stress organization, discrimination, and elaboration as central to the encoding process. But Hirst admits that applying these principles is an art and probably requires a better than average person. His point is that the training should be devised so that the processes important to memorization occur outside the trainees' own volition.

As noted, the heart of his proposal is that what matters in memory is what you do, not what your capacity is. On the surface this contention also seems to contradict Kosslyn, who argues for individual differences in the capacity of the subsystems of imaging; however, once Kosslyn's argument is carefully stated, the apparent contradiction disappears. For Kosslyn individual differences are traceable to the subsystems of imaging and not imaging per se. This same point can be mapped onto memory. Memory is not a unified process any more than imaging is. It also involves a host of systems and processes, any one of which could be better or worse than the average. So to speak of someone as having a good or bad memory makes no more sense than to speak of someone as being a good or bad imager. One may have trouble classifying or elaborating on a list of equipment terms, and these difficulties may lead to poor recall or recognition. This statement, however, is different from saying that a person has a bad memory.

Thus, although Hirst may be right that the ability of someone to memorize and remember may not be limited by a general memory capacity, a far more subtle limitation on the components or modules underlying memorizing and remembering may exist. Kosslyn has systematically studied individual differences in imaging with an emphasis on the subsystems, not overall performance. As far as we know, a similar program has not been carried out in the area of memory. Such a program of research may provide a better understanding of how to assess students or trainees and what kind of instruction they should receive.

Uniting Mind and Brain

The background research reports from which this summary has been made possible is divided into two parts: the first concerned with neuroscience research and the second concerned with psychological research. Although both areas are concerned with memory, their vocabulary and methodology are quite different. It would be nice if there was some way to unite the two. The one chapter that we have not discussed so far offers a possible means of filling the gap. The parallel distributed processing (PDP) discussed by Sejnowski and Rosenberg offers a powerful tool for modeling psychological function in physiological terms. The model that they discuss consists of a network of nodes with weighted connections. The network is capable of learning through changes in the weighting of the connections. Such changes are accomplished by an algorithmic process known as backward propagation, which reassigns weight values in parallel across the network. Consequently the resulting representation of the learned material is distributed across the network. Early network models consisted of only two layers: a set of input nodes connected to a set of output nodes. The current generation of PDPs contain a third layer, the hidden layer, in which nodes mediate between input and output.

Sejnowski and Rosenberg focus on a particular PDP model, called NETtalk. Here the input is a set of orthographic features of words and the output is the phonological features of the words. As Sejnowski and Rosenberg show, NETtalk can learn to "read" text and in other learning contexts demonstrate such well-known psychological effects as spacing.

Clearly a network consisting of connected nodes can show quite impressive learning if the weights of the connections are allowed to change with experience. From a neuroscience perspective there is an easy analogy between massive parallel networks and neuronal circuitry, at least on an abstract level. Nodes and connections can be thought of as neurons and synapses, respectively. Changes in neuronal structure occurring with experience can be thought of as changes in the weights of the connections between nodes. From a psychological perspective massively parallel networks can not only learn but seem to learn in psychologically real ways, as Sejnowski and

Rosenberg imply when discussing the spacing effect. So at one level we have a plausible psychological model; at another level, a plausible physiological model.

Yet it may be premature to cast an unreserved vote for PDP models. First, it is not known whether PDP models, as currently conceived, should elegantly capture all sorts of processing. As any cognitive psychologist knows, some processing is done serially, whereas other processing is done in parallel. For instance, there is much evidence that search through short-term memory is done serially and that search through a large array of trigrams is parallel. Should both forms of search be modeled as parallel processing? In the early days of information processing psychology, when the von Neumann computers dominated theorizing, many investigators were tempted to model every process as serial. Now, many workers on PDP try to model every process as parallel. Clearly some thought must be given to the best medium for modeling a particular process.

Moreover, although NETtalk mimics the spacing effect, there is more to memory than the spacing effect. For instance, NETtalk learns things more quickly than the psychological study of reading would suggest. Moreover, it learns to read in a qualitatively different way from the way a child learns to read. NETtalk heard the same passage over and over again. Children may hear the same story over and over again, and they do not generalize from this one story and begin reading words in new contexts. Indeed, a child learns by working with simple text first and then slowly building up to complex task. NETtalk jumps right into the complex texts. The situation here may be similar to that in early work on transformational grammar, when it was pointed out that transformational grammar as a formal mechanism was much too powerful (Peters and Richie 1973). It could describe not only natural language but also a much larger class of languages—natural and unnatural language if you like. The PDP models currently being considered may also be too powerful. They may not only model human processing and human representation of information; at the extreme they may also be able to capture a much larger class of processing, human and nonhuman alike.

A related problem is that, as a biological model of brain processing, the current crop of PDP models do not fully capture the complexity of biological phenomena. As already noted, there are limits on the extent of nonstructural LTP and time constraints on neuronal growth. These time parameters are presumably important in the formation of memories; yet there is nothing in the general conception of PDP that takes into account these limitations. Similarly, as several of the chapters make clear, the workings of the synapse are quite complex and presumably have implications for learning. Again, the PDP models do not capture this complexity. Finally, Shepherd pointed out several principles governing the structure of neural microcircuitry. These same principles should presumably be built into the microcircuitry of PDP networks.

It would seem, then, that if the PDP models are to bear more than a superficial similarity to biological models of memory, they must be constrained in both biological and psychological terms. And there is no reason why they cannot be. As currently conceived, the models are indeed general, but there is nothing inherent in their formulation that forces this generality. Indeed, one of the promises of PDP modeling is that it may provide a good language in which to test formally the consequences of many observable biological constraints.

Moving from Present Research to Future Research: Implications for Training and Instruction

So far we have summarized chapters from this book on the state of both neuroscience and psychology of memory. Research has to date proved quite fruitful, and each area reviewed shows promise for future growth. But at some point one must move beyond the specifics of current research efforts and address what motivated this collection in the first place. One wants to know whether any of these specifics can translate into possible application. Particularly, three questions should be addressed: (1) How can the current state of the art of memory research be expanded so that it can bear on issues of instruction and training? (2) What aspects of training and instruction might be affected by current research and the proposed research? (3) How close is current and future research to actual application?

Let us say at the outset that at some level these are impossible questions to answer. Both neuroscience and the cognitive science of memory are in their infancy. They certainly have not developed to the point that applications fall out willy-nilly. There is still much art in going from the tentative findings of these fields to applications in learning and training. To some extent the difficulty has to do with the way information is presented in the field of memory. Neuroscientists and cognitive scientists are not driven by the need for application. They simply want to know how memory works. As a consequence, they are more likely to structure their information around questions central to a full understanding of the functioning of memory than around questions central to its improvement. To be sure, a full understanding of how memory works would probably tell us how to improve memory, but the field is far from such a complete story. Until it is available, applications of research findings might seem like glib translations of complex work to many working scientists.

Embedded in the three questions concerning application is a deeper issue. Even if there is agreement that application is not around the corner, how can the long-term goals of the neuroscientists and cognitive scientists be structured so that applications will seem more apparent or at least be closer to fruition? Could either the neuroscience or psychology of memory benefit by tying theory development and experimental research to practical concerns?

Neuroscience

Let us first consider neuroscience. When discussing neuroscience, we often reduce practical applications to pharmacological aids. The chapters in this volume did not address issues of pharmacology directly. They emphasized the development of theory about the biological foundations of memory. Although the theories chart rapid progress, theory development has not reached the stage at which it is possible to predict what drugs will improve memory. For that to happen, there must be a theory of how information is encoded, stored, and retrieved. Such a theory does not presently exist.

Should the research goal of developing a theory of the biology of memory be changed so that research focuses on finding a pharmacological aid to memory rather than a theory of how memory works? The emphasis on theory has led to simulations such as Sejnowski and Rosenberg's, studies of synaptic organization, as typified by Shepherd's work, and research on neurotransmitters, as in Black's work. A goal of discovering a memory drug might narrow the scope of discussion, with an emphasis on neurotransmitters.

But beyond narrowing the focus of discussion, emphasis on finding a drug for memory improvement might not be as productive as one might hope. The pharmacology of memory has not been terribly successful and has only revealed immense complications. It is worth reviewing briefly.

The Quick Fix Those researchers concerned foremost with the discovery of a memory drug are guided by the work on L-DOPA and Parkinson's disease. For many years, investigators of Parkinson's knew that an impoverished level of the neurotransmitter dopamine was associated with the disease, and consequently they tried to use dopaminergic drugs to cure the disease. The problem was finding the right drug and the right dosage. For years they tried without success. It was only when Cotzias administered L-DOPA at what appeared to be a senselessly high dosage that the "cure" was found. Similar

serendipity is prayed for in the work with so-called memory drugs. But an examination of the literature suggests that this trial and error approach has little chance of succeeding.

First, many different neurotransmitters appear to be involved in memorial processes. In a recent review of the literature on drugs and memory, Squire and Davis (1981) cautioned that they were limiting the range of their discussion but then discussed an incredibly large number of different neurotransmitter systems, reviewing drugs that affect the cortical level of neurotransmitters and related substances, including acetylcholine, norepinephrine, dopamine, ACTH, vasopressin, endorphins, and the opioid peptides. Clearly, unlike investigators of Parkinson's who were able to limit themselves to the dopaminergic system, memory researchers must investigate many more neurotransmitter systems.

But even if the number of systems is small, there are still problems and complications. Many drugs with an apparent effect on memory performance may act indirectly (by increasing attention, for instance) rather than directly on the mechanisms underlying memory.

Consider the work with vasopressin. Work with animals suggests that vasopressin is important for successful memory. For instance, administration of vasopressin such as lysine-8-vasopressin (LVP) to rats facilitates long-term retention of passive avoidance training (Ader and de Wied 1972; Bohus et al. 1978; Krejci and Kepkova 1978; Leshner and Roche 1977; Gold and van Bushkirk 1976) and sexually motivated learning (Bohus 1977). But this positive effect with animals does not easily transfer to humans. To be sure, positive effects of vasopressin on memory can be found. Four depressed patients treated with a long-acting analogue of vasopressin showed improved memory scores on a range of tests (Gold et al. 1979), and a group of normal adults aged 50 to 65 showed similar improvement when given a daily regimen of 16 IU of vasopressin nasal spray (Legros et al. 1978). In both cases, however, the observable memory improvement was also accompanied by improvement in other cognitive activities that could indirectly affect memory. The depressed patients showed improved affect, and the normal adults performed better on tests of perceptual-motor speed and attention. Thus apparent effects on memory may often have little to do with changes in the actual biological mechanism of memory.

Of course, careful research can separate the indirect effects on memory from the direct effects. And again, even if the effort were made and putative direct effects could be found, the right drug would probably still not be at hand. Consider acetylcholine (ACh). The brains of Alzheimer's patients with severe memory problems have an abnormally low level of ACh (Davies and Maloney 1976). Inasmuch as ACh facilitates transmission of nerve impulses across synapses, this abnormality is thought to bear on the Alzheimer's memory problems [see Baddeley (1976) for a discussion]. If the ACh level could be increased in these patients, the reasoning goes, then, like the work with L-DOPA, memory problems should be alleviated.

Unfortunately the scenario proved much more complicated than this. First, although nerve transmission will not occur if there is too little ACh, it also will not occur if there is too much ACh. For example, administration of an anticholinesterase, such as physostigmine or diisopropyl fluorophosphate (DFP), will inhibit the breakdown of ACh at the synapse and thereby increase the amount of ACh present for nerve impulse transmission. As the dosage of anticholinesterase increases, memory performance will first improve and then decline, reflecting the level of ACh (Deutsch et al. 1966).

But matters become even more complex, because the natural production of ACh appears to follow a complicated path when a nerve is stimulated. According to Deutsch (1971), the level of ACh declines between 30 minutes and 1 day after training and then gradually increases. This production schedule is not as counterintuitive as it may seem at first. When rats' maze learning was studied, their memory performance paralleled this putative production schedule, with performance dipping when tested between 30 minutes and 1 day after training and then gradually improving (Huppert and Deutsch 1969).

This ACh production schedule makes detailed and confirmed predictions about the effect of DFP a heartening prospect for anyone interested in empirical science. Thus DFP administered 30 minutes after training will facilitate memory performance because the level of ACh is low and the DFP increases it. However, by day 5, administration of DFP will actually hinder memory performance. At this point in time the level of ACh at the synapse has increased so much that any addition supplied through the mechanisms of DFP would raise the quantity to harmful levels.

The level of ACh cannot increase unendingly. At some point it must gradually decline, as the memory itself weakens. Indeed, Deutsch et al. (1966) found that, if DFP is administered late enough after training, it will once again have a facilitatory effect. In one study administrations after 14 and 28 days were compared. An inhibitory effect was observed after 14 days, but a facilitative effect was found after 28 days. This and other studies suggest that the effect of DFP depends on the strength of the trace and on the amount of time after training. Along the same lines it has been found that slow learners respond to physostigmine differently from fast learners.

Thus one simply cannot administer a cholinergic drug and expect a general improvement in memory performance. The success of the drug depends on the abilities of the learner, the age of the memory, and how well the memory was learned in the first place. This complicated picture must be viewed within a larger perspective to see how disastrous it is. As noted, the work on cholinergic drugs began because some patients with Alzheimer's disease have an impoverished level of ACh. What ACh has to do with memory is unclear. To be sure, under the right circumstances it enhances the transmission of nerve impulses, but does enhanced nerve transmission have anything to do with memory? No one knows. It could equally have an effect on task speed. In other words, ACh, even given its complexity, is investigated simply because people with memory problems have an impoverished level in the brain, not because it is known to have anything to do with memory.

Given this complicated and tenuous set of findings and connections, it is not surprising that work with cholinergic drugs in the clinic has met with little success. Researchers administer either anticholinesterases, which raise ACh level by preventing its breakdown, or choline and lecithin, which raise ACh level by aiding in the manufacture of ACh. Either set of drugs yields small but positive effects in some studies and no effects in others (Pearce 1984). Little order can be given to these studies, except that any positive effects that can be found depend on the severity of the amnesia.

Recently, more encouraging results have been reported using 1,2,3,4-tetrahydro-9-amniocridine (THA), a potent centrally acting anticholinesterase (Summers et al. 1986). When Alzheimer's patients of varying severity are placed on a daily regimen of THA, dramatic positive effects were seen for patients with mild dementia. Patients who had been forced to retire returned to their jobs, and others who could not manage daily tasks were once again responsible enough to govern their own affairs. Like so many drugs, however, THA improved not only memory but also a host of other cognitive activities. Whether it is a "memory" drug or a general cognitive enhancer is not known. Of course, from a clinical viewpoint, this distinction might not matter, and THA and other drugs, such as the cognitive enhancers being developed by Ayerst and Squibb, should be further explored. But will THA or any other cognitive enhancer help the subpopulation of people with memory problems, but without the more general cognitive deficits found in dementia?

Back to Theory Thus the pharmacology of memory has not fulfilled its promise. The few success stories, such as that with THA, may not even involve a memory drug but a general cognitive enhancer. The serendipity and trial and error quality that characterizes the search for a memory drug will probably continue so long as it is not guided by a theory of biology of memory. For that reason we would encourage future research on neuroscience to continue to be directed by theoretical concerns rather than applications. Advances in neuroscience have been swift, and the necessary theoretical framework for application may lie in the foreseeable future. We should, however, emphasize that, by turning away from the "quick fix" of pharmacology, we are not saying that such research should be avoided. It has its place, in small doses. But the main thrust of future research on the neuroscience of memory should continue to be driven by theory.

Cognitive Science

Discovering Subsystems Although the neuroscience of memory might not benefit from a strong focus on application, the cognitive science of memory may. In examining Hirst's and Kosslyn's chapters, a common theme emerges. To a large extent the folk psychology terms for talking about memory and imagery may not be the correct ones. Any cognitive system, be it the memory system or the imagery system, consists of separate subsystems. Cognitive science and cognitive neuroscience have focused energies on discovering the exact nature of these subsystems. Indeed, one of the major findings in cognitive science is that skills and abilities that had been viewed as unitary and undifferentiated, especially in folk psychology, have in fact a rich and complex underlying structure.

Kosslyn is particularly clear about this point when discussing work on imagery. Rather than being a single ability or trait, imagery appears to be a host of abilities. In his review Kosslyn mentions several abilities, including image retention, image generation, and image transformation. Individuals do not differ in their ability to image but in their ability to retain, generate, and transform images. Each of these processes is a separate unit of the cognitive system. The processes are called on when people image. The resulting phenomenological experience of imaging is in reality a complex combination of these processes.

This work on imagery illustrates the kind of research we think should be encouraged in cognitive science, especially memory research. Cognitive scientists must be encouraged to discover the subsystems of memory. Almost all researchers in the area of memory agree that memory is not a unitary faculty but consists of subsystems. The list of possible subsystems is quite long. People have differentiated separate storehouses in which memory can be stored. There are short-term and long-term memories (Atkinson and Shiffrin 1968), episodic and semantic memories (Tulving 1983), procedural and declarative memories (Cohen 1984; Squire 1982), and memories with and without awareness (Jacoby and Whiterspoon 1982). Researchers have also distinguished qualitatively different kinds of encoding. For instance, information about the spatial location, temporal order, and frequency of occurrence of an event may be encoded automatically, whereas information about the content of an event must be encoded effortfully (Hasher and Zacks 1979). Several distinctions have also been made in discussions of retrieval. Investigators, for instance, distinguish direct access to an item, as observed in a recognition task, from the retrieval of an item, as observed in a recall test [see Kintsch (1970)]. Others separate the processes on which familiarity judgments are based from the processes that guide retrieval (Mandler 1980). These storage, encoding, and retrieval distinctions are only beginning to be assembled into a coherent model of the subsystems of memory. Clearly more research must be done.

Methodology for Further Research Of course, before a model can be proposed, one must know which of the subsystems proposed by various researchers truly reflects the architecture of human memory. The subsystems mentioned in the preceding paragraph have already received some attention, especially the difference between short-term memory and long-term memory and between episodic memory and semantic memory. The distinctions between procedural and declarative memories and between direct and indirect memories are less well-understood (Hirst 1987; Schacter 1987a). As for the encoding and retrieval distinctions, these are only now beginning to receive extensive research interest. Furthermore, as far as we know, each of the distinctions tends to be treated without regard to its relation to the other distinctions. In other words, no one has taken the putative subsystems and fitted them into a single model of memory. For instance, no one has explored the relation between the formation of episodic memories and the automatic encoding of spatiotemporal information.

To a large extent the discovery of subsystems depends on discovering dissociations. But it is often easy to supply alternative interpretations for any putative dissociation. For that reason a single piece of evidence is rarely conclusive. Converging evidence based on a wide variety of techniques is necessary.

Psychologists argue for separate subsystems when performance in a task that requires subsystem A but not subsystem B is stochastically independent of performance in a task that requires subsystem B but not subsystem A. Thus short- and long-term memories are assumed to be separate components of memory because 15 seconds of distraction will obliterate a short-term memory but will have little or no effect on a long-term memory.

Such dissociations can be found in experiments using normal subjects when two tasks have differential effects on memory performance. One can also look at individual differences. Two tasks tap different subsystems if individual variation in performance in one task is stochastically independent of individual variation in performance in the second task. The study of brain-damaged patients is also important in the discovering of memory subsystems. For instance, the distinction between procedural and declarative memories is based to an overwhelming extent on the discovery that amnesiacs can learn perceptual-motor skills at the same rate as normal subjects, even though they cannot remember the fact of learning the skill. Kutas illustrates how the use of evoked potentials can support putative dissociations by showing that behavior in one task elicits different ERPs from those elicited in other tasks. Other imaging techniques, such as positron emission tomography (PET), magnoencephalography, and magnetic resonance imagery (MRI), can also provide support for dissociation between varying subsystems.

Thus it is possible to show that two subsystems exist because people behave differently when they perform tasks that require the two subsystems and because their brains also behave differently when different parts of the brain mediate different subsystems. Taken together, a compelling story can usually be assembled in favor of the distinct subsystems. The search for subsystems must take a multidisciplinary approach. Research on subsystems must search for interactions between memory tasks done by normal individuals, for stochastically independent variation among individuals on memory tasks, for dissociations among patients with varying degrees of brain damage, and for differences in brain activity for different memory tasks, as imaged by techniques such as ERPs, PET scans, MRI, and magnoencephalography.

A Distinction in Search of Articulation What putative subsystems of memory appear to be most deserving of continued research effort? One of the more interesting distinctions in the literature and one of the most poorly understood is the one that variously goes under the labels "procedural" versus "declarative" memories, "explicit" versus "implicit" memories, "semantic" versus "episodic" memories, and memory "with awareness" or "without awareness." This distinction

may reflect the difference between skill learning and fact learning, or it may reflect the difference between verbalizable and nonverbalizable knowledge, or there may be an alternative way to phrase the distinction.

To a large extent this set of distinctions has been introduced in the context of the study of amnesia. In order to appreciate what phenomenon researchers are trying to explain and why there are so many different versions, we review briefly the relevant work.

Although amnesia is often thought of as a failure to remember events shortly after they occur, the memory failure is rarely, if ever, complete. Insights into the subsystems of memory can be garnered from a full understanding of what aspects of memory are preserved and disrupted with amnesia. Initial studies of amnesia by Milner et al. (1968) supported the distinction between long-term memory and short-term memory. They found that amnesiacs could retain information for a short period without rehearsal but would quickly forget new information if distracted. That is, their short-term memory was intact and their long-term memory was disrupted. Although the distinction between short-term memory and long-term memory may account for this dissociation, it cannot explain more recent findings. Cohen and Squire (1980) reported that amnesiacs learn to read mirror images of words at the same rate as normal subjects, even though they do not remember learning the task. Clearly some kinds of long-term memories are preserved, whereas other kinds are disrupted. Cohen and Squire suggest the distinction between procedural and declarative memories, which, although intuitively appealing, cannot readily account for work on priming in amnesiacs.

If exposed to the word "garbage" and then later asked to say the first word that comes to mind that completes the stem "gar," amnesiacs are as likely to say "garbage" as are normal subjects. However, if asked instead to complete "gar" with a studied word, amnesiacs are much less likely to say "garbage" than are normal subjects [see Schacter (1987a) for a review of the relevant literature]. Cermak et al. (1985) have used such results to argue that declarative memory is subdivided into semantic and episodic memories and that the amnesiac's deficit rests with episodic memory. According to this model the memory system would be a three-tier hierarchy, with the memory system first divided into long-term memory and short-term memory, long-term memory then divided into declarative memory and procedural memory, and declarative memory finally divided into episodic and semantic memory.

Graf and Schacter (1985) have argued that this account cannot explain normal associate priming in amnesiacs. That is, when subjects study the paired associate "tulip–garbage" and are later asked to complete the stem "gar," they are more likely to complete it with "garbage" if the stem is given in the same context in which it was learned ("tulip") than if given in a novel context. This effect is as robust for amnesiacs as it is for normal subjects. Inasmuch as the context is semantically unrelated to the target, the improved priming in amnesiacs cannot be accounted for by simple spreading activation in an intact semantic memory. Graf and Schacter offer a discriptive dichotomy—between explicit memory and implicit memory. Amnesiacs have trouble with tasks involving the explicit use of memory but not with those involving the implicit use of memory.

Hirst et al. (1986) have argued that even this descriptive dichotomy cannot fully account for what is preserved with amnesia. They showed that amnesiacs can recognize information better than one might expect given their poor recall. Moreover, this relatively preserved recognition cannot be attributed to their intact priming. Hirst argued that the amnesiac deficit is best described in terms of a breakdown in one kind of processing over another rather than in a disruption of particular storehouses of memory. Some of the mnemonic processes that supply the glue that holds individual events together and creates a unified representation are lost with amnesia. Without this glue memories would consist of a collection of individual traces of past events unconnected to one another. An amnesiac might be able to obtain direct access to these memories if provided with the appropriate probe (as in a priming or recognition task) but would not be able to search through memory in a systematic fashion. Hirst (1982) specifically builds on the distinction between the automatic encoding of context and the effortful encoding of content to suggest that amnesiacs do not encode the context of to-be-remembered events. As a consequence, events are not connected to each other along spatiotemporal dimensions.

Although each proposal—declarative versus procedural memories, explicit versus implicit learning, episodic versus semantic memories, encoding context versus encoding content—has been debated, the basic data has not. Research should be directed toward explaining these data. Each proposed distinction appears to be focusing on something that is as central to any theory of memory as the distinction between short-term and long-term memory. A full understanding of the implications of these distinctions is essential for the building of a model of memory. Research that would aggressively determine the empirical basis for such distinctions as procedural versus declarative memory should be encouraged. As noted, Hirst et al.'s (1986) results suggest that some aspect of declarative memory

may be intact with amnesia. Thus the amnesia syndrome does not place a clear cleavage between declarative and procedural memory. Hirst et al. (1986) suggest that it may be wrong to argue for qualitatively different forms of memory. Rather, they propose that emphasis should be placed on understanding qualitative differences in the way normal subjects and amnesiacs encode, retrieve, and represent information. Amnesiacs may be able to learn perceptual-motor skills and not a list of words because learning of the latter requires them to store the context in which the words were learned and to form associations between them, whereas the encoding of context and the formation of interlist associations is less important for the learning of perceptual-motor skills or the learning involved in priming or recognition. Research that goes beyond the discovery of dissociations and the establishment of different memory systems should be encouraged. A better understanding of the way amnesiacs encode, retrieve, and represent information is now needed.

A fuller understanding of amnesiac encoding, retrieval, and representation will permit a finer discussion of qualitatively different structures and processes in the memory system. The simple distinction between procedural and declarative memory is interesting, but it does not provide the kind of foothold needed by someone interested in training. If a person has difficulty learning declarative memories and not procedural memories, it does not help to know that these two might be qualitatively different when developing a training strategy. It would be better if one knew what about the way a person encoded, retrieved, and represented information made it difficult for them to learn declarative memories. For instance, if the problem is with the use of contextual information to aid retrieval, then training could be focused on ameliorating this deficit, either by overcoming it or by finding ways to circumvent it. Thus a fuller understanding of the amnesia syndrome may not only clarify the distinction between procedural and declarative memory but may also provide more refined categories on which a more sophisticated training strategy may be developed.

Implications for a Subsystem Approach for Training and Instruction Careful articulation of the subsystems of memory has implications for training and instruction. Any training and instructional program should consist of two phases: evaluation and training. In the evaluation phase tests should be given to determine individuals' strengths and weaknesses. These evaluations can then be used to assign trainees to the second phase: the instruction per se. The trainer could use the evaluation to assign trainees to an instructional program that stresses their strengths. People with good quantitative ability might be assigned to a program that teaches them to program computers. When there is no job or training program that fits a person's strengths, the evaluation can be used to target weaknesses and the trainee can be placed in a program that works with the weaknesses.

The success of this two-phase procedure will depend on the degree to which the evaluation procedure actually taps distinct abilities and the degree to which one can relate the distinct abilities to the tasks taught in the training. If the evaluations are not based on a componential analysis but on categories from folk psychology, it might not effectively measure individuals' strengths and weaknesses. Again, Kosslyn illustrates the point. Consider a test that examines subjects' ability to image rather than their ability to transform, preserve, or generate an image. The test might indicate that a particular trainee has poor imaging abilities, and, as a consequence, that trainee might not be put into a training program that requires good image generation. But this person might be quite good at generating images. He or she might have difficulty with image retention. The evaluation would fail to pick this up, and a good trainee would be lost. Similarly a trainee with putatively poor imaging skills might be put into a training program that focuses on building general imaging skills. This assignment would also be a mistake, inasmuch as the trainee does not need training on image generation.

The same procedure can be followed for the evaluation of memory strengths and weaknesses. For an effective evaluation strategy one must determine the subsystems underlying memory abilities and then use this information to construct evaluation procedures that assess each subsystem. Of course, one does not have to worry about the relevance of the discovery of mnemonic systems to training and instruction in order to do research on this issue. Simply wanting to know about the architecture of memory is motivation enough. So, as far as the discovery of subsystems goes, such discoveries may have implications for evaluation and hence training and instruction, but scientists doing the work do not have to worry about applications in order to produce useful knowledge. It is probably best to leave cognitive scientists to their own devices.

Determining the Subsystems Used When Completing a Task Tying research goals to practical concerns and applications becomes important when trying to determine the circumstances under which various subsystems are used. It is not enough to simply know what the subsystems are; one must also know how and when they are used. Processing subsystems probably do not mimic the overt behavior internally. Two tasks that seem quite similar on the surface may involve quite different subsystems. Alternatively two quite different tasks may draw on many of the same processing subsystems. Ice and water may seem quite different on the surface, but their underlying structure is quite similar. For example, to build on Kosslyn's chapter once again, image transformation and image generation may share a subsystem, even though on the surface they seem quite different. Transforming an object in an image (for example, imaging an object rotating) may be a noisy process, and hence a subsystem that realigns scrambled parts may be used. This subsystem in turn may draw on another subsystem that looks up in memory the proper alignment of the parts [see Kosslyn (this volume) for further details]. Image generation involves the process of forming a short-term-memory visual representation on the basis of information stored in long-term memory. Thus both image generation and image transformation require a person to look up in memory the spatial relations among parts of an object. Kosslyn has argued that much progress has been made in understanding what imagery tasks call on what subsystems, although clearly much research still needs to be done.

Schacter (personal communication) makes a similar point when he discusses a patient with a severe memory deficit following a stroke. The company the patient worked for decided to transfer her to the mailroom. They thought that a mail clerk performed menial tasks, and hence the patient could easily handle the job. She could not. It turned out that she performed well as a keypunch operator. Mail clerks use their memory constantly—in sorting the mail, in deciphering telegraphic addresses, and in plotting out their delivery route. A keypunch operator, however, places few demands on memory. Thus it is not enough to classify a job as menial or "higher level." A job that may be higher on the pay scale may make fewer demands on memory than a lesser paying job. Only a detailed task analysis can determine the demand characteristics of the job.

Cognitive scientists are task analyzers par excellence. They specialize in reducing a task to its components, and they have the requisite skills for determining which subsystems are involved in a particular task. But this analysis is not straightforward and requires careful experimentation and model building, often involving computer simulation. Thus, if there are tasks that need examination, it pays to emphasize them.

At present, such contact does not exist. If a trainee must learn to identify different ships, then there is no reason why ships cannot serve as the stimulus material for an experiment instead of nonsense shapes, pictures of faces, or words. If one starts with understanding the circumstances under which various subsystems of memory are used to memorize faces, one must still study whether the principles learned in the face study generalize to ships. This extra step could be avoided if the aim of the study was made clear from the start.

Designing Instruction Programs We have suggested that cognitive scientists be urged to continue to discover the subsystems of memory. A careful articulation of these subsystems should include a means of evaluating individual differences. Moreover, it should involve the articulation of circumstances under which these subsystems are used, particularly in tasks relevant to the aims of an instructional institution. Such knowledge can be gathered by asking cognitive scientists to keep these aims in mind while designing research tasks. The research program that we have in mind, then, has several stages. For instance, cognitive scientists might be told that trainees must learn the procedure for repairing complex equipment. The scientists first determine what subsystems are needed for such learning and then develop a means of evaluating strengths and weaknesses of these subsystems in individuals. The training institution can then evaluate recruits and assign them to training programs emphasizing their strengths. If people with the requisite strengths cannot be found, then the weaknesses of the potential trainees should be assessed and instruction should be targeted to overcome these weaknesses. This research program has practical consequences—effective evaluation—but is based on research squarely focused on the discovery of the architecture of memory. Such close interaction between the development of evaluative instruments and work on the mental architecture has heretofore been missing and should serve as a major aim for future research.

We still have not said a great deal about instruction. We have emphasized that instruction should be geared to trainees' strengths or weaknesses. That is, if the trainees are bad at image generation and the task you want them to learn involves image generation, then instruction should focus on improving their ability to generate images. How this is done will, of course, depend on what has to be taught. There is, of course, a vast research enterprise concerned with education and instruction, and we will not comment on this line of research. Hirst makes the suggestion that memorization conducted involuntarily is as effective as memorization conducted voluntarily. He illustrates his points with a number of experiments from cognitive psychology but admits that more research needs to be done. We concur. He also suggests that the to-be-learned task be structured so that it is easily memorable. Hirst is arguing that human factors engineers consider memory and training demands and performance demands when designing equipment. Again, we concur.

In making these points, Hirst in essence is claiming that the expectation of a memory test has little bearing on subsequent memory for past events. Hirst asserts that, in teaching people to identify different ships, it does not matter whether you explicitly teach them to identify the features of the ship or embed the feature-learning task in a game that requires feature detection for success. But clearly there must be an interaction between the way something is taught and the way it is tested. For instance, a person who explicitly learns the features of ships may more easily be able to verbalize what these features are than a person who learns the features in the context of a game. Thus, in designing instruction, one must first decide not only what knowledge trainees must acquire but also how they must produce this knowledge.

This distinction between knowing and manifesting this knowledge is similar to the one that we argued was deserving of further research, that is, the one between explicit and implicit knowledge or between procedural and declarative knowledge. Many of the relevant theoretical questions are also similar. In particular, we must determine whether verbalizable knowledge is stored in different locations, processed differently, or represented differently from nonverbalizable knowledge. Information is usually taught explicitly in the classroom, and subsequently testing usually requires verbalization. But it is rarely the case that this knowledge must be verbalized on the job. It just has to be used. For this reason it is important to determine whether explicit instruction aids, interferes, or has no effect on the actual conduct of the job for which the employee is trained. Does verbalizable knowledge help when learning a skill? Will it help people to learn to identify ships if they first learn explicative features of the ships? Or does such verbalizable knowledge function independently of the knowledge that governs the actual identification? Moreover, are there certain conditions under which verbalizable knowledge does interact with the knowledge governing skills? And in drawing the distinction between these two forms of knowledge is it best to draw the line by considering verbalizations or by considering the degree to which the information is explicitly recalled? As far as we know, these questions still await an answer.

Final Considerations and Specific Recommendations

Although lists are usually dangerous affairs, because they are so naked and explicit, several themes do emerge from our present discussion and the chapters themselves. Consequently we conclude this discussion with a brief statement of each of the themes. They essentially point to the direction present research on memory is taking or should take.

1. *There should be a continued development of a theory of the biological basis of memory.* The neuroscience of memory has progressed rapidly in the past few years, and research along the same lines, with the same priorities, should continue to prove fruitful.

2. *To a limited extent, research directed toward finding a drug to improve memory should be done. Emphasis, however, should be given to so-called cognitive enhancers as opposed to specific "memory drugs."* Although it has proven difficult to develop drugs that are specifically targeted to improving memory, drugs that have a more general effect on cognition may prove useful as a memory aid. The relevance of the new cognitive enhancers currently being developed at Squibb and Ayerst should be investigated.

3. *A close tie between the neuroscience of memory and the cognitive science of memory should be developed.* Research should be done to establish whether the various biochemical and physiological models of memory do indeed have anything to do with memory. This research should include behavioral work investigating conditioning and simple learning in vertebrates and should examine more complex aspects of memory, especially what psychologists concerned with human memory call free recall. The research on long-term potentiation should be encouraged. But more important, the connection between long-term potentiation and memory must be better mapped out. Of course, before this mapping can be done convincingly, there must be some general understanding about what memory is or should include. To this end, emphasis should be placed on interdisciplinary work between cognitive scientists and neuroscientists that focuses on essential definitional issues. As far as a unified theory of memory goes, it is probably premature to expect a theory of memory that accounts for both biological and psychological data. Nevertheless, there is enough interest among neuroscientists and cognitive scientists in the goal of a unified theory that discussions among them, no matter how preliminary, should be encouraged. Currently the fields are in the tunnel together, and some workers can see the light at the end of the tunnel. Nevertheless, there is still a great distance to travel, and the light is still quite dim.

4. *The development of connectionist models that have a strong tie to research in neuroscience should be encouraged.* As a better understanding develops of the cognitive science and neuroscience of memory, research should be done to study the means by which various anatomical and physiological constraints on the nervous system can be incorporated into parallel distributed networks. Along the same lines, PDP models should be used to study the function of these anatomical and physiological constraints. At present, PDP models do not have the tight connection with the work of neuroscience that their structure suggests. More must be done to make neuroscience computational and to make computational cognitive science constrained by neuroscience.

5. *Research on the subsystems of memory, especially the subsystems that underlie procedural and declarative memories and implicit and explicit learning, should be encouraged.* As we have argued, it is necessary for both instruction and evaluation to know what the subsystems of memory are. Although much work has been done already, the distinction between procedural and declarative memory (and related distinctions) still needs a strong theoretical foundation if it is to prove useful. Research on the architecture of memory should include standard cognitive psychological experiments, neuropsychological experiments with brain-damaged patients, and experiments using various brain imaging techniques. Model building may involve computer simulations.

6. *Work on the subsystems of imagery should be continued.* Work on imagery has guided much of the research program we have outlined for memory. For this reason imagery research should be vigorously extended. It has reached the stage at which knowledge about the architecture of imagery can now be applied to practical problems, such as the use of imagery in navigation.

7. *Evaluative instruments should be designed in accordance with new knowledge about the architecture of the memory system.* Much of the current means of evaluation is based on outdated knowledge about what the memory system is like. These means should be assessed on the basis of new information in cognitive science, and, where they no longer provide the kind of fine-grained assessment dictated by a subsystems approach, the instruments should be redesigned.

8. *Researchers should be made aware of the kinds of tasks used by the work force. Research should then be directed toward understanding what subsystems are involved in successful performance of these tasks.* It is not enough to discover what the subsystems are. One must also understand how and when they are used. Here it pays to tailor the experimental task to the practical concerns of the relevant funding agency.

9. *The relation between implicit and explicit learning or verbalizable and nonverbalizable knowledge should be understood, and the impact of the research results on instruction should be investigated.* Instruction often involves teaching trainees to verbalize knowledge that is relevant to some task. The ability to verbalize knowledge may have little to do with the actual performance of the task. Research must determine if and when explicit knowledge bears on the performance of a task.

References

Ader, R., and D. de Wied. 1972. "Effects of lysine vasopressin on passive avoidance learning." *Psychonomic Science* 29:46–48.

Atkinson, R. C., and R. M. Shiffrin. 1968. "Human memory: A proposed system and its control processes," in *The Psychology of Learning and Motivation: Advances in Research and Theory,* K. W. Spence and J. T. Spence, eds. New York: Academic Press, vol. 2, 89–195.

Baddeley, A. D. 1976. *The Psychology of Memory.* New York: Basic Books.

Baddeley, A. D. 1982. "Domains of recollection." *Psychological Review* 89:708–729.

Bohus, B. 1977. "Effect of desglycinamide-lysine vasopressin (DG-LVP) on sexually motivated T-maze behavior of the male rate." *Hormones and Behavior* 8:509–517.

Bohus, B., G. Kovacs, and D. de Wied. 1978. "Oxytocin, vasopressin and memory: Opposite effects on consolidation and retrieval processes." *Brain Research* 157:414–417.

Cermak, L. S., N. Talbot, K. Chandler, and L. R. Wolbarst. 1985. "The perceptual priming phenomenon in amnesia." *Neuropsychologia* 23:615–622.

Cohen, N. J. 1984. "Amnesia and the distinction between procedural and declarative knowledge," in *The Neuropsychology of Memory,* N. Butters and L. R. Squire, eds. New York: Guilford Press, 83–103.

Cohen, N. J., and L. R. Squire. 1980. "Preserved learning and retention of pattern-analyzing skill in amnesia: Dissociation between 'knowing how' and 'knowing that.'" *Science* 210:207–209.

Davies, P., and A. J. F. Maloney. 1976. "Selective loss of central cholinergic neurons in Alzheimer's disease." *Lancet* 2:1403.

Deutsch, J. A. 1971. "The cholinergic synapse and the site of memory." *Science* 174:788–794.

Deutsch, J. A., M. D. Hamburg, and H. Dahl. 1966. "Anticholinesterase-induced amnesia and its temporal aspects." *Science* 151:211–223.

Eich, E. 1985. "Context, memory, and integrated item/context imagery." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11:764–770.

Gold, P. E., and R. van Bushkirk. 1976. "Effects of posttrial hormone injections on memory processes." *Hormones and Behavior* 7:509–517.

Gold, P. W., H. Wiengartner, J. C. Ballenger, F. K. Goddwin, and R. M. Post. 1979. "Effects of 1-desamo-8-*d*-arginine vasopressin on behavior and cognition in primary affective disorder." *Lancet* 2:992–994.

Graf, P., and D. L. Schacter. 1985. "Implicit and explicit memory for new associations in normal and amnesia subjects." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11:508–518.

Hasher, L., and R. Zacks. 1979. "Automatic and effortful processes in memory." *Journal of Experimental Psychology: General* 108:356–368.

Hebb, D. O. 1949. *The Organization of Behavior.* New York: Wiley.

Herrnstein, R. J., R. S. Nickerson, M. de Sanchez, and J. A. Swets. 1986. "Teaching thinking skills." *American Psychologist* 41:1279–1289.

Hirst, W. 1982. "The amnesic syndrome: Descriptions and explanations." *Psychological Bulletin* 91:435–460.

Hirst, W. 1987. "Cognitive psychologists become interested in neuroscience," in *The Making of Cognitive Science,* W. Hirst, ed. New York: Cambridge University Press.

Hirst, W., and B. T. Volpe. 1982. "Temporal order judgments with amnesia." *Brain and Cognition* 1:294–306.

Hirst, W., and B. T. Volpe. 1984a. "Automatic and effortful encoding in amnesia," in *The Handbook of Cognitive Neuroscience,* M. S. Gazzaniga, ed. New York: Plenum, 369–386.

Hirst, W., and B. T. Volpe. 1984b. "Encoding of spatial relations with amnesia." *Neuropsychologia* 22:631–634.

Hirst, W., E. A. Phelps, M. K. Johnson, and B. T. Volpe. 1987. "The effectiveness of different retrieval cues in amnesics."

Hirst, W., M. K. Johnson, J. K. Kim, E. A. Phelps, G. Risse, and B. T. Volpe. 1986. "Recognition and recall in amnesics." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 12:445–451.

Huppert, F., and J. A. Deutsch. 1969. "Improvement in memory with time." *Quarterly Journal of Experimental Psychology* 21:267–271.

Jacoby, L. L., and D. Whiterspoon. 1982. "Remembering without awareness." *Canadian Journal of Psychology* 36:300–324.

Karis, D., M. Fabiani, and E. Donchin. 1984. "P300 and memory: Individual differences in the von Restorff effect." *Cognitive Psychology* 16:177–216.

Kintsch, W. 1970. *Learning, Memory and Conceptual Processes.* New York: Wiley.

Krejci, I., and B. Kepkova. 1978. "Effects of vasopressin analogues on passive avoidance behavior." *Activitas Nervosa Superior* 20:11–12.

Legros, J. J., P. Gilot, X. Seron, J. Claessens, A. Adams, J. M. Moeglen, A. Audibert, and P. Bercheir. 1978. "Influence of vasopressin on learning and memory." *Lancet* 1:41–42.

Leshner, A. I., and K. E. Roche. 1977. "Comparisons of the effects of ACTH and lysine vasopressin on avoidance of attack in mice." *Physiology and Behavior* 18:879–883.

Mandler, G. 1980. "Recognizing: The judgment of previous occurrence." *Psychological Review* 87:252–271.

Mayes, A. R., P. R. Meudell, and A. Pickering. 1985. "Is organic amnesia caused by a selective deficit in remembering contextual information?" *Cortex* 21:167–202.

Milner, B., S. Corkin, and H. L. Teuber. 1968. "Further analysis of the hippocampal amnesic syndrome. Fourteen year follow-up study of H.M." *Neuropsychologia* 6:215–234.

Pearce, J. M. S., ed. 1984. *Dementia: A Clinical Approach.* Oxford: Blackwell Scientific.

Peters, S., and R. Richie. 1973. "On the generative power of transformational grammars." *Information Sciences* 6:49–83.

Schacter, D. L. 1987a. "Implicit and explicit learning" *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13:501–518.

Schacter, D. 1987b. "Memory, amnesia, and frontal lobe dysfunction." *Psychobiology* 15:21–36.

Schacter, D. L., J. Harbluck, and D. R. McLachlan. 1984. "Retrieval without recollection: An experimental analysis of source amnesia." *Journal of Verbal Learning and Verbal Behavior* 23:593–611.

Smith, M. L., and B. Milner. 1981. "The role of the right hippocampus in the recall of spatial location." *Neuropsychologia* 19:781–794.

Smith, S. M. 1979. "Remembering in and out of context." *Journal of Experimental Psychology: Human Learning and Memory* 5:460–471.

Squire, L. R. 1982. "The neuropsychology of human memory." *Annual Review of Neuroscience* 5:241–273.

Squire, L. R., and H. P. Davis. 1981. "The pharmacology of memory: A neurobiological perspective." *Annual Review of Pharmacology and Toxicology* 21:323–356.

Summers, W. K., L. V. Majovski, G. M. Marsch, K. Tachiki, and A. Kling. 1986. "Tetrahydroaminoacridine in long-term treatment of senile dementia, Alzheimer-type." *New England Journal of Medicine* 315:1241–1245.

Tulving, E. 1983. *Elements of Episodic Memory.* New York: Oxford University Press.

Warrington, E., and L. Weiskrantz. 1971. "Amnesic syndrome: Consolidation or retrieval?" *Nature* 228:628–630.

Index