2 *Conscious Contents Provide the Nervous System with Coherent, Global Information*

BERNARD J. BAARS

Consciousness is what we might expect in an organ, added for the sake of steering a nervous system too complex to handle itself.
William James (1890)

I. BRIEF OVERVIEW

Renewed interest in consciousness is evident in contemporary cognitive psychology. While there is reasonable agreement on the empirical constraints on a theory of consciousness, there is less consensus on the shape of a theory. This paper specifies a number of empirical constraints, stated as pairs of *conscious-unconscious contrasts*, and suggests a rather small set of principles that can organize these constraints in a rather straightforward way. These principles include the following:

First, the nervous system is viewed as a "distributed" information processing system, in which highly complex and efficient processing is performed by specialized processors in a relatively independent way. These processors may be "data driven"—i.e., they may decide by their own criteria what is worth processing, so that a central mechanism is not needed to exercise executive power over the specialized processors. However, these specialists do require a "central information exchange" in order to interact with each other. This central interchange has been

BERNARD J. BAARS • Department of Psychology, State University of New York, Stony Brook, New York. The development of this paper was supported by the Cognitive Science Program, UCSD Center for Human Information Processing, supported by the Alfred P. Sloan Foundation, and by USPHS Biomedical Research Support Grant 5 S07 RR 07067-12 to the State University of New York at Stony Brook.

R. J. Davidson et al. (eds.), *Consciousness and Self-Regulation*

called a global data base. In operation, a global data base bears a striking resemblance to "working memory."

Conscious contents are thought to reflect a special operating mode of a global data base, namely one in which there is a stable and coherent *global representation* that provides *information* to the nervous system as a whole. This implies that any specialized processor in the nervous system can receive the global information, depending upon its own internal criteria. If the global information is relevant to some specialized processor, it will make predictions regarding it, and if the predictions fail, it will work to reduce the mismatch. That is, it will tend to *adapt* to the global information. Each specialist can also engage in *local* information processing without the benefit of the global data base.

A group of specialists can specify a global representation by cooperating or competing with each other until a consistent representation emerges. Those components of a global representation that are entirely stable can be called a *context* because they will influence other components to organize themselves in a way that fits their constraints. Contexts are not necessarily complete. They can leave a number of degrees of freedom to be filled in. If there are processors in the system that are able to complete a set of stable global constraints (a context), they will tend to do so. This provides a mechanism whereby a global context will be able to pose a problem to the system as a whole, such that distributed processors will act to provide a solution to the problem.

Two sets of empirical constraints fit this analysis. First a set of *capability constraints* are used to support the idea that consciousness reflects an operating mode of a global data base. Secondly, a set of *boundary constraints* show the limits of our experience of some conscious content. Together, these empirical constraints place strong limits on any possible theory.

The theoretical discussion precedes the empirical analysis. Some readers may wish to read the empirical sections (III and IV) before the theory section (II).

II. INTRODUCTION

In recent years, psychologists have begun to approach the issue of consciousness quite pragmatically, largely free from the theoretical obstacles that restricted the scope of earlier attempts. Recent authors agree reasonably well on the phenomena that a theory of consciousness must explain, and it appears that some theoretical consensus is beginning to emerge as well. For example, authors like Posner and Warren (1972), Mandler (1975a,b), and Straight (1979) agree that consciousness is closely associated with short-term memory and the limited-capacity components of the nervous system. Others are impressed by the very large range of the phenomena that are consciously accessible (e.g., Smith, 1969; Hilgard, 1976, 1977; Tart, 1975), and still others relate consciousness to the control of action, executive functions, and automaticity (e.g., Shallice, 1972, 1978; Norman, 1976; Shiffrin & Schneider, 1977; Norman & Shallice, 1980). This paper describes a unified theoretical approach to a large set of phenomena associated with consciousness, and it suggests that a rather simple set of principles can account for many of these phenomena. This approach is further extended in Baars and Mattson (1981), and Baars and Kramer (1982).

Whenever the words *consciousness* or *awareness* are used in this essay, they refer to "consciousness of some content," as opposed to other meanings such as "waking consciousness" (e.g., Natsoulas, 1979). To focus on this sense of consciousness, the empirical constraints shown in the tables are stated in terms of *contrasting pairs of statements* about some conscious content compared with a similar unconscious one (Tables 1 and 2).

For example, one would like to know why we tend to lose consciousness of predictable stimuli, such as the pressure of our clothing, the ambient light, or background noise. Or we would like to know why we are not usually conscious of certain higher-level constancies, such as discourse presuppositions. It is important to note that at times, we can become quite conscious of discourse presuppositions, and we can often make habituated patterns conscious as well.

There are many more examples of our being conscious of something at one point but not another. In the cocktail-party effect, we are largely unconscious of nonattended information, but we can make any stream of information conscious "at will." Further, we are unaware of subthreshold energies, although under the right conditions, these same energy patterns are consciously available. Long-term memory is typically unconscious, but an extraordinary number of things in memory can be made conscious. There are many other such empirical conscious—unconscious contrasts, which generally evoke very little controversy. The job of theory is to find an explanation that will fit all the empirical constraints simultaneously. This approach can lead to an explanation of conscious and unconscious functioning in perception, voluntary action, problem solving, memory retrieval, and attention.

The tables show the empirical constraints on consciousness that this paper attempts to explain. These tables contrast each conscious event or property with a corresponding *un*conscious one, as described above. Table 1 (p. 54) shows a number of empirical *capability constraints*, which contrast the capabilities of conscious and unconscious processes.

Table 2 (p. 56) shows some *boundary constraints*, specifying the synchronic and diachronic boundaries of conscious contents. Together, these sets of constraints serve to exclude a very large number of possible theories of awareness. Note again that the constraints on conscious phenomena described in the tables are entirely empirical: they derive from consistent evidence regarding the experiences of very many people. The claims made in the tables regarding unconscious processes are naturally more inferential, but they seem to generate little disagreement. Thus, the empirical constraints seem to be quite firm.

These empirical constraints fit a small number of rather simple ideas that form the conceptual core of this paper. The following sections explore these ideas in detail and give specific arguments in favor of applying them to a theory of consciousness.

A. The Need for a Central Information Exchange in a Distributed Processing System: Some Theoretical Definitions

Following are some theoretical definitions that will prove useful in exploring this perspective. These definitions are largely conceptual, and the remainder of this chapter shows how they can be used to illuminate the issue of consciousness.

The nervous system resembles in many respects a "distributed" information-processing system, in which highly complex and detailed processing is performed by specialized subsystems in a relatively independent way. There is extensive psychological and neurophysiological evidence for such specialized and relatively autonomous processors (e.g., Geschwind, 1979; Hilgard, 1976, 1977; La Berge, 1974; Shiffrin & Schneider, 1977). Of course, there is also sound evidence for molar, nonspecialized activity (e.g., John, 1976; Mandler, 1975b). In the present point of view, consciousness is that which unites specialized and nonspecialized processes.

Perhaps the key concept is the idea of distributed information processing, a kind of systems organization that is currently being studied intensively in computer science. In such a system, a large number of specialized processors may be "data-driven"; that is, they may themselves retain the processing initiative. The specialists themselves decide by their own criteria what is and what is not of interest, so that a central mechanism does not necessarily have executive power over the special-purpose systems. However, these specialists do require some means of interacting with each other.

It is somewhat difficult conceptually to think of our nervous system as a distributed system because we tend to attach great importance to executive processes. Indeed, we tend to identify with some core executive component of our system, and we tend to believe that this component is in charge of all the others. No doubt, there is some truth to this commonsense view, but the idea of a distributed system can account for a great deal of the evidence. Actually, these ideas are not contradictory. One can speak of executive systems that operate in a domain defined by distributed operations, much as a government may be viewed as an executive operating in a domain defined by multitudinous interactions between individual human beings. But we are in the habit of thinking hierarchically about the nervous system, and it takes some reorganization to think distributively.

If we accept the analogy of the nervous system as a distributed society of specialized processors, some of which try to act as a governmental executive toward the others, then consciousness is much like a publicity organ in this society (see Figure 1). Consciousness seems to be closely associated with a mechanism that permits interaction between specialized, dedicated processors. This mechanism behaves remarkably like a global data base—a "central information exchange" used by artificial intelligence workers to permit any set of specialized processors to cooperate or compete in order to solve some central problem (Kaplan,

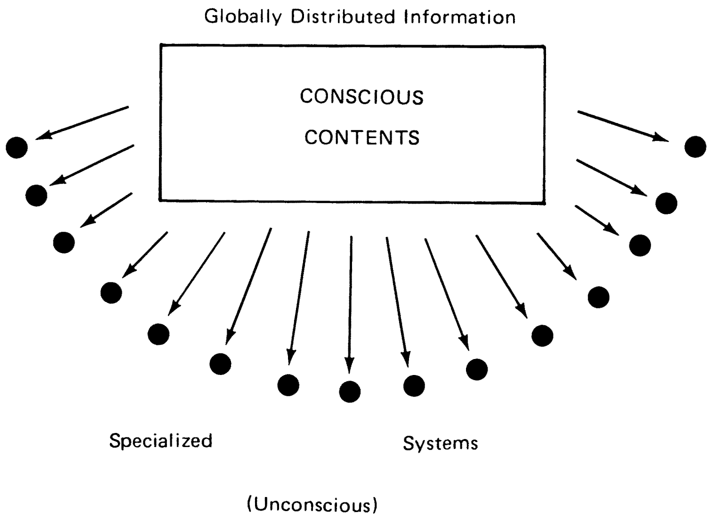

Figure 1. This diagram provides a first approximation to a theory of consciousness. The rectangle symbolizes a global data base, a memory whose contents are available to any specialized processor in the system, and that is in turn controlled by some subset of specialists. Only the global information is conscious—the operation of the specialists is not normally conscious, unless it is displayed by the global data base.

1973; Reddy & Newell, 1974; Erman & Lesser, 1975). A global data base is essentially a memory to which all processors in the system have potential access, and from which all can potentially receive input (see Section IIB). Any representation in the global data base is distributed to the specialized processors much as a television broadcast is distributed to a viewing audience. But unlike an ordinary television audience, some subset of the specialized processors can act on the global data base in return, to propose hypotheses that can then be broadcast to any of the others.

Each of the specialists in the audience can decide on the relevance of the global representation for its own domain. Specialists are assumed to be triggered by the mismatches between the global representation and their own internal representation of their domain. Thus, syntactic specialists are sensitive to linguistic input, and spatial specialists to visual input. However, if the input is already known, if there is no mismatch, the global representation is simply redundant. Further, if the global information is in some other domain that is irrelevant to a specialist, it simply fails to respond. In this sense, any specialist decides by its own criteria whether to process the global representation.

Thus, all specialists are potentially responsive to global input, but the word *global* does not mean that all processors must always accept all global information. It would be useless for a syntactic processor to accept visual input, and vice versa. Rather, *global* implies that a representation is available to *any* processor that has a relevant domain of specialization, providing that it finds the global input to be informative. The word *global* applies whenever we can make this "any" argument, that is, whenever we can plausibly say that some representation is available to any processor, or that any processor can act on it. Generally speaking, the ability to distribute information globally is especially useful if one cannot decide ahead of time which one of the specialized processors needs some information.

What constitutes a specialized processor? This point is rather tricky, because the extent to which a processor behaves autonomously may depend on the task (e.g., Hyde & Jenkins, 1973; Jenkins, 1974). Thus, in riding a bicycle with perfect automatic skill, one may wish to speak of a "bicycle-riding system" with some justification. Certainly, in the beginning of learning to ride a bicycle, this would not be justified. And certainly, when the skill runs into problems—if the right pedal falls off—it must be decomposed, so that control systems for the right foot can learn to behave differently. Yet, we would be in trouble if we had to recruit all the components of bicycle riding each time we leaped on a bicycle; normally, it is better to access the skill as a whole, and for this reason, automatic bicycle riding may be treated as a unified, relatively autonomous specialized processor. This kind of flexibility may discomfit a builder of neat minitheories, who would like to have building bricks that remain stable in all circumstances. But the nervous system probably profits by this potential for flexibility.

As we noted above, global representations are distributed to specialized processors much as a television program is distributed to a large number of viewers. Each viewer has the option of processing or not processing the television program. If the viewer already knows the information, or if it is irrelevant, it may be ignored. If a global representation is neither redundant nor irrelevant to some specialist, it will attempt to *adapt* to the global information. Adaptation is defined as an attempt by the specialized processor to match the global information in its domain, to *reduce the mismatch* that triggered it in the first place. At a physiological level, there is extensive evidence for processes like this: both neurons and systems of neurons *habituate selectively* to input, ceasing to fire when the input is absorbed. But if a change occurs in the habituated pattern (that is, in conditions of mismatch with the previous adaptation), these systems activate again until the new input has become redundant, equilibrium is restored, and they cease firing (Sokolov, 1963; Asratyan, 1965). Note that selective habituation to current input is, in fact, a way in which neurons can store information about any current state of the input. (See Figure 2.)

Any global representation that triggers widespread adaptation can

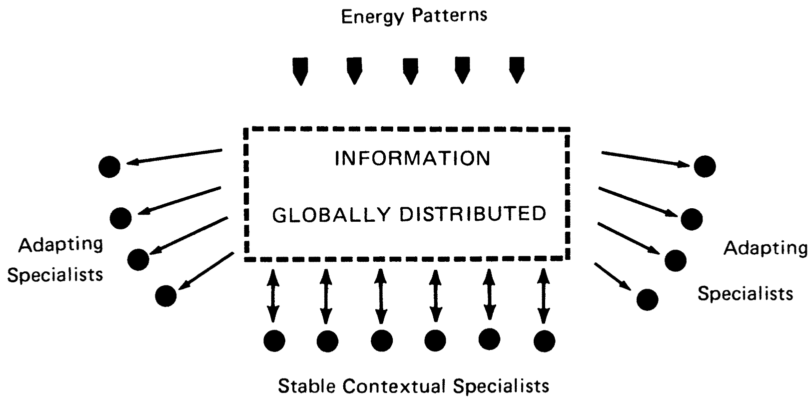

Figure 2. A somewhat more complete diagram, showing that the global information is, in fact, the result of an interaction between incoming energy patterns and a stable set of specialized systems, which provide a context. The resulting global representation, when it has become stable, triggers adaptation in the remaining specialists in the system. Each of the specialists attempts to reduce mismatch between global information in its specialized domain and its own model of the input. Since the global data base exists in a "distributed" system, the processing initiative is left to the specialists themselves.

be said to provide *global information* (MacKay, 1969). Neurophysiological evidence on this point is quite good: new events in the environment, those that are psychologically most likely to be conscious, cause extensive neural activity, far beyond the anatomical pathways of the sensory modality involved (John, 1976). But when this same input is presented repeatedly, the pattern of activity becomes much more localized and limited to special pathways. Redundant input ceases to be global.

Thus, global information is defined as being information to the system as a whole—and here again, one may use the television analogy. If everyone in a country tuned in to a television news program at the same time, one might similarly speak of the content of the news program as providing *global* information to the country *as a whole*. Global information is to be distinguished from *local* information, which is information that is processed within a single specialist without the benefit of the global data base.

B. *The System in Operation*

A number of different processors may cooperate or compete in sending hypotheses to the global data base by acting to confirm or disconfirm global hypotheses until all competition is resolved. If some global hypothesis proposed by one specialized processor is immediately contradicted by another, the hypothesis will have only a transient existence. In order to establish a *stable* global representation, a number of processors must cooperate; that is, they must create a *context*.[1] A context is defined as being a set of stable constraints on a global representation,

[1] The word *context* is often rightly criticized for being theoretically and empirically empty. In this paper, however, it has a number of very specific implications. Theoretically, it is defined as a set of stable, global constraints, which serve to guide and define inputs to the global data base. Empirically, contextual factors are defined as being those factors that can change conscious contents without themselves being conscious. Thus, in any experimental situation, it is clear what factors are and are not part of context, though, of course, this definition does not tell us ahead of time which factors will be contextual in any particular case. Nevertheless, it is quite possible to make some rather strong empirical claims, which include the following: Any conscious content can become a contextual constraint when the system is globally habituated to it. As such, the formerly conscious content will constrain the interpretation of future conscious contents. Components of a context must always be mutually consistent, otherwise competition would occur between them. Incomplete contexts serve to mobilize processors able to complete them, and whenever any component of a context is strongly violated, it will tend to become conscious. All these predictions can be tested in appropriate experimental situations. Thus, the word *context*, as used here, is nonempty both theoretically and empirically.

provided by a set of cooperating processors. That definition implies, of course, that these constraints are consistent with each other; if this were not so, the processors providing the constraints would begin to compete, and the global representation would lose stability. One can view a context as consisting of the set of constraints that is relevant to some particular process to which the system as a whole has *already* adapted.

A set of stable global constraints (a context) is not necessarily complete; there may be many degrees of freedom left, so further constraints can be added to the context. Indeed, some of the constraints that define a context may be changed by incoming information. Finally, incoming information may be able to fill in certain parameter values in dimensions that are specified by the context. Thus, the existence of a stable, coherent set of contextual constraints is not incompatible with the acquisition of new information.

So far, all these definitions have been purely conceptual. Our psychological argument will be that in the nervous system, a stable, global representation becomes conscious *if it provides global information* (Section IV). Conscious representations provide information to the system as a whole, or alternatively, one may say that they trigger adaptation in the system as a whole. Thus, any specialized processor can respond to conscious information relevant to its domain.

It is important to be aware that a global data base is not an executive, though it may be *used by* systems acting in an executive capacity. Indeed, the power and usefulness of distributed information-processing derives from its decentralized organization (Greene, 1972; Turvey, 1977). Again, it is more accurate to compare the global data base with a broadcasting station, which can send information to a vast number of processors, and which can, in turn, be controlled by some processors. In much the same way, a government can control a broadcasting station—but it is the government that acts as an executive, while the broadcasting facility is merely a medium. Consciousness is viewed in this paper as a certain operating mode of this medium, and consciousness can likewise be *used by* processors acting as executives, without itself being an executive.

In many ways, this approach is not new. The global data base has been used by a number of researchers in artificial intelligence (Kaplan, 1973; Erman & Lesser, 1975), and it has a clear similarity to the idea of working memory. Mandler (1975b) has pointed out the close relation of consciousness to short-term memory. Furthermore, Shallice (1972, 1978) has proposed that action systems may dominate the limited-capacity system in much the way that this chapter suggests that consistent sets of processors may dominate the global data base. Others have discussed this kind of a system in a number of different contexts (Lindsay & Norman, 1976; Arbib, 1982; Tart, 1975; Hilgard, 1977; Baars & Kramer, 1982). Nonetheless, the present approach does suggest new ways of viewing the psychological implications of such a system.

Given this perspective, a large amount of evidence falls into line. In Section III, detailed evidence is discussed for associating consciousness with the psychological equivalent of a global data base, and Section IV presents evidence concerning the boundaries of conscious contents.

C. *Advantages of the Global Data Base*

The special advantages and disadvantages of global data bases can be enumerated. First, the advantages:

1. Global information is distributed to all relevant processors, so that if there exists some specialist able to handle it in a fast, efficient way, it can be found immediately.

2. In a problem space that is uncertain or badly understood, a global data base can unite information from many incomplete sources to produce greater certainty than any individual specialist could produce by itself (Erman & Lesser, 1975).

3. A distributed processing system with a global data base would seem to be an ideal learning device. In our version, it is inherently an adapting system since global information is displayed to many different specialists, which are assumed to adapt to those new aspects of the global display that are within their purview. Indeed, we argue in Section IV that we have an *experience* of some event whenever the system as a whole is adapting to a stable, global representation of the event.

4. A global data base can optimize the fundamental trade-off between structure and flexibility. This is a general problem for large systems: on the one hand, it is vital to have specialized, structured solutions available for standard problems, and on the other hand, such structured solutions can be a drawback when the system is faced with really new conditions. In new situations, *flexibility* is at a premium. A global data base permits one to change from a highly structured approach to a highly flexible one. One can have the advantage of structure if the problem is in the province of specialization of a particular processor, along with the advantage of flexibility in choosing between alternative processors, and in the possibility of having a number of specialized processors cooperate in solving some problem.

5. Though a global data base is not an executive, it can be *used by* executive systems. Executive systems can use the global data base to distribute information to control other systems, and to receive feedback from subordinate systems.

6. New processors can be added in a modular fashion. The system can grow without serious disruption, since specialized processors can be added without having to change the previous set of processors. Indeed, the entire configuration may be used to *develop* new specialized processors: as joint information from partial knowledge sources becomes more and more determinate, a new rule-set is defined. This new rule-set may become autonomous and may begin to behave as a specialized processor in its own right. Karmiloff-Smith (1979) has observed a process very much like this in the acquisition of language and other representational systems by children.

7. The same processor may be used in different tasks. For instance, speech perception and speech production have many components in common; thus, it may be that speaking and listening, in fact, involve many of the same processors, which are merely organized differently for speech input and output. Along these lines, Geschwind (1979) claimed that "the primary motor and sensory areas are specialized in the sense that each one is dedicated to a specific function, but the functions themselves are of general utility, and the areas are called on in a great variety of activities."

8. In terms of *content addressability*, one does not need to call any particular processor from a global data base: it is necessary only to present the conditions that the processor finds unambiguously informative. The "name" of any processor is, in a real sense, the information to which it is responsive. This property has considerable advantages and corresponds well to what is known about human cognition (e.g., Norman, 1976).

9. A global data base can solve the need for a mental *lingua franca* (Dennett, 1978). In Western thought, the idea that there must be a common mental code so that one sense modality can communicate with the others goes back to Aristotle's notion of a "common sense." The problem is, of course, that visual information is to some extent unique and different from auditory information, which differs, in turn, from the motor code that controls speech articulation (e.g., Norman, 1976). The unique properties associated with the efficient control of speech are simply not directly translatable into a visual code. Thus, on the one hand, a *lingua franca* is desirable; on the other, it would vitiate the advantages of specialized languages.

A global data base operating in a distributed system obviates this need for a *lingua franca*. Only those systems will respond to a global representation that already "speak the language" of the global representation. Other systems simply don't look at this information, because they do not speak the language. This is perhaps a poor metaphor because the system has not an arbitrary code like natural language, but a content-addressable code more like a semantic network. If the content of the global data base changes, or if in the process of adaptation the specialists change so that they become sensitive to new dimensions of the global information, one might say that they are "learning to speak" the language of the global representation.

There may be one common code able to access all processors, and that is temporal simultaneity. Work on evoked potentials (John, 1976) shows that for a brief time, a prominent new event in any sensory modality reverberates widely throughout the nervous system, far beyond the special anatomical areas and pathways associated with the sensory input. Further, we know from biofeedback research (Schwartz, 1975) that an enormous variety of specialized processors can respond intelligently to those widespread events that affect the whole nervous system simultaneously for a short period of time. It is well established that conscious feedback can be associated with remarkably specific events in the nervous system, events that are presumably controlled by specialized processors. Note that we are not claiming that these events—EEG rhythms, autonomic functions, or single motor neurons—are controlled consciously. In fact, we do not control anything consciously in the literal sense of knowing precisely how we do things (Baars & Mattson, 1981). Rather, the idea is that in biofeedback, specialized systems that control EEG rhythms, or single motor units can independently decide to respond and adapt to widely broadcast information if there is effective temporal simultaneity.

10. A final advantage of this kind of system has been pointed out by Hayes-Roth and Lesser (1976): it consists of flexibility of access to the global data base. Given the same set of specialists, it is possible to experiment with various strategies to control access to the global data base. Certain specialized processors may be given a higher priority than others. This "focus-of-attention" problem has intriguing parallels to the psychological issue of attention. Furthermore, it provides a theoretical mechanism whereby certain potential conscious contents can be avoided, so that in principle, one could incorporate ideas of repression and the dynamic unconscious.

D. Disadvantages of the Global Data Base

No system design is without drawbacks, and the global data base has some obvious ones. For one, it uses a large number of processing resources because all specialists must continually monitor the central information relevant to their domain. Further, global problem-solving is relatively slow, certainly when compared with the fast and efficient action of a specialist that knows how to solve a specific problem. Many different processors must learn to cooperate in order to produce a solution to the global problem. Whenever possible, the global data base should relegate some problem to a specialist (i.e., as soon as a determinate solution is found).

As we point out below (Section IIIA), all these disadvantages have parallels in the "computational inefficiencies" of consciousness: consciousness, too, seems to demand a very great number of resources; it, too, is slow compared with unconscious information-processing. When a conscious solution to a problem is discovered, it is also quickly relegated to unconscious processors: it becomes habituated or automatic.

This paper does *not* claim that consciousness is identical to the operation of a global data base in the nervous system. Rather, it seems that we are conscious of some content when there exists an internal representation that meets three criteria: it must be global, stable, and informative (Section IV).

We turn now to a set of arguments to show that conscious processes are closely associated with a system that acts very much like a distributed system with a global data base.

III. CAPABILITY CONSTRAINTS: ARGUMENTS FOR ASSOCIATING CONSCIOUSNESS WITH A GLOBAL DATA BASE

How do we know when someone is conscious of something? Most obviously we are willing to infer that someone is conscious of an object if the observer can describe it. But consciousness cannot be limited to verbal description—first, because that would confound the construct of consciousness with the evidence that is used to infer it, and second, because that would include talking parrots and computers while excluding babies, aphasics, and ourselves when we are not talking! Adequate measures for any construct *result* as much *from* good theory as they *lead to* it. Thus, any initial definition may need to be changed as the theory is developed. But as a first approximation, we may say that we are willing to infer consciousness when someone can potentially *act discriminatively* toward some internal representation, especially when the internal representation is nonroutine. This inference would include babies, aphasics, at least some animals, and ourselves in our more silent moments; it would exclude even ourselves if we were engaged in extremely routine tasks or were processing very routine representations of things. It would include the objects of perception, which can reasonably be thought of as represented in the nervous system because one can do sensitive memory tests afterwards, for example, and find quite remarkable recognition memory for these perceptions. Also, it would include images among conscious experiences, since all cognitive measures of imagery ask the subject to behave discriminatively toward the image (e.g., Paivio, 1975). Note that a subject does not have to act discriminatively at all times toward some internal representation to be considered conscious of the object. As long as the subject can *potentially* do this, we may consider him or her to be conscious because of this potential.

In practice, we capitalize on the fact that people can consistently answer questions like, "Are you conscious of the words in front of you?" "Before you read this question, were you conscious of the feel of your chair, of the presuppositions of this question, of the breakfast you ate yesterday?" Such questions are answered so consistently by so many people that when it comes to collecting empirical constraints on a theory of consciousness (Tables 1 and 2), no practical obstacles arise.

An initial definition like this may not help us to decide on difficult cases like trance states, automatic writing, cases of multiple personality (e.g., Hilgard, 1976, 1977), or reports of "consciousness without content" (e.g., Naranjo & Ornstein, 1971; Globus & Franklin, 1980). Further, it is possible that people can be conscious of some things so fleetingly as to be unable to answer these questions accurately, as suggested, for example, by the well-known Sperling (1960) phenomenon. But in theory building, as in law, hard cases make bad laws. The great bulk of phenomena we wish to capture can be incorporated in the definition. Perhaps further insight can be gained by considering the more typical phenomena first so that we can then approach these other fascinating questions more intelligently.

Table 1 compares the capabilities and limits of conscious and unconscious processes. These comparisons are purely relative. For example, when we claim that entirely conscious processes are computationally inefficient, this is only in comparison to the evident efficiency

TABLE 1
Capability Constraints on a Theory of Consciousness[a]

Conscious processes	Unconscious processors
1. Computationally inefficient.	Highly efficient in special tasks.
2. Great range and relational capacity.	Limited domains and relative autonomy.
3. Apparent unity, seriality, and limited capacity.	Highly diverse, can operate in parallel, and together have great capacity.

[a] The capability constraints provide one set of conscious-unconscious contrasts that is quite well established and uncontroversial. Yet, these constraints place considerable limits on possible theories.

of unconscious processors. Note, by the way, that the first column refers to unitary conscious process*es*, while the second column refers to multiple unconscious process*ors*. This is itself, of course, a theoretical claim that is supported by some of the arguments made below.

In this section, we discuss each of these points in detail and show how, together, the capability constraints lead to the notion that consciousness reflects the functioning of a global data base in a distributed processing system.

A. Computational Inefficiency

If by *computational efficiency* we mean the ability to compute some standard algorithm quickly and without error, then it is clear that conscious processes lack computational efficiency, while unconscious processors are often remarkably efficient.

1. Some Limitations of Conscious Phenomena

Virtually all authors have noted that the vast preponderance of truly effective human information processing is *not* normally open to awareness (e.g., Miller, 1962; Mandler, 1975a,b; Shallice, 1972, 1978; Straight, 1979). In recent years, we have grown increasingly aware of the awesome complexity of processes needed in the comprehension of even a single sentence, the analysis of a visual scene, or the control of a single motor gesture. People can point to correct and incorrect instances of these events, but they cannot specify the rules involved or explain how they work. This lack of awareness of processing details is universal: it applies to perception and memory retrieval, to most of problem solving, to the control of speech and action, and so on. Where we *are* able to carry out some mental algorithm in a largely conscious way (as in mental arithmetic), the process is often slow and prone to error. Indeed, as conscious processes become more and more proficient, they also become less and less consciously available.

Some people may be tempted to conclude from these observations that consciousness is unimportant or even "epiphenomenal", that is, that it has no *functional* role to play in the nervous system. That is not the solution advanced here. However, it is clear that a functional role for consciousness cannot be found in its computational efficiency.

2. Comparison with Unconscious Processors

The claim that unconscious processors are *highly efficient and specialized for routine tasks* is clearly inferential, but not unreasonable. We may be aware of the sound of a sentence, and of the words, but we are certainly not aware of the fast, complex, and generally error-free processes that mediate between the awareness of sound and the awareness of words. We get some inkling of this complexity when we first acquire some knowledge or skill, but once the new ability is learned to the point of proficiency, it drops out of consciousness. It is then no less complex than before, and it is processed a good deal more efficiently than when it was first acquired—yet, it *seems* easier, presumably because the processing is handled by specialized systems that make little demand on our conscious processing capacity.

There is well-known physiological evidence for independent special-purpose processors in speech, spatial analysis, emotion, metabolic control, and even music (e.g., Geschwind, 1979). On the other hand, there is also a large amount of well-established evidence to show that neural activity in response to new or significant information is extremely widespread, going far beyond the classical neuroanatomical pathways of each sensory modality (John, 1976). (Notice that consciousness of some content is also associated with new or significant events; see Table 2.) But there is no contradiction between localized, dedicated processors and global activities: the tasks performed by the brain require both specialization and global coordination.

It makes sense to suppose that all the truly efficient processors are specialized. The rule systems for spatial analysis are different from those involved in moral judgments, which, in turn, differ from the rules of syntax. Different, yes, but not absolutely autonomous. If we want to understand property law, we must understand how spatial relationships such as boundaries, thoroughfares, and surface features interact with considerations of morality; and, of course, the law has its own syntactic devices to make these considerations explicit. Thus, these three highly cohesive rule systems—morality, spatial relations, and syntax—must

TABLE 2
Boundary Constraints on the Contents of Consciousness[a]

Conscious events	Events that do not become conscious
Synchronic phenomena:	
1. Percepts.	Context required to organize percepts.
2. Input consistent with context.	Input inconsistent with context.
Diachronic phenomena:	
3. Percepts.	Preperceptual processes.
4. Any change in a habituated stimulus.	Habituated percepts.

[a] The boundary constraints suggest that conscious contents are coherent (and hence stable) and also informative.

interact in some cases. More generally, one can claim that for any two apparently separate and internally cohesive rule-systems one can discover a context in which they must interact. Consciousness is presumably involved in facilitating the interaction between such rule systems, until their interaction becomes routine. Once it has become routine, one could speak of a single moral–spatial–syntactic processor with a considerable degree of internal coherence.

3. *Relevance of These Points*

All of these points are consistent with the view that consciousness is associated with a global data base in a distributed processing system. In such a system, purely global processes would be slow and inefficient, because they require the cooperation of a number of otherwise separate processors. By comparison, once some specialist is able to compute a standard algorithm for some known problem, it can do so with very great speed and efficiency.

B. Some Advantages of Conscious Processes

While completely conscious processes are computationally inefficient, the contents of consciousness have extraordinary *range, relational capacity,* and *context sensitivity*. By contrast, unconscious processors by themselves have *limited domains* and are *relatively autonomous* (Table 1).

1. *Conscious Phenomena*

(a) Range. Consciousness seems to participate in all known mental processes at some time. This extraordinary range of conscious contents is one of those "obvious" facts that become puzzling only on further examination. If, for instance, perceptual experience were merely the result of energy transduction, we might explain the vast variety of conscious percepts quite simply: we need only suppose that many different kinds of energy are transduced into some common electrochemical form. But sensation and perception seem to require highly intelligent algorithms, which are probably so complex that each different form of perception demands a specialized set of processors. And if perceiving one kind of object or event demands a specialized processor, how many processors do we need to account for the vast range of percepts, thoughts, feelings, and intentions to which we have conscious access?

Indeed, we seem to have access to an astonishing variety of events in the nervous system. Under optimal conditions, sensory sensitivity approaches the theoretical minimum in which a single retinal receptor may be stimulated by a single photon, or (in the case of hearing) a hair cell may be stimulated by the Brownian motion of molecules in the ear canal. Electrical stimulation of cortical neurons can sometimes be experienced, and by means of conscious feedback, the firing patterns of a single spinal motor unit can come under voluntary control (Schwartz, 1975). This does not mean that the biofeedback subject is conscious of what she or he is doing; rather, it means that the system controlling the motor unit behaves as a specialized distributed processor, able to look at global information. The conscious feedback signal presumably represents such global information. Further, stimulation to which we have become habituated can become conscious (as when we are reminded of background noise, of the effects of gravity, of the feel of a chair, and so on). The unconscious presuppositions of perception, comprehension, and action can become conscious when they are strongly violated (e.g., Offir, 1973; Hornby, 1974). And memories we had long thought lost can reappear in a variety of circumstances (Williams & Hollan, 1981; Hilgard, 1977).

Contrary to widespread opinion, the contents of consciousness are not limited to the so-called higher mental functions. Indeed, there is a striking ability to range far and wide between the most abstract conceptual representations and supposedly more "concrete" sensory-perceptual elements.

What kind of a system could model this extraordinary range of contents? Certainly, a single task-specific processor could not do it. Insofar as a processor is good at some particular task, it is likely to be limited in doing other tasks. No such limitation appears for consciousness.

(b) Relational Capacity. We can relate the contents of consciousness to each other almost without limit. Many decades of research on conditioning indicate that people and even animals have a remarkable capacity for learning arbitrary relationships between different stimuli and between stimuli and responses, though it is easier, of course, to learn nonarbitrary relations.

Although there seems to be evidence that humans can sometimes absorb input without awareness (Dixon, 1971), there is little or no evidence that we can acquire new relationships between inputs without awareness (Brewer, 1974). Indeed, several decades of experimental attempts have shown how difficult it is to demonstrate this. Clearly, the acquisition of new relationships generally involves awareness. Indeed, Smith (1969) has proposed that the capacity to arbitrarily relate any stimulus to any other, or to any response, *is* the criterion-like property of consciousness.

In humans, there is one relational capacity that is of special importance, and that is, of course, our ability to operate on conscious contents with highly specialized linguistic processors. It is important to note again that the existence of some conscious content does not depend on our capacity to express it in words. Verbal report is usually good evidence of some conscious content, but it is not the same as the conscious content.

(c) Context Sensitivity. This is a special kind of relational capacity, and it is of fundamental importance to the present discussion. It differs from the kind of relational capacity that is shown in cases like conditioning. During classical conditioning, for example, we learn that a tone signals the coming of a shock; both of these events are experienced *as* events. They each "stand out as figure from ground." But when we speak of context sensitivity in this chapter, we mean that an experience is affected by factors that do not stand out in this manner. It is a truism that experimental tasks are always affected by a host of variables of which the subjects are not aware. Many of these variables may have been conscious at one time, but they often have their effect long after they have become unconscious. This kind of context sensitivity is basic to the arguments made in this chapter (see Section IV).

While context sensitivity implies that conscious contents are affected by a variety of unconscious factors, one can equally well make the complementary point that conscious contents have widespread *effects* that are themselves not conscious. Both of these observations make sense if we think of consciousness as an operating mode of a global data base in a distributed processing system. Specialized processors sensitive to contextual factors can affect conscious representations, and, in turn, the conscious representations can have a widespread effect throughout the system.

2. *Corresponding Unconscious Phenomena*

Compare the wide range, the relational capacity, and the context sensitivity of conscious events to what we claim to be the relatively *limited domain of unconscious processors and their relative autonomy*. What is the evidence for these claims?

In this view, unconscious processors *by themselves* (i.e., without the intervention of the global data base) are relatively limited and autonomous. The only problem with this claim is that the people we study are generally conscious, so that the limitations of unconscious processors are seldom exhibited overtly. But in the case of involuntary slips of speech or action, we are privileged to see some unconscious processors in a relatively uncontrolled way. Involuntary slips can apparently violate any level of linguistic control (Fromkin, 1973, 1981), and the same observation applies to slips of complex action (Norman, 1981). We can define involuntary slips as those actions that are surprising to the actor: they are not consistent with his or her own previous plans. In extreme cases, slips like this can violate rules that the actor is highly motivated to follow. But in all cases, one can argue that *some specialized rule system, which should have anticipated and prevented the slip, was momentarily decoupled*. Since slips can violate *any* kind of rule system, it follows that any particular system can be momentarily decoupled from the others.

Presumably, if the actor had only a little more time to think about the action, or if his or her attention had been drawn to the relevant dimension, he or she would have anticipated the problem and prevented it (e.g., Baars, Motley, & MacKay, 1975). People are invariably surprised at having made a slip *when they become aware of it*; but this surprise implies that after the slip has been committed, the straying processor must have become coupled again (otherwise the slip would not be recognized as a slip). Indeed, it seems that the process of becoming conscious of an action has the effect of relating the action to its proper context. This again makes sense from the viewpoint of a global data base: if some representation is globally distributed, all the relevant factors can operate on it and respond to it. Conversely, as long as a processor is isolated from the global data base, it can violate rules imposed by other specialized processors.

That unconscious rule systems are relatively autonomous also follows from our frequent inability to exert lasting voluntary control over undesired habits. Most people seem to have automatisms that they would like to eliminate but that seem to be quite autonomous and resistant to external considerations. The more overlearned they are, the less they are conscious, and the harder it may be to exert voluntary control over them. They seem to appear especially when we are consciously distracted or overloaded.

For another example, consider the apparent autonomy of inputs to which we have become habituated. Suppose we have an air conditioner that emits a constant hum of which we rapidly lose awareness. If we need to leave the house and want to shut off the air conditioner (i.e., if the context changes so that we need to operate on the source of the habituated stimulus), we need to *become aware* of the fact that the air conditioner is on. If we fail to bring this fact to awareness, we are likely to leave the air conditioner on, because the habituated representation of this information is not sensitive to changes in context (i.e., to the fact that we are leaving the house).

3. *Relevance of These Points*

Like consciousness, a global data base in a distributed processing system has enormous range, relational capacity, and context sensitivity. By comparison, each specialized processor has a relatively limited domain and is relatively autonomous.

C. *Apparent Unity, Seriality, and Limited Capacity*

1. *Conscious Phenomena*

The impressive relational capacity and context sensitivity of conscious contents should not suggest that there are no limits on our conscious relational ability. However, these limits are of a very interesting kind and seem to depend exclusively on the mutual *informational* compatibility of the mental contents. There are many well-known demonstrations suggesting that we cannot simultaneously experience two mutually exclusive organizations of input (Gregory, 1966; Bransford & Johnson, 1973; Bransford & McCarrell, 1974).

Along very similar lines, there is an extensive lore in the history of science regarding the inability of scientists working within one paradigm to understand a competing paradigm (Kuhn, 1970). Comparable demonstrations of "fixedness' in problem solving go back to Luchins (1942) and Duncker (1945), illustrating the very general fact that a problem cannot be solved if approached within a framework that resists the correct solution. Similarly, Levine (1971) has demonstrated that an extremely simple discrimination task cannot be solved, even under "ideal S-R reinforcement contingencies," if subjects approach it with the wrong set. In the area of conditioning, Dawson and Furedy (1976) have shown that human galvanic skin response (GSR) conditioning will not take place if people are given an explanation of the conditioned-stimulus–unconditioned-stimulus relation that "masks" the contingency between these events. Similar conclusions emerge from work on ambiguous stimuli (e.g., MacKay, 1970) and on brain damage (e.g., Gazzaniga & LeDoux, 1978). If one can safely generalize over a large literature involving such disparate experimental techniques, one might say that any two pieces of information can be consciously related to each other, provided that they can coexist within a single, coherent framework. Facts like the above suggest that conscious organization demands unity, even if the unity is spurious.

The "unity of consciousness" fits quite well with the global-database notion. Any global representation that is not consistent with some processor will quickly encounter competition, so that it will be very unstable. Stable global representations must be coherent at any one time, though they may be contradicted at some later time by another stable global representation.

If the contents of consciousness must be coherent, this requirement also implies that incompatible contents must become conscious *serially*, and that there is a *limited capacity* for competing contents. Thus, the apparent unity, seriality, and limited capacity of conscious contents seem to belong together as a set of related phenomena.

Shaffer (1976) has pointed out that people can do a number of tasks (such as conversing and playing the piano) simultaneously, though we would ordinarily consider them contradictory. It is to be noted that Shaffer's subject is extremely well practiced at these tasks, so that her *conscious* capacity is not likely to be taxed. In general, it appears that tasks that compete when they are new stop competing after enough practice has been gained (i.e., after the tasks are taken over by efficient and relatively autonomous specialized processors). Thus, it may well be that two otherwise "incompatible" tasks can go on in the nervous system, provided that they are not conscious at the same time. Hilgard (1977) provided some spectacular examples of such apparently incompatible processes.

2. *Comparison with Unconscious Processes*

Compared with the unity, seriality, and apparent limited capacity of conscious processes, it appears that unconscious processors are *highly diverse*, that they can operate *in parallel* (unless they need to interact in some way), and that together, the set of specialized unconscious processors has a very great processing capacity. What is the evidence for these claims?

The diversity of processors follows immediately from the idea that they are highly specialized, combined with the incontestable observation that the nervous systems does an enormous number of different things. Parallelism and the idea of a very large unconscious processing-capacity can be viewed in the following way.

(a) Capacity. Consider the physiological facts. There are some 10^{10} neurons in the cerebral cortex alone, firing at an average rate of perhaps 40 Hz (Eccles, 1973). Thus we have about 40×10^{10} events taking place each second, or roughly one-half trillion. This certainly seems like a system with very great capacity—yet we know that conscious capacity seems extraordinarily limited. We can store only 7 plus or minus 2 isolated items even with conscious rehearsal, we can process only one stream of speech at a time, and it takes us at least 100 msec to react to

a conscious stimulus. Unless we assume that most neurons are firing away "epiphenomenally," so that their activity has little effect on psychologically interesting variables, we must somehow reconcile this picture of frenetic activity and relatively ponderous conscious processing. The idea of a global data base appears to reconcile this apparent conflict. As a distributed system, it is no surprise that the great amount of processing activity is not global but is relegated to dedicated processors. Because a global representation requires the cooperation of a number of specialists, it must change much more slowly than any single specialist. This view helps to resolve what seems to be a contradiction between the limited capacity of consciousness and the enormous amount of processing activity that we observe at the same time.

(b) Parallel Processing. The neurophysiology also suggests that "the organization of the brain implies parallel processing" (Thompson, 1976; Anderson, Silverstein, Ritz, & Jones, 1974). However, we well know that conscious processes seem to be quite serial (e.g., Newell & Simon, 1972). How can these different impressions be reconciled?

Consider some of the general properties of parallel information processing. Superficially, it would seem that one could get much more accomplished by parallel than by serial processing. However, parallel processors are restricted if there is some contingency between one process and another—and intelligent processes often involve a series of contingent decisions. In particular, if processors operating in parallel need to interact, there exists a contingency between otherwise independent systems. The result is a bottleneck, which behaves in an apparently serial fashion. These considerations are especially relevant to this discussion, because *a global data base can be viewed precisely as a device that facilitates interaction between otherwise independent, parallel systems.* Marslen-Wilson and Welsh (1978) have provided evidence that some components of speech perception are, in fact, mediated by a parallel-interactive system of this kind.

3. *Summary: How the Global Data Base Fits the Capability Constraints of Table 1*

The relevance of this theoretical metaphor for consciousness is now beginning to emerge. A global data base is not itself an efficient computational device; rather, it permits a multitude of efficient processors to communicate in some commensurable way. Hence, it must be able to display an enormous range of representations, and to relate any two arbitrary representations to each other, so that distributed processors can help to specify the relationship between any two global representations. At any one time, the global data base can display only a single *coherent* content; incoherent representations will swiftly decay because of internal competition. Indeed, a *context* has previously been defined as a stable, coherent set of constraints on a global representation. This need for unity at any one time will make it appear that the global data base has a very limited capacity, and for this reason, competing contents must appear serially. In short, all of the capability constraints of Table 1 very naturally "fall out of" the concept of a global data base.

By contrast, the specialized processors in a distributed system are highly efficient in their special tasks; necessarily, their domains of specialization are also limited and relatively autonomous. They are highly diverse, they can operate in parallel (provided they do not need to interact), and together, they possess great processing capacity.

Thus, there is a close association between consciousness and the kind of systems configuration discussed here, but not an identity. In the following section, arguments are given that conscious representations must be global, and also stable and informative.

IV. THE BOUNDARIES OF CONSCIOUS CONTENTS

What does it mean to experience something? What are the boundary conditions of *conscious contents*? More is known about perception than about any other kind of conscious content, and most detailed examples will be perceptual. However, the conclusions to be drawn from these examples should suggest an approach to such conscious contents as images, which are not dependent on external input. The arguments given here depend on a detailed consideration of the boundary constraints (Table 2, p. 56) which show under what conditions conscious events become unconscious, and vice versa. Following is a short summary of the arguments, followed by more detailed considerations.

Two kinds of boundary conditions may be called synchronic, since they exist at the same time as any conscious experience, though they are not themselves conscious. First, we know that there must be internal representations of the context within which a percept is defined, but that this contextual representation is *not* conscious. In Section II, a *context* was defined as a stable set of constraints on any global representation. Thus, the claim is made that those properties of a global representation that are *entirely constrained* are not conscious. Until we encounter an Ames trapezoidal room, we are not conscious of the fact that we interpret trapezoidal shapes as rectangles in our carpentered world, and until we encounter someone of a different ancestry and culture, we are not conscious of the fact that we have certain assumptions about people's appearance, dress, and mannerisms. These contextual assumptions are clearly used constantly to make sense of the world, but they are not conscious. Second, it is also the case that sensory input that is *not interpretable within the current context* is not conscious. When we hear a word in a meaningful sentence, we are typically conscious of only one meaning, even though a glance at a dictionary should convince anyone that all words have more than one meaning. We are not conscious of any other meaning until we are in the right context. When we are confronted with a foreign language, especially one with a very different phonology, the sounds of the language do not become conscious until the phonology is spelled out, or until we practice it, or until we hear minimally contrastive pairs of words. We can listen to rock music many times without understanding the words, until someone tells them to us; after that, the words seem limpidly clear. The examples can be multiplied indefinitely. All the phenomena discussed under the heading of "Apparent Unity, Seriality, and Limited Capacity" (Section IIIC), represent cases where some input is perfectly interpretable within one context, but incoherent in another one.

These two points can be summed up as follows: Context, taken by itself, is unconscious, and input, taken by itself and in the absence of the appropriate context, is also unconscious. Only when *both* of these conditions exist—that is, when there is input that can be organized within a current context—are we conscious of some percept.

Next, there are two kinds of unconscious representation that exist diachronically, before and after a conscious representation. The first of these diachronic representations involves *preperceptual* processes, which are clearly representational in nature but not conscious. Second, and of very great importance to the present argument, conscious percepts *habituate* rather quickly, if the input remains predictable. Habituation in all its forms is something of a stepchild in the psychological literature on learning and memory. It is often treated as a rudimentary kind of learning, but one that is not very interesting because it is not associative. Perhaps it is a result of "neural fatigue," etc. In any case, habituation is viewed as something of a by-product, of limited interest.

But in the present approach, habituation is thought to be an intelligent matching of input by any neural system—perhaps the fundamental form of learning, therefore. Neural systems stop working when they complete an internal match of the input, as suggested by Sokolov (1963), and this decrement of responding may resemble a kind of fatigue from the outside. But the decrement simply reflects the fact that after the system actively matches the input, the input becomes redundant with respect to the system. Habituation of *consciousness* of some object is treated in this chapter as *global redundancy*, indicating that the system as a whole has adapted to some stable global representation (see also

Nadel, 1981). In the same sense, actions that have become automatic can be considered globally redundant, though they do require *local* information processing (see Section II). The notion of global redundancy emerges most clearly from a consideration of habituation of awareness, but once established, it can be seen in a number of other cognitive phenomena.

What follows now is a detailed presentation of these arguments.

A. Contexts

It is a very general fact that the perception of some object or event requires a stable context or framework (e.g., Minsky, 1975; Asratyan, 1965) and, indeed, that without this stable context a percept cannot be established. A change in the relevant context will produce a different experience or render the input pattern incoherent. This is true whether we speak of linguistic presuppositions, of assumptions made about space in order to process the visual world, or of the set of assumptions that produce fixedness in problem solving.

Contextual constraints of any experience must be represented with great precision, *yet they are invariably unconscious*. That does not mean that we cannot become aware of some contextual assumption, but that as we become aware of it, it ceases to be context and requires some other unconscious assumptions to be comprehended. The context *as* context always escapes our awareness.

One way to bring a contextual constraint to awareness is by violating it strongly (e.g., Offir, 1973; Hornby, 1974). In the Ames trapezoidal room, such a violation occurs when the observer tosses a ball against the wall, and the ball bounces back in an unexpected way. As the observer becomes aware of the trapezoidal shape of the room, his or her contextual assumptions—the stable constraints on his or her conscious representation of the room—go through a transformation. Indeed, one may argue that contextual factors can become conscious *only* when they are challenged.

Cognitive psychologists can often avoid dealing with stable contextual factors by operating *within* a given experimental situation. In other disciplines, such as anthropology, developmental psychology, or the history of ideas, this is not possible. As a result, scholars in these fields are often acutely aware of the effects of changes of context. But it seems that in any situation, the invisible context contains the most powerful factors in the situation, so that this contextual frame is well worth the attention of cognitive psychologists.

Are we defining *context* circularly? Ultimately, we need to work with well-defined theories that make explicit exactly the content of contextual frameworks. Fortunately, over the past 20 years, workers in artificial intelligence have made great strides in the direction of specifying the knowledge needed to understand everyday situations. It seems likely that over the next few decades, this work will yield theories rich enough, and explicit enough, so that we can follow someone around a supermarket, perhaps, talking about being conscious of this and that while taking for granted active but unconscious contextual knowledge.

B. Undefined Inputs

There is a second synchronic boundary to a conscious experience, which becomes especially clear in the case of perceptual input. Consider a perceptual demonstration of a hidden figure, such as the well-known "Dalmatian in the park," which shows a spotted Dalmatian in a shadow-flecked park. Since the entire picture is in black and white, it is initially very difficult to spot the hidden figure. Until the right organization is discovered, it does not become conscious. (It is noteworthy that contextual factors, in the sense defined before, can suggest the correct organization.) However, once proficiency is gained in spotting the hidden figure, the process of spotting the figure becomes fast and efficient to the point where it becomes very difficult to *avoid* seeing the dog. Now the situation is reversed: it becomes difficult to see the spots *as* spots, unless the image of the dog is further obscured (for instance, by turning the picture upside down). Thus the hidden figure is bistable, much like an ambiguous stimulus, but it is nonreversible; once the hidden dog is discovered, the newer, more coherent representation tends to prevail.

This example is really a paradigm case for much of this discussion. The idea of a "hidden pattern" may be generalized to cover a multitude of psychological tasks. For example, in word perception, there is, in a reasonable sense, a "hidden figure" in the pattern of sound. Before one makes a scientific discovery, there is, so to speak, a hidden pattern in the evidence. The reader can no doubt supply further examples of the search for hidden patterns. The analysis of this example may be applied to many similar cases.

The input pattern specifying the Dalmation is unconscious *until* we find the right context—the right set of stable constraints—within which to interpret the input. Thus, we can now state a second very general boundary condition on consciousness of input: available energy patterns are unconscious if they are not defined coherently within the current context. The problem of organizing a pattern in the right way is the problem of finding the right set of stable contextual constraints for that pattern.

The first and second boundary conditions are really different sides of the same coin. The first claims that, to be conscious, any input demands some contextual constraints that are themselves unconscious. The second one states that any pattern is unconscious unless the right framework is available to organize it in a coherent fashion. In the absence of the right context, any input pattern is unconscious. *Thus, context alone is unconscious and input alone is unconscious.* Only when these two unconscious components interact in the right way are we conscious of some event. This is a fundamental claim.

The third and fourth boundary conditions on conscious contents add another requirement. Not only must a conscious representation be coherent and stable, but it must also provide *global information*; that is, it must make a difference to the system as a whole. Arguments for this requirement are discussed next.

C. Global Information and Global Redundancy

Consider the very general phenomenon of habituation of awareness, the third boundary condition. When some stimulus is repeated or continued past a certain point, it is no longer experienced. This effect is as general as perception itself: it occurs in all sensory modalities, with any kind of stimulation.[2] Indeed, at any time, there are a large number

[2] It may be objected that one can have the "same" experience many times without losing awareness of it. For example, one may travel the same road to work each day without a complete loss of awareness of the road. There are a number of answers to this objection. First, we lose complete awareness only of stimuli that are entirely predictable and, in particular, of stimuli with entirely predictable temporal properties. Very often, we do not adapt completely to some conscious content; we simply shift to a different content. Nevertheless, it is easily conceivable that we shift away if some particular dimension of the content has been absorbed, especially if that dimension is relevant to a current task. Second, we are not merely exposed to the same road each day; we interact with it, so that often we can voluntarily override the tendency to lose awareness of the information. Third, whenever some change takes place in a well-known environment to which we have become habituated, the change must be integrated into a larger set of events. If someone is exposed to a regularly repeated burst of noise, for instance, and we change only one aspect of the noise burst, all the other properties of the stimulus must be reevaluated. Thus, if the onset ramp of the habituated noise is changed, the subject will not become aware of the onset ramp *in abstracto* but will become aware of the whole noise burst. Similarly, if one aspect of the road to work changes, other aspects must become conscious as well. In sum, our continued awareness of routine events does not constitute a counterexample to the claim made here.

of predictable energy patterns impinging on us, from gravity to the ambient light, sound, and temperature, to the pressure of our clothing. All these energy patterns are typically unconscious. It may be that the visual system is especially protected against premature habituation to constant inputs by means of physiological nystagmus (the constant high-frequency tremor of the eyes), which causes light edges in the visual field to fall always on a slightly different part of the retina. Without this feature, awareness of the visual input is lost within seconds. Clearly in all other sensory modalities, awareness of some redundant input tends to habituate quite rapidly.

Habituation of neural structures occurs at all levels, from single cells to complex structures. As we noted above, habituation has not been thought to be a full-fledged kind of learning in the literature on learning and memory, even though, as Sokolov (1963) pointed out, habituation cannot be a form of fatigue because it is stimulus-specific, and because habituated animals will dishabituate to the *absence* of the repeated stimulus. That is, dishabituation is a response not to energy input but to new information. And if dishabituation (orienting, etc.) is a response to information, it is not too great a leap to suppose that habituation is a response to redundancy (Asratyan, 1965; Nadel, 1981). In the present perspective, it is assumed that all specialized processors attempt to model global input that is relevant to them, and that having done so, they cease responding to this input. Thus, habituation for these processors is a sign that learning has occurred; that is, the input matches the local representation sufficiently so that no further adaptation is required. In this view, *habituation of awareness* means that the system as a whole has adapted to the input, so that the input no longer provides global information. (Recall that *global information* has been defined as a global representation that triggers adaptation in the system as a whole, so that any relevant processor can adapt to it.) Thus, habituation of awareness to any input may be considered a sign of global redundancy, which is simply the absence of global information.

But now we can go one important step further. By definition (see Section II), a context is a set of stable, global constraints—constraints to

---

What about the impact of significant stimuli on habituation? It seems obvious that awareness of significant events is lost more slowly than awareness of insignificant ones. In some cases, like chronic pain, one could indeed maintain that awareness is never permanently lost. It seems useful to treat the effect of significance in informational terms. It is highly plausible to think that significant events require more adaptation throughout the system; thus, in a very strict sense, significant events are more "informative" than other events. Significant changes must propagate more widely throughout the system before adaptation can take place. Presumably significant events also demand more problem solving (Baars & Kramer, 1982) before the system achieves complete adaptation.

which the system as a whole has adapted. But this is, of course, equivalent to saying that these constraints are globally redundant. Anything that is globally redundant can therefore become part of a new context, able to affect the way relevant new stimuli will be experienced. We may say, then, that conscious representations that are lost from awareness because of habituation do not disappear: they continue to provide a context within which future related representations are defined.

This is not surprising if we consider well-established facts on gross perceptual readaptation (e.g., Köhler, 1962). Consider an everyday example. When we first step onto a small sailboat, we are very much aware of the movements of the boat, but most of us adapt fairly rapidly, so that the movements of the boat become highly predictable and are lost from awareness. They become globally redundant, in the sense suggested above. What previously constituted information has habituated and is now a part of the context. It is easy to show that this habituated information is not completely lost; it merely becomes the framework within which new information is defined. Thus, on returning to dry land, what is usually in the background now becomes information. That is to say, it now seems as if the *world* is reeling drunkenly: we make false predictions regarding our orientation to gravity and attribute the source of this information not to ourselves but to the world. What was globally redundant (and unconscious) at sea now becomes globally informative and hence conscious, until we once again adapt to land, and the spatial context once again fades into the background.

Thus, habituation of conscious contents can create new contextual constraints that can affect the way conscious information is structured. This is true not only of relatively gross properties, such as our orientation to gravity. Rather subtle properties of our perceptual experience can also be affected by a habituation phenomenon like selective adaptation (e.g., Eimas & Corbit, 1973). In the remainder of this chapter, the idea that habituated conscious contents can create the context for future conscious contents will be considered a general and very important property of the system.

Note that even though a certain process may be globally redundant, it can still require some *local* information-processing. Consider automaticity, which is the counterpart of habituation in the case of a proficient skill (LaBerge, 1974). As we walk around the world, we are largely unaware of the fast, complex, and subtle details of balancing and moving. That is, the action of walking is largely globally redundant, in the sense defined above. Yet we cannot claim that these fast-moving details of walking are nowhere computed; rather, we may say that they do not require global information-processing, because they are essentially routine and predictable at a global level. This is presumably not true for babies just beginning to walk, nor is it true for bed-ridden individuals who are readapting to walking, nor for people just stepping off a small sailboat. For all these individuals, walking requires global information-processing.

Again, the question may be raised whether we are defining *information* and *redundancy* in a circular way, by reference to the phenomena they are supposed to explain. We know that when we can control stimulation, we can repeat a stimulus over and over again, and it will disappear from consciousness. Further, we know that if we change only one dimension of the stimulus—its amplitude, energy envelope, quality, or temporal parameters—the stimulus will become conscious again. These are not circular claims if we have experimental control. But this explanation is still far from satisfactory. We cannot predict the informativeness of a new stimulus in many situations, and until we have adequate theoretical representations of these situations, we will not be able to specify what is informative and what is not. Work in artificial intelligence does seem to be moving rapidly in the direction of such explicit representations for everyday situations.

So far, the argument applies to entirely predictable inputs, such as the ambient light or temperature. However, the notion of global redundancy can also apply to highly predictable *components* of input. For example, when we hear a series of paraphrases of a single sentence, there is little doubt that we will begin to ignore the meaning rather quickly, even though the physical input is continuously changing. Indeed, it will be difficult to attend to the meaning at all: it has become globally redundant. More generally, it seems likely that as soon as we fully understand the meaning of any sentence, any repetition of the meaning will be globally redundant. One may interpret the "click" of comprehension as that moment in which the meaning of a sentence becomes globally redundant, so that now it can be used to interpret new, incoming information. Thus, the notion of global redundancy can be extended beyond the pure case of completely predictable input.

The idea that globally redundant constraints are not conscious adds a very important qualification to the discussion so far. Such things as predictable stimuli and automatic skills must be represented in the nervous system in a coherent fashion, but of course they are not conscious. But this means that the coherence of a representation is a necessary but not a sufficient condition for the representation to be conscious. A conscious representation must be new or *globally informative* as well as coherent. We can now add this to the set of boundary conditions specified so far and claim that *conscious experience of some content involves an inter-*

action between an energy pattern and contextual constraints, resulting in a coherent representation that provides global information to the system. This is one of the major conclusions of this chapter.

We now have three sets of constraints along with a reasonable theoretical interpretation. Conscious contents always seem to involve an interaction between some energy pattern and a set of contextual assumptions, so that the interaction results in a coherent representation. But if this coherent representation provides no news, if it is not informative in some sense, it is not conscious. This observation suggests that after habituation of awareness, contextual processors have adapted to the representation so that the news has been absorbed and has indeed become a part of the system itself.

D. Preperceptual Processing

Consider now the fourth boundary condition of perception. The idea is widely accepted that input information is preprocessed for a few tenths of a second before it becomes conscious (e.g., Neisser, 1967). Preperceptual processing is usually viewed as a kind of hypothesis testing, in which many different hypotheses are brought to bear on the problem of representing the input. Hypotheses are representations, of course, and we must explain why this kind of input representation is not conscious. This is the fourth boundary condition for conscious experience, and it is "diachronic" because it refers to a stage of stimulus representation that comes before the conscious experience.

Indeed, in practice, the global-data-base configuration has been used primarily for the kind of hypothesis testing that presumably takes place before the establishment of a percept (Erman & Lesser, 1975). The global data base is useful when the processors needed to analyze the input are unknown, so that *any* relevant processor can be brought to bear on any global hypothesis. Because the global data base is very useful in broadcasting information to relevant but unpredictable processors, it seems plausible to assume that the nervous system makes use of something like a global data base during preperceptual processing. But this means that there must be global hypotheses that are *not conscious.* Now explain the difference between conscious hypothesis-testing and the unconscious hypothesis-testing that presumably takes place before perception?

That explanation is really already available. It was previously shown that input patterns not defined within the current context are not conscious. But preperceptual processes involve precisely a set of hypotheses that are undefined within the current context, because they are unstable and mutually competitive. By the time they cooperate sufficiently to establish a coherent context, they become conscious. It is therefore not surprising that the preprocessing of input is not conscious.

If we look at the hidden-figure example again, we can observe this process in slow motion. In the beginning, it is difficult to find the Dalmatian at all, but given the right context (which can be induced by external hints, etc.), the input can be interpreted in the right way. On repeated exposure, it is indeed difficult to avoid seeing the dog, to access the less-coherent interpretation of the input. But even when the pattern is analyzed quite efficiently, we can presume that preprocessing still takes place. It is then no longer conscious, presumably because the processing has speeded up and the lower-level hypotheses are fleeting and unstable.

This point is further supported by a set of examples that show the opposite effect. That is to say, what would happen if we were to *slow down* preperceptual processing? Would the previously unconscious process of hypothesis testing become conscious? This has in fact been tried (Norman, 1976; Kolers, 1978; Bruner & Potter, 1964). For instance, when we read a sentence such as this one, fast and efficient unconscious processes take care of letter and word recognition. But try reading a sentence upside down: suddenly we begin to test *conscious* hypotheses about letters and words. Similarly, when we see an unfocused slide that is gradually brought into focus, we begin "spontaneously" to generate conscious hypotheses. These conscious hypotheses are probably very similar to the unconscious ones that occur preperceptually, although they may be more coherent than the unconscious hypotheses.

There is thus a rather thin dividing line between unconscious hypothesis-testing and conscious hypothesis-testing. It is not clear whether the global representation becomes conscious because it is coherent or because it is stable, because stability and coherence tend to covary perfectly. Certainly, in the model developed so far, specialists that compete with each other can display a global hypothesis for only a very short time before it is contradicted. Therefore, any hypothesis that is coherent will also be stable, and vice versa. Thus, it is safer to restate the conclusion reached above as follows: "Conscious experience of some event always involves an interaction between some energy pattern and a contextual framework that results in a coherent *and stable* representation that provides global information to the system."

This is the fundamental perspective on conscious contents that is advanced here. It seems to account for the empirical constraints in an economical fashion. The model is still a thinking tool—not a permanent position to cling to. But it appears to be both adequate and economical. For instance, it is theoretically pleasing to see that the four general boundary conditions actually reduce to only two properties of conscious contents. First, preprocessing and undefined input turn out to be unconscious for the same reason. Preperceptual processes are unconscious because they lack a stable and coherent context, and similarly, undefined inputs lack such a context. Further, context and highly predictable input patterns are both unconscious because they both involve stable, global representations to which the system as a whole has *already adapted*. Thus, four classes of phenomena reduce to two theoretical terms: events are unconscious if they are undefined in the current context, or if they are so stable as to be part of the context. This theoretical economy is encouraging and suggests that the analysis is on the right track.

V. Summary and Conclusions

The metaphor of a global "information exchange" in a distributed processing system helps to explain a number of psychological phenomena. It appears that we are conscious of something when there is an interaction between input and context, resulting in a stable and coherent global representation that provides information to the nervous system as a whole. this description fits the empirical constraints of Tables 1 and 2 and makes a good deal of functional sense as well. When we are conscious of something, we are adapting to it in a global way.

A. Extensions of the Theory

Because of limitations of space, we can only suggest several extensions of this theoretical framework to incorporate further empirical constraints (Baars, 1980; Baars & Mattson, 1981). For example, the theory can incorporate the classic "insight" phenomena in problem solving, and it suggests a role for the conscious components of intentional action. These points are briefly summarized.

In both spontaneous problem-solving and intentional action, a global context serves to guide specialized processors able to complete the context. Although we are not directly conscious of this global context, any of its components can become conscious when contradicted in some way, by either internal or external influences. In problem solving, one accumulates a set of constraints that are at first fully conscious, and as the system adapts to these constraints, they become components of the problem context. As part of a context, these constraints are, of course, not conscious unless they are violated in some way.

An *intention* may be considered a special kind of problem context, one that serves as a global goal to mobilize and organize a large, diverse set of action specialists to prepare and execute an action. Further, the intention has a timing component that permits the action to run off when it is ready.

When there is a conflict between different intentions, or when some action violates a component of its intention, the appropriate component of the intention seems to become fully conscious. All these cases involve changes in the intention context and therefore fit our previous characterization of the conditions under which components of a context become conscious. The "internal monologue" that we often use to command ourselves can be viewed as a way in which a processor able to broadcast a conscious command can trigger the creation of an intention by other processors. The intention then can serve to prepare and execute an action. It is to be noted that such conscious commands are never as complete as the intention: they seem to involve only what is new and different about an intention. Most of the content of the command is tacitly understood, just as the bulk of our communications with each other are tacitly understood (viz., Baars & Mattson, 1981).

Finally, we can deal with the closely related issue of attention. Attention involves a set of systems able to *select and maintain* some particular conscious content, either voluntarily or involuntarily. In voluntary control of attention, we may give a conscious command that triggers an intention that can control future conscious contents. In the involuntary case, specialized processors act to control the contents of consciousness. The experience of mental effort is thought to result from conflict between voluntary and involuntary means of controlling consciousness. Finally, the ideas of the dynamic or "affective" unconscious fit in naturally with the notion of attention, with the difference that some specialized processors may exercise control of access to the blackboard in order to avoid certain conscious contents, while others may seek to display certain other contents.

B. Conclusion

A great deal of work needs to be done to expand and clarify this approach, and to test it for theoretical adequacy and consistency. There are very many empirical implications that have not been discussed in this paper and that must be considered in detail elsewhere. Nevertheless, the theory sketched here fits a large number of facts about consciousness.

No theory at this stage can be more than a thinking tool, to be falsified and changed as our understanding grows. If the present paper serves to define some of the issues with more precision, and if it helps to develop a vigorous and pointed debate about them, a large part of its purpose will have been achieved.

ACKNOWLEDGMENTS

The author gratefully acknowledges Michael A. Wapner (CSULA) for a number of the ideas presented in this paper, and Donald A. Norman (UCSD) for his welcome insistence on brevity and clarity, which helped to improve the paper immensely. Others whose comments were most helpful include Robert Buhr (UCSD, Cognitive Science Program), Lynn Nadel (UC Irvine), Michael Posner (University of Oregon), David Cross (SUNY Stony Brook), and Michael T. Motley (UC Davis). The comments of Robert Monk (Chicago), Richard J. Davidson (SUNY Purchase), Daniel Dennett (Tufts University), and Walter Ritter (Einstein College of Medicine) were most helpful as well.

REFERENCES

ANDERSON, J. A., SILVERSTEIN, J. W., RITZ, S. A., & JONES, R. S. Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review*, 1973, *84*(5), 413–451.

ARBIB, M. A. Perceptual structures and distributed motor control. In V. B. BROOKS (Ed.), *Handbook of physiology* (Vol. 3). American Physiological Society, 1982.

ASRATYAN, E. A. *Compensatory adaptations, reflex activity, and the brain*. Oxford: Pergamon Press, 1965.

BAARS, B. J. *What is the role of consciousness*? Unpublished manuscript, 1980.

BAARS, B. J., & MATTSON, M. E. Consciousness and intention: A framework and some evidence. *Cognition and Brain Theory*, 1981, *4*(3), 247–263.

BAARS, B. J., & KRAMER, D. N. Conscious and unconscious components of intentional control. *Proceedings of the Fourth Annual Cognitive Science Conference*. Ann Arbor: University of Michigan Program in Cognitive Science, 1982.

BAARS, B. J., MOTLEY, M. T., & MACKAY, D. G. Output editing for lexical status in artificially elicited slips of the tongue. *Journal of Verbal Learning & Verbal Behavior, 14*, 1975, 382–391.

BRANSFORD, J. D., & MCCARRELL, N. S. A sketch of a cognitive approach to comprehension: Some thoughts about what it means to comprehend. In W. B. WEIMER & D. S. PALERMO (Eds.), *Cognition and the symbolic processes*. New York: Wiley, 1974.

BRANSFORD, J. D., & JOHNSON, M. K. Consideration of some problems of comprehension. In W. G. CHASE (Ed.), *Visual information processing*. New York: Academic Press, 1973.

BREWER, W. F. There is no convincing evidence for operant or classical conditioning in adult humans. In W. B. WEIMER & D. S. PALERMO (Eds.), *Cognition and the symbolic processes*. Hillsdale, N.J.: Lawrence Erlbaum, 1974.

BRUNER, J. S., & POTTER, M. C. Interference in visual recognition. *Science*, 1964, *144*, 424–425.

DAWSON, M. E., & FUREDY, J. J. the role of awareness in human differential autonomic classical conditioning: The necessary-gate hypothesis. *Psychophysiology, 13*(1), 1976, 50–53.

DENNETT, D. C. Toward a cognitive theory of consciousness. In D. C. DENNETT (Ed.), *Brainstorms*. San Francisco, Calif.: Bradford Press, 1978.

DIXON, N. F. *Subliminal perception: The nature of a controversy*. London: McGraw-Hill, 1971.

DUNCKER, K. On problem solving. *Psychological Monographs*, 1945, *58* (5, Whole No. 270).

ECCLES, J. C. Brain, speech, and consciousness. *Naturwissenschaften*, 1973, *60*, 167–176.

EIMAS, P. D., & CORBIT, J. Selective adaptation of linguistic feature detectors. *Cognitive Psychology*, 1973, *4*, 99–109.

ERMAN, L. D., & LESSER, V. R. A multi-level organization for problem solving using many, diverse, cooperating sources of knowledge. *Proceedings of the 4th International Joint Computer Conference*, Georgia, USSR, 1975, 483–490.

FROMKIN, V. A. (Ed.). *Speech errors as linguistic evidence*. The Hague: Mouton, 1973.

FROMKIN, V. A. (Ed.). *Errors of speech and hearing*. New York: Academic Press, 1981.

GAZZANIGA, M., & LEDOUX, J. P. *The integrated mind*. New York: Plenum Press, 1978.

GESCHWIND, N. Specializations of the human brain. *Scientific American*, 1979, *241*(3), 180–201.

GHISELIN, B. *The creative process*. New York: Mentor, 1952.

GLOBUS, G., & FRANKLIN, S. Prospects for a scientific observer of perceptual consciousness. In J. M. DAVIDSON & R. J. DAVIDSON (Eds.), *The Psychobiology of consciousness*. New York: Plenum Press, 1980.

GREENE, P. H. Problems of organization of motor systems. In R. ROSEN & F. M. SNELL (Eds.), *Progress in theoretical biology* (Vol. 2). New York: Academic Press, 1972.

GREGORY, R. L. *Eye and brain: The psychology of seeing*. New York: McGraw-Hill, 1966.

HAYES-ROTH, F., & LESSER, V. R. Focus of attention in the Hearsay II speech understanding system. *Proceedings 5th International Joint Computer Conference*, Cambridge, Mass., 1976, 27–35.

HILGARD, E. R. Neodissociation theory of multiple cognitive control systems. In G. E. SCHWARTZ & D. SHAPIRO (Eds.), *Consciousness and self-regulation: Advances in research* (Vol. 1). New York: Plenum, 1976.

HILGARD, E. R. *Divided consciousness: Multiple controls in human thought and action*. New York: Wiley, 1977.

HORNBY, P. A. Surface structure and presupposition. *Journal of Verbal Learning Verbal Behaviour*, 1974, *13*, 530–538.

JAMES, W. *The principles of psychology*. New York: Holt, 1890. (Reprinted by Dover, New York, 1950.)

JENKINS, J. J. Remember that old theory of memory? Well, forget it! *American Psychologist*, 1974, *29*(11), 785–95.

JOHN, E. R. A model of consciousness. In G. E. SCHWARTZ & D. SHAPIRO (Eds.), *Consciousness and self-regulation: Advances in research* (Vol. 1). New York: Plenum Press, 1976.

KAPLAN, R. M. A general syntactic processor. In R. RUSTIN (Ed.), *Natural language processing*. New York: Algorithmics Press, 1973.

KÖHLER, I. Experiments with goggles. *Scientific American*, 1962, *206*, 62–75.

KUHN, T. S. *The structure of scientific revolutions* (2nd ed.). Chicago: University of Chicago Press, 1970.

LABERGE, D. Acquisition of automatic processing in perceptual and associative learning. In P. M. A. RABBITT & S. DORNIC (Eds.), *Attention and Performance* (Vol. 4). London: Academic Press, 1974.

LEVINE, M. Hypothesis theory and nonlearning despite ideal S-R reinforcement contingencies. *Psychological Review*, 1971, *78*, 130–140.

LEWIS, F. L. Semantic processing of unattended messages using dichotic listening. *Journal of Experimental Psychology,* 1970, *85,* 225–228.
LINDSAY, P. H., & NORMAN, D. A. *Human information processing* (2nd ed.). New York: Academic Press, 1976.
LUCHINS, A. S. Mechanization in problem solving. *Psychological Monographs,* 1942, *54*(6, Whole No. 248).
MACKAY, D. G. Mental diplopia: Towards a model of speech perception at the semantic level. In G. B. FORES D'ARCAIS & W. J. M. LEVELT (Eds.), *Advances in psycholinguistics.* Amsterdam: North Holland, 1970.
MACKAY, D. M. *Information, mechanism and meaning.* Cambridge, Mass.: MIT Press, 1969.
MANDLER, G. Consciousness: Respectable, useful, and probably necessary. In R. SOLSO (Ed.), *Information processing and cognition: The Loyola Symposium.* Hillsdale, N.J.: Lawrence Elrbaum, 1975.(a)
MANDLER, G. *Mind and emotion.*New York: Wiley, 1975.(b)
MILLER, G. A. *Psychology: The science of mental life.* New York: Harper & Row, 1962.
MINSKY, M. A framework for representing knowledge. In P. H. WINSTON (Ed.), *The psychology of computer vision.* New York: McGraw-Hill, 1975.
NADEL, L. Cognitive and neural maps. In P. JUSCYZK & R. KLEIN (Eds.), *Essays on the nature of thought: In honor of D. O. Hebb.* Hillsdale, N.J.: Lawrence Erlbaum, 1981.
NARANJO, C., & ORNSTEIN, R. E. *On the psychology of meditation.* New York: Viking, 1971.
NEISSER, U. *Cognitive psychology.* New York: Appleton-Century-Crofts, 1967.
NEWELL, A., & SIMON, H. A. *Human problem solving.* Englewood Cliffs, N.J.: Prentice-Hall, 1972.
NORMAN, D. A. *Memory and attention: An introduction to human information processing* (2nd ed.). New York: Wiley, 1976.
NORMAN, D. A. Categorization of action slips. *Psychological Review,* 1981, *88*(1), 1–15.
NORMAN, D. A., & SHALLICE, T. Attention & action: Willed and automatic control of behavior. Unpublished paper, UCSD, La Jolla, Calif., 1980.
OFFIR, C. E. Memory for presuppositions of relative clause sentences. *Journal of Verbal Learning and Verbal Behavior,* 1973, *12*(6), 636–643.
PAIVIO, A. Perceptual comparisons through the mind's eye. *Memory and Cognition,* 1975, *3,* 635–647.
POSNER, M., & WARREN, R. E. Traces, concepts and conscious constructions. In A. W. MELTON & E. MARTIN (Eds.), *Coding processes in human memory.* Washington, D.C.: Winston, 1972.
REDDY, R., & NEWELL, A. Knowledge and its representation in a speech understanding system. In L. W. GREGG (Ed.), *Knowledge and cognition.* Potomac, Md: Lawrence Erlbaum, 1974.
SCHWARTZ, G. E. Biofeedback and physical patterning in human emotion and consciousness. *American Scientist,* 1975, *63,* 314–324.
SHAFFER, L. H. Intention and performance. *Psychological Review,* 1976, *83*(5), 375–393.
SHALLICE, T. Dual functions of consciousness. *Psychological Review,* 1972, *79*(5) 383–393.
SHALLICE, T. The Dominant action system: An information-processing approach to consciousness. In K. S. POPE & J. L. SINGER (Eds.), *The stream of consciousness: Scientific investigations into the flow of experience.* New York: Plenum Press, 1978.
SHIFFRIN, R. M., & SCHNEIDER, W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review,* 1977, *84,* 127–190.
SMITH, K. *Behavior and conscious experience: A conceptual analysis.* Athens, Ohio: Ohio University Press, 1969.
SOKOLOV, E. N. *Perception and the orienting reflex.* New York: MacMillan, 1963.

SPERLING, G. The information available in brief conscious presentations. *Psychological Monographs*, 1960, *74* (Whole No. 11).

STRAIGHT, S. H. Consciousness as a workspace. *SISTM Quarterly*, 1979, *1*, 11–14.

TART, C. *States of consciousness.* New York: Dutton, 1975.

THOMPSON, R. F. The search for the engram. *American Psychologist*, 1976, *31*, 209–227.

TURVEY, M. T. Prelimenaries to the theory of action with reference to vision. In R. SHAW & J. D. BRANSFORD (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology.* Hillsdale, Md.: Erlbaum, 1977.

WILLIAMS, M. D., & HOLLAN, J. D. The process of retrieval from very long-term memory. *Cognitive Science*, 1981, *5*(2), 87–119.