

Twelve Issues for Cognitive Science

DONALD A. NORMAN

University of California, San Diego

I am struck by how little is known about so much of cognition. One goal of this paper is to argue for the need to consider a rich set of interlocking issues in the study of cognition. Mainstream work in cognition—including my own—ignores many critical aspects of animate cognitive systems. Perhaps one reason that existing theories say so little relevant to real world activities is the neglect of social and cultural factors, of emotion, and of the major points that distinguish an animate cognitive system from an artificial one: the need to survive, to regulate its own operation, to maintain itself, to exist in the environment, to change from a small, uneducated, immature system to an adult, developed, knowledgeable one.

Human cognition is not the same as artificial cognition, if only because the human organism must also be concerned with the problems of life, of development, of survival. There must be a regulatory system that interacts with the cognitive component. And it may well be that it is the cognitive component that is subservient, evolved primarily for the benefit of the regulatory system, working through the emotions, through affect.

I argue that several concepts must become fundamental parts of the study of cognition, including the roles of culture, of social interaction, of emotions, and of motivation. I argue that there are at least 12 issues that should comprise the study of cognition, and thereby, the field of Cognitive Science. We need to study a wide variety of behavior before we can hope to understand a single class. Cognitive scientists as a whole ought to make more use of evidence from the neurosciences, from brain damage and mental illness, from cognitive sociology and anthropology, and from clinical studies of the human. These must be accompanied, of course, with the study of language, of the psychological aspects of human processing structures, and of artificially intelligent mechanisms. The study of Cognitive Science requires a complex interaction among different issues of concern, an interaction that will not be properly understood until all parts are understood, with no part independent of the others, the whole requiring the parts, and the parts the whole.

HUMAN INFORMATION PROCESSING: THE CONVENTIONAL VIEW

When I first began the study of psychology, I was interested in mechanisms. The task seemed straightforward enough—difficult, yes, but well defined. The human is an animate being, functioning in the environment. The human has certain biological facets, physical facets, intellectual facets. The basic conceptualization went like this: Intellectual processes are the result of the operation of several separable systems: sensory-perceptual systems, central processing (thought), memory, and response output (motor control). Sensory transducers feed a steady stream of information about the environment to some central processing structures where that information is analyzed, interpreted and fed to a response system which controls body movements and speech sounds.

Considerations of this sort led to the view—the reasonably well accepted view in psychology—of the human as composed of separable subsystems of information processing mechanisms: perceptual systems (including pattern recognition), motor or output systems, memory systems, and systems for internal reasoning and deduction, which includes thought, problem solving, and language. A summary of the components and a rough sketch of their interactions is shown in Figure 1, which might be considered to be a modern updating of the conventional flow chart of the human information processing system. The figure summarizes what is known today about the “Pure Cognitive System,” the system built around pure cognitive functioning, with a physical symbol system as its central component.

Different workers might put more weight on one aspect of this system than on another, but on the whole, this has come to be a fairly well accepted view of things. I will not review for you the history of this and other approaches to the study of the human information processing system, but I will discuss some aspects of it. Basically, I believe that although this view is accurate, it is but one of many possible views. Taken alone, this view is both inadequate and misleading.

In recent years I have become more and more dissatisfied with the conventional view of information processing. The source of the dissatisfaction was not obvious: each of the components of Figure 1 seemed reasonable, and although one might (and did) argue about the details, the powerful arguments for physical symbol systems seemed persuasive. The problem seemed to be in the lack of consideration of other aspects of human behavior, of interaction with other people and with the environment, of the influence of the history of the person, or even the culture, and of the lack of consideration of the special problems and issues confronting an animate organism that must survive as both an individual and as a species, so that intellectual functioning might perhaps be placed in a proper perspective. These considerations have accumulated until they finally have forced themselves upon me. The human is a physical symbol system, yes,

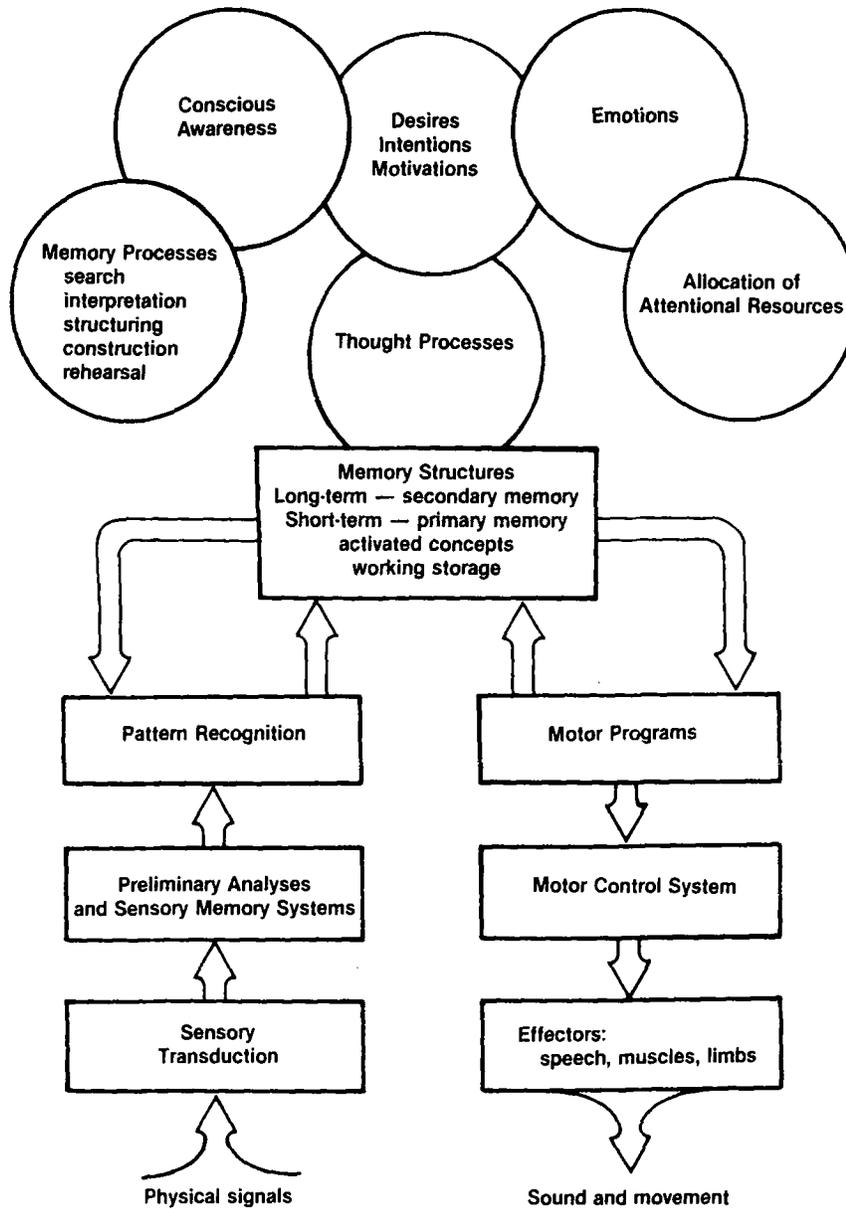


Figure 1. A modern version of the conventional flow chart of the human information processing system. The basic components are a series of processing mechanisms that take in information about the environment, perform general central processing operations, and control motor output. The central processing is complex, with various sources of knowledge interacting with one another, controlled by an as-yet little understood processing structure which allows for some simultaneous operation, self awareness, consciousness of some of the processes. The stuff in the central part of the figure is sufficiently vague as to allow for a large number of interpretations of its nature.

with a component of pure cognition describable by mechanisms of the sort illustrated in Figure 1. But the human is more: The human is an animate organism, with a biological basis and an evolutionary and cultural history. Moreover, the human is a social animal, interacting with others, with the environment, and with itself. The core disciplines of cognitive science have tended to ignore these aspects of behavior. The results have been considerable progress on some fronts, but sterility overall, for the organism we are analyzing is conceived as pure intellect, communicating with one another in logical dialogue, perceiving, remembering, thinking where appropriate, reasoning its way through the well-formed problems that are encountered in the day. Alas, that description does not fit actual behavior.

These objections are not novel. They are raised in one way or another by other contributors to these proceedings. Simon reminded us that behavior is always relative to the environment, so in the study of human behavior we are really studying social behavior. I agree, but feel he did not go far enough: there is more to interaction than social interaction. Geschwind reminded us of our biological origins, with emotional systems playing a central role, not just in overall behavior, but perhaps even in such "pure" cognitive functions as memory. We have wired-in, specialized subsystems for doing what might seem to be general processes, such as recognizing faces. Johnson-Laird and Lakoff reminded us that thought may not proceed smoothly through logical constructions, but may instead rely upon metaphorical modeling of the current situation as an instance of a past experience, so that experience colors thought in fundamental ways. And Winograd argued for a much richer analysis of our history and our social and cultural interactions as a prelude to the understanding of language.

I expand upon these arguments, for although I sympathize with them, I do not think that even they went far enough. Each of the papers of these proceedings presents one point of view, each view appropriate for some aspect of intelligent behavior. I wish to take yet another view, one that attempts to put the others into proper perspective. Let me illustrate by several examples. One is a brief description of an airplane accident, another the view of classroom behavior. These two examples are followed by some general discussion of human functioning and then by a re-evaluation of the role of pure cognition. I conclude that there is more to human intelligence than the pure cognitive system, and that a science of Cognition cannot afford to ignore these other aspects.

Tenerife

In March of 1977, two Boeing 747 airliners collided on a runway at Tenerife, in the Canary Islands. The crash killed 582 people. What caused the accident? No single factor. The crash resulted from a complex interaction of events, including problems of attentional focus, the effects of expectation upon language understanding that combined with an inability to communicate effectively over a

technically limited communication channel when there were major difficulties in language (although all involved were speaking English), the subtle effects of differences of social structure among the participants, the effects of stress, economic responsibilities and social and cultural factors upon decision making. All in all, it is a fascinating—if horrifying—story for Cognitive Science.

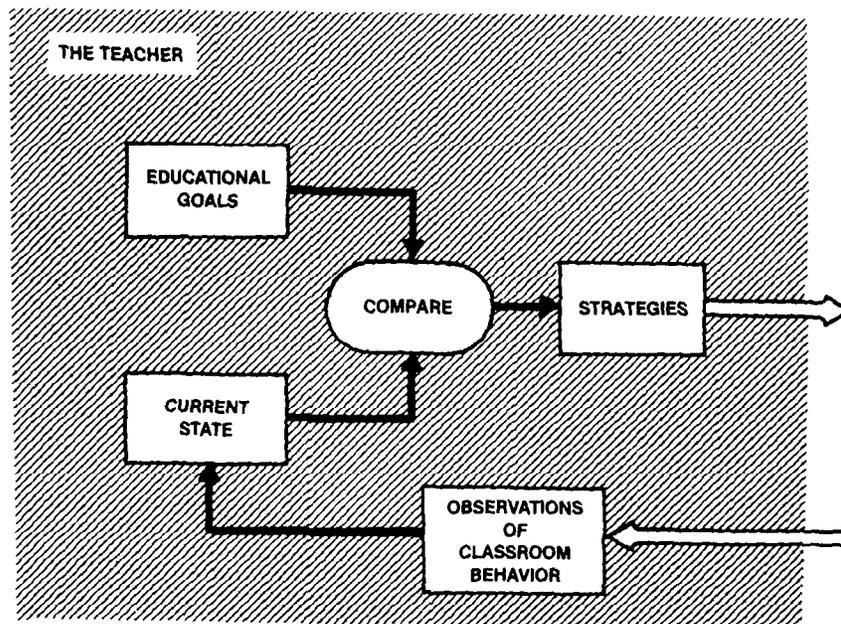


Figure 2. An information processing model of the teacher. Starting with educational goals, the teacher compares those goals with the current state of classroom behavior and knowledge and uses an instructional strategy appropriate to the situation. The teacher continually monitors classroom behavior and modifies the instructional strategy, or the knowledge being taught accordingly. This is a feedback model of instruction. The "current state" implies, among other things, a model of student knowledge and behavior. This model of a teacher is common to modern instructional theory (including my own). It is probably necessary, but by itself, it fails to be useful in the prediction of classroom behavior.

A View of the Classroom

Consider the classroom situation, especially the early grades of school. The teacher has a point to make, a body of information to get across. This aspect of teaching has been receiving considerable attention in recent years. The teacher must construct a mental model of student knowledge, match the model of the student with that of the desired endpoint, determine some strategy for presenting the information not yet currently held by the students, and go forth and teach. Figure 2 shows a possible model of the teacher. Don't worry about the details, just think of the model as an attempt to summarize how the teacher determines the appropriate way to transmit the topic matter to the class.

The individual students must themselves be represented by models similar to that of Figure 2, except complementary in that they respond to the new information about the topic matter and construct mental memory structures to accommodate them. Each student has some knowledge and as the student interacts with the teacher, the student knowledge is altered and enriched in appropriate ways. If questioned by the teacher, the student can apply the new knowledge in order to answer the query, thereby providing feedback to the teacher about the state of learning.

Alas, anyone who has actually taught in a classroom (especially an elementary school classroom) knows that this description provides only the most idealistic view of the real behavior. Some of the description is appropriate, but there is much more happening. In some classrooms, it would be difficult to find any evidence that teaching—in the sense just described—was ever taking place. The students are in a social setting, interacting with one another, acutely aware of each other and of the overall classroom behavior. Individual students tailor their behavior for the other students to some degree, sometimes entirely for the other students. The behavior of the teacher and the individual students is responsive to the events of the classroom, but the classroom events are the results of the combination of behaviors of the teachers and the students.

Cybernetics and Behavior. Cybernetics. A term connotating engineering models of servomechanism systems, the sort of systems one would expect for motor control, or for homeostatic body functioning. Why do I introduce it here?

I use the term “cybernetics” to mean a feedback system, one in which the operation of the system depends upon interaction with the environment. This is what Norbert Weiner meant when he coined the term. The concept has been lost from most of cognitive studies, in part because of the lack of study of output and of performance (more on this later). Without output, there is no feedback. Without global views of functioning, the question of the role of feedback does not arise.

Much social interaction can be viewed as a cybernetic system. Each person is responsive to the environment. Each person is a human information processing system, consisting of something like the components of Figure 1, each behaving something like the model of the teacher presented in Figure 2. But the overall behavior is the result of all the participants, and the participants, in turn, respond to the total behavior. The overall view is something like Figure 3—a view that works for both the classroom and for the Tenerife situation as well.

Suppose we are interested in classroom instruction. In this case we need to understand classroom interaction, the classroom behavior. We must take a view that is something like that shown in Figure 3. We need to understand the several different interactive themes that are simultaneously active within the classroom: the social interactions among the students, the sociolinguistics of their language use, the status differences among the students and between students and teacher.

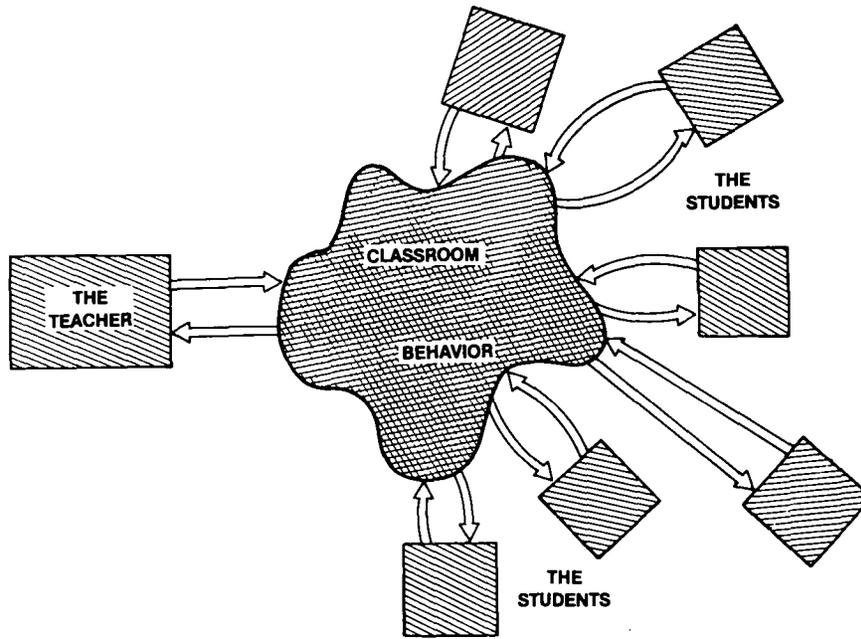


Figure 3. The classroom behavior is the result of a combination of interactions. Each student responds to the behavior of the classroom, as well as to the internal goals of satisfying other students, the teacher, and self needs. So too with the teacher. The teacher and students are all modeled by something akin to Figure 2, but understanding of their behavior requires understanding of the entire interaction. The classroom is a system of individual cybernetic actors. (This basic picture of interaction applies to a variety of situations in addition to the classroom. Only the labels need be changed.)

These all color the use of language and participation. Even the seating pattern and room arrangement will turn out to matter.

Obviously one also needs to know of the motives that drive the teacher, the lesson that is to be taught, the time constraints that must be obeyed, the kind of classroom interaction the teacher desires, and the kind that the teacher will tolerate.

Now, if one wishes to understand the particular responses of the teacher or of an individual child to a particular classroom event, then it is going to be necessary to have an information processing view of the person, somewhat of the form of Figure 2. But the model is only going to be useful if it is coupled with an understanding of the several simultaneous (and possibly conflicting) goals and motivations of the various participants.

My point is not an indictment of any particular approach to the study of learning and teaching. On the contrary, it is a statement that all approaches are necessary. The information processing psychologist who studies the transfer of

knowledge from teacher to student is contributing some understanding of the classroom situation. The person who studies the sociological influences upon the students' behavior and their tolerance to classroom discipline is also contributing some understanding, of an entirely different kind. My argument is that the situation is not going to be understood until all these different points of view are combined, for the overall classroom behavior is a result of all these forces, no one more fundamental than another.

How Much Does Cognitive Science Know?

I am struck by how little is known about so much of cognition. The crash at Tenerife and the interactions of the classroom are but two examples of the complex interactions of cognitive factors that can play important roles in our lives. But there are much simpler examples.

Memory. I have studied memory for years, yet am unable to answer even simple questions about the use of memory in everyday life. Mental activity. The study of thought processes has concentrated upon logical, systematic behavior, one step at a time. What about the processing deeply buried within the subconscious where it can go on without awareness for hours, days (months?).

Slips, Freudian and Otherwise. People make slips of the tongue, slips of action. Some are undoubtedly easy to explain: confusions, lack of knowledge, or obvious sources of the error. But others require much more subtle analyses, involving the nature of memory and retrieval, activation and stress, or conflicting simultaneous thoughts. Freud had a theory, one that I suspect is much more appropriate than Cognitive Scientists tend to give credit today. At least Freud did worry about the relationships among emotions, conscious and subconscious events (we would say "processing"), and how the subconscious is manifested in behavior.

Performance. Consider the highly skilled typist, producing over 100 words per minute, about 10 keystrokes per second. Or the professional pianist playing 25 notes per second in a Chopin Nocturne. Motor skills are fascinating aspects of our behavior, little understood, little studied (in comparison with, say, language). How does one hit a baseball that is travelling at great speed, or steer a speeding automobile through narrow spaces, or control a large crane, making precise movements at the end of a boom a hundred meters long with controls that seem to have little relevance to the actions being performed?

Language. If you think we understand language, well, how about real language, the language between two people in casual conversation? By the rules of formal language, such language is often ungrammatical and it should be

unintelligible. Indeed, as an inveterate eavesdropper on the conversations of strangers (in the name of Science, of course), I can attest that one-minute fragments of other people's conversations are unintelligible and remarkably free of content. The conversants would not agree. They have established sufficient bonds that they can relax the normal constraints of language. Gesture, timing, intonation can carry as much weight as the formal content of the words. This observation is not meant to be a surprise: We are all aware of the phenomenon. But not as scientists: we do not understand how.

You will not be surprised if I tell you that we understand little of the interactions of social groups, or of society, or of cultures, especially of the mechanisms of that interaction. You might be surprised if I claim that these factors play a large role in everyday behavior, even in performance on our abstract tasks within the laboratory. Perhaps one reason that our theories of the separable components of information processing say so little about real world activities is the neglect of social and cultural factors, among other things.

One goal of this paper is to convince you that the study of cognition requires the consideration of all these different aspects of the entire system, including the parts that are both internal and external to the cognizer. (By "internal," I mean the knowledge, the processing mechanisms, the rules, strategies, and control mechanisms. By "external" I mean the environment, the society, culture, and the interaction of all these with one another.) Of course no one can study everything all at the same time, but I argue that we cannot ignore these things either, else the individual pieces that we study in such detail will not fit together in the absence of some thought about the whole.

ON THE DIFFERENCES BETWEEN ANIMATE AND ARTIFICIAL COGNITIVE SYSTEMS

Intelligence, thought, cognition—these are central topics in the study of Cognitive Science. So let us start by considering the elements of a cognitive system. Suppose we concentrate on the intellectual functioning—what are the essential elements of a cognitive system? Let me go through the arguments of the necessary components, starting with a reasonably traditional view (I will end quite differently). The basic picture of the human information processing structure, in its modern format, has been presented as Figure 1.

Now consider, if you will, an intelligent artificial system, one that might be the goal of your favorite robotologist. What does an artificial system need? Obviously there are several possible answers. If we consider only the *Pure Cognitive System* (henceforth, PCS), then we see obvious differences in structure between natural and artificial systems, between electronics and biology. Nerve cells convey their signals through electrical potentials, by chemical transmission. They are affected by biological chemicals (hormones, nutritive fluids, transmitter

substances). And natural systems have wiring diagrams that are not yet understood, that seem to be adaptive, that have billions of interconnections. But despite the obvious differences, there are no obvious differences at the level of functional mechanisms. Presumably, the biological system has memory structures, perceptual structures, and so on, and in principle, if we wish to and knew enough we could build artificial systems whose operations mimicked the biological ones. We would need to learn a considerable amount more than we currently know about the functioning of such a system, but the "in principle" point is what is critical for those of us who pursue the study of psychological mechanisms.

But wait. The difference between natural and artificial devices is not simply that they are constructed of different stuff; their basic functions differ. Humans survive, get nourishment from the environment, protect themselves against physical insult, form families and societies, reproduce themselves and protect and educate the young. Much of this is handled with the aid of biological structures that I will call the *Regulatory System* (RS). Consider how the RS interacts with the cognitive system—something like Figure 4 emerges.

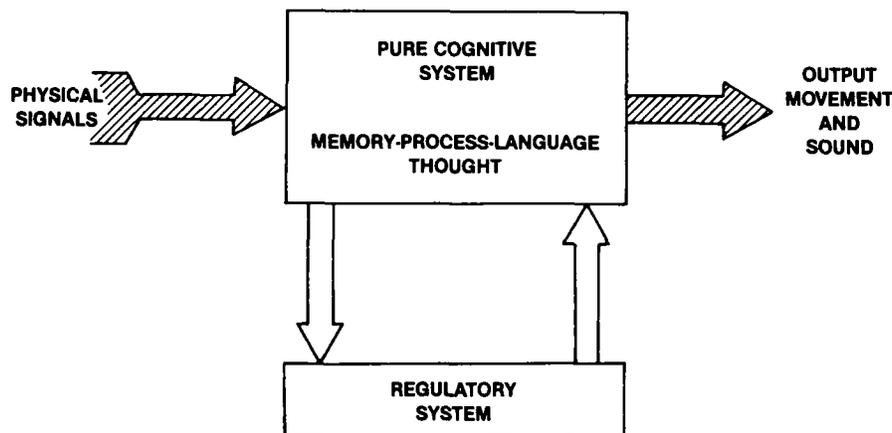


Figure 4. To the Pure Cognitive System of Figure 1 we must add the properties of the Regulatory System. In this view of things, the Cognitive System dominates. This view is an obvious one, but probably wrong.

Consider the implications of Figure 4. (Yes, even such a simple diagram does have implications.) Dangerous situations require immediate attention, immediate response. If potential danger is to be discovered quickly, there must be continual monitoring of possible sources of evidence. Moreover, when danger is detected, then the organism must be alerted to the problem and it must allocate sufficient resources to deal with it. It is easy to understand how this might be done when dealing with things like changes in body states, such as temperature, blood sugar level or fatigue. Environmental situations that lead to pain or other-

wise send sensory signals that can be monitored are also easy to understand with the framework shown in Figure 4.

The issues are not so simple when we consider how to respond to events that must be interpreted: dangerous heights, the sight of a wild tiger, fire, the sound of an explosion, or the airline pilot's announcement that two engines have failed. For these events, perception, knowledge, and language must be called into play—essentially, all the mechanisms of the Pure Cognitive System. But these interpretations must operate with immediacy, interrupting whatever primary task was going on. The problem here is that it takes the cognitive system to do the interpretation for the maintenance system, thereby allowing the maintenance system to interrupt the cognitive system. It can't work.

We need to rethink the organization implied in Figure 4: maybe the PCS is not the pinnacle of human functioning. It is comforting to think so, that the focal point is PCS, with the RS serving to maintain both the body and the PCS. This egotistical point of view is especially nice for intellectuals, but it doesn't hold up. It is always dangerous to invent and then to rely on biological principles and evolutionary causation, but it is also useful. Did the evolutionary sequence that produced superior cognitive systems do so to permit professors to exist, to publish, to hold conferences? One suspects not, that the regulatory system was first, that the cognitive system grew out of the requirements of that system. To determine that a limb should be withdrawn from a painful stimulus did not require much cognition: to avoid the situation in the first place did.

The point is simply that the functions and the requirements of animate systems include the problem of survival, and that this problem requires a regulatory system of considerable complexity, one in which considerable cognitive power is required. And so, the cognitive system is apt to be the servant of the regulatory system, not the other way around, as shown in Figure 5. Emotional systems might very well be an interplay between the two, so that perceptual analysis (done by the PCS) might at times cause the RS to create the necessary emotional arousal to alert the PCS.

If the RS dominates, with the cognitive system its servant, interesting implications follow. Perhaps PCS is a myth, with intellectual thought an outgrowth of the use of biological function for purposes somewhat foreign to the original need. Cognitive systems might be the result of the generally increasing demands of the regulatory system for an intelligent component. Perhaps when the cognitive side reached some critical mass, it then possessed sufficient computational power to have its own existence and to establish its own goals and functions. But only afterwards, grafted on, if you will, to the functions of supporting one's own life.

What about emotions? Are they superfluous to cognitive functioning? Most of us—and I include myself in the "us"—would prefer to believe this. Contemporary theories of cognitive functioning—no matter the discipline—seem to be theories of pure reason. Emotions have to do with something else, perhaps an

evolutionary carry over from an earlier time when the demands upon human functioning were different. Well, the novelist, the playwright, the clinical psychologist and psychiatrist know differently. If I am correct in my assertion that the cognitive system is subservient to the regulatory system, with pure cognition an artificial situation grafted on to a biological organism, then emotions play a critical role in behavior.

A summary of these arguments about the nature of cognitive systems in general and of animate cognitive systems in particular is presented in Table 1. I believe that we should reconsider the functioning of human processing. Some things will not change: our observations and theories will still apply. But I suspect some aspects will change in fundamental ways. I cannot now tell you what will change and what will not: we must wait and see.

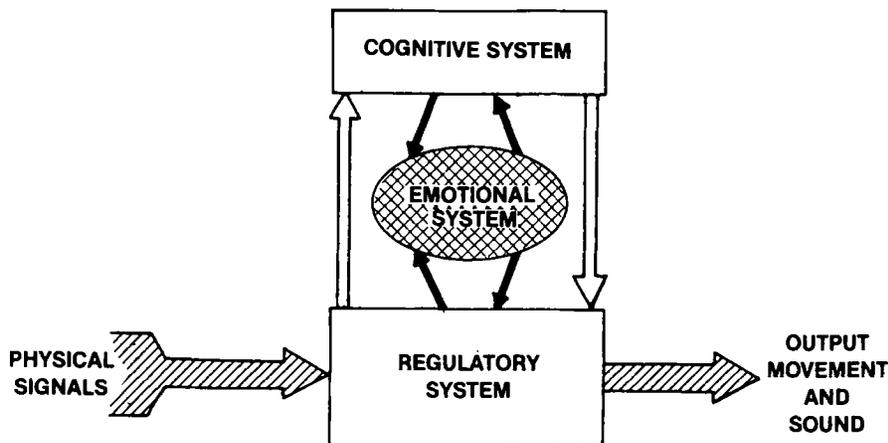


Figure 5. The Regulatory System is here given primacy over the Pure Cognitive System. Compare with Figure 4: the basic format is the same, except that sensory inputs and motor outputs now leave and enter RS rather than PCS. An emotional system stands between. And the relative sizes of the boxes that symbolize the systems have been changed to mark the change in emphasis.

SOME ISSUES FOR COGNITIVE SCIENCE

The arguments of the preceding pages suggest that we must broaden the issues considered by the discipline of Cognitive Science. In fact, twelve major issues attract my attention. These twelve are neither independent of one another nor equal in importance. I do claim, however, that these twelve are among a core group of issues along which we must progress if our field is to make substantive advances.

I believe in the value of multiple philosophies, multiple viewpoints, multiple approaches to common issues. I believe a virtue of Cognitive Science is that it

TABLE 1
Essential Elements of Cognitive Systems:
In General, and in Animate Organisms

All Cognitive Systems

All cognitive systems, animate and artificial, must have the following:

A way of receiving information about the world: receptors

A way of performing actions upon the world: motor control

Cognitive processes, which include:

a means of interpreting and identifying information received by the receptors

a means of controlling the actions to be performed

a means of guiding the allocation of cognitive resources when more needs to be done than can immediately be done (this can be derived from the fact that a finite system must have finite resources)

a memory for the history of actions and experiences

These cognitive processes imply that:

because resources are finite, there will be times when more is being attempted than can be accomplished; some means of resource allocation (attention) will be required

because there will be synchronization problems with events in the environment and internal events, buffer (short-term) memories are required

There must be basic operations, an interpreter, and some feedback mechanisms that can observe the effect of operations upon the world and change accordingly

There must be some way to devise plans and then to monitor their operation; this requires levels of knowledge—meta-knowledge

For intelligent interaction, there must be a model of the environment, of one's self, and of others

There must be learning, changing one's behavior and knowledge in fundamental ways (as opposed to simple adaptation), and this will probably require a system capable of inferring causality, inter-relations among concepts and events, and self-observation

Animate Systems

A major difference between animate and inanimate systems is that an animate system maintains itself, protects itself, regulates its own operation, and reproduces itself. A newly born organism requires considerable physical, biological, and educational maturation, which takes place through a protracted time course of infancy, childhood, and adulthood. The organism at birth differs from adulthood:

It is smaller, both physically and in the amount of its cognitive (neural) structure;

It has less knowledge;

Its regulatory system is not fully developed.

An animate system must survive, which means it must be alert for unexpected occurrences: its regulatory and cognitive systems must interact. The regulatory system is a homeostatic system, designed to maintain life. It must interact with the cognitive system, for interpretations are required of the situation and actions are required to maintain homeostasis, comfort, and safety.

An animate system has goals, desires, purposes. The system is motivated to perform some activities. There must be a means of selecting "interesting" and goal-related tasks from among those that could possibly be done, controlling the amount of effort devoted to that task, and scheduling the initiation and termination of the various activities. Long term goals and issues related to survival receive dominance, although the mechanisms for accomplishing these may not be part of the self-awareness of the organism.

brings together heretofore disparate disciplines to work on common themes. My reason for discussing these twelve issues is the hope that I can focus some efforts upon them. I introduce and discuss these issues from my own perspective, which is primarily that of a psychologist interested in the workings of the mind. I treat the twelve briefly. My intention is to raise them, not discuss them in detail: that is done elsewhere. Alternative points of view are possible, welcome, and necessary.

Issues Are Not Levels

The issues are topic matters that are to be studied. They are not the names of disciplines nor prescriptions for methods of study. Each issue should be addressed from different directions, yielding different levels of explanation.

When we come to describe the mechanisms of cognition, the explanations should be couched at several different levels. The psychologist talks of functional mechanisms and of behavior, the neuroscientist talks of cells and neural systems. The anthropologist and the sociologist each have their levels of analysis. Linguists, philosophers, and computer scientists each view cognition from their special perspectives, each different, yet complementary. I believe that the complete science must have all of these different levels represented. We need to know about the neurological and biological basis for animate cognitive systems, about the mathematical and philosophical basis for cognitive systems, about the mechanistic basis for artificial systems. But that is not the point of this paper. The issues I discuss are not statements about the philosophy or level of approach. Rather they are issues, or problem areas, that should be considered.

The Twelve Issues

I give you the following issues:

Belief systems	Emotion	Learning	Performance
Consciousness	Interaction	Memory	Skill
Development	Language	Perception	Thought

“What a strange list,” you must be thinking. Not what you expected. Emotion? Skill? On the same level as language and memory? Aren’t learning and memory and skill and performance all the same, or at least highly related? What about motivation, or representation, or whatever your favorite topic? Wait, I will clarify some of the problems (though not all). Remember, these 12 issues are ones that I see as key to the development of a science of cognition. Not all are recognized by everyone as being relevant. Not all are thought to be important. Some are well studied, but not normally thought to be a part of cognition. Some issues are seen as simply subsets of others. I disagree: all are essential.

A BRIEF TOUR OF THE TWELVE ISSUES

Belief Systems

I start with Belief Systems, accidentally the first in my alphabetized list of issues, but deserving of primacy under other criteria as well. For belief systems mark the merger of the traditional domain of cognitive science—the study of knowledge—with the domains of those who study real world interaction of humans—the anthropologists, the social psychologists and the sociologists. This issue could perhaps more easily be called “knowledge,” or perhaps, “world knowledge.” I do not use these labels in order to emphasize the merger of several different classes of knowledge, including culture, belief, and world knowledge of several sorts. The basic concept here is that we acquire a lot of knowledge over our lifetime which then colors our interactions with others, with the environment, and even our internal processing. A major component of anthropology and sociology is concerned with the examination of these belief structures.

Cultural knowledge is that special subset of general knowledge that is passed on from generation to generation, taught in the family, or in the schools, or (more commonly) not so much taught as experienced. Styles of dress, social interaction, rankings of social groups, interaction patterns including conversational (discourse) rules, social deference, and other patterns are included here. The physical shape of the environment is altered through culture. The style of buildings, paths, transportation—our technology.

The belief systems go beyond obvious cultural interactions, however. They carry over to such things as rules for memory and thought. You will come to believe these statements more the more you believe that thought and memory are done through reference to real world experience. Suppose that logical inference is normally done by setting up a mental model of a concrete analogy to the problem, using experience to guide the solution of that concrete analogy, then interpreting the result for the problem at hand. If this is the case, then belief systems are of critical importance in determining the basis for much of thought.

Similar statements can be made about memory, perception, problem solving, the interpretation of texts and the conduct of dialogs, legal negotiations, and so on, and on, and on. Many of you are familiar with delusional belief systems that result from mental abnormalities (paranoia being the most fashionable to discuss, for it seems most directly tied to the development of a rich, delusional system, self consistent in its own way, but a great danger to the possessor).

So, belief systems are important, both as interesting items of study in their own right but also as important determiners of much of the rest of cognitive behavior. At the moment, the tools for formal analysis of such structures are just being developed. There is much talk among cognitive anthropologists and sociologists of scripts and schemas, of story grammars, and representational issues. I have an obvious interest in this direction of work, having myself urged the importance of the study of representation and the utility of the study of

structured memory units (schemas). The major issue, though, is not yet the one of representation. Rather, we must first lay out the development of the problem itself, examine the nature of belief systems in general, and determine what the implications are for cognitive behavior. My brief excursions into this area have left me impressed with how much my own hidden belief structures influence my "pure" logical inference, memory processes, and social interactions. I suspect that we will find that more of our behavior is thus determined, not less.

Consciousness

"Everyone knows what attention is." So said William James in 1890, and so too have I said repeatedly in my courses and lectures on attention. But the statement is false, quite false. We really do not know about attention, to a large extent because we do not know about consciousness. Studies of attention have restricted themselves to a small segment of the phenomena of which James wrote:

It is the taking possession of the mind, in clear and vivid form, of one out of what seems several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others (James, 1890, Vol. I., p. 403–404).

Consciousness, under which I include the issues of conscious and subconscious thought, the problem of self awareness, attention, the control structures of cognition, the formation of intentions. Here too are such issues as the phenomenological states of awareness, states of consciousness. Hypnosis: a powerful force, potentially a powerful tool for the investigation of consciousness, but little understood, not sufficiently well explored.

It should not be necessary to talk about consciousness to a group concerned with cognition. But consciousness is a peculiar stepchild of our discipline, agreed to be important, but little explored in research or theory. There are legitimate reasons for the relative neglect. This is a most difficult topic, one for which it is very difficult to get the hard, sensible evidence that experimental disciplines require. We have little idea of the real nature of consciousness, of the functions it might serve, of the nature of the subconscious. We are just beginning to get a glimmer of the phenomenology of consciousness, of different states of awareness and different phenomenological experiences (though most of this comes from nontraditional sources).

It is exactly the description given in the quote from James that we do not understand, cannot understand until we come to a better appreciation of the working of the mind, of the several simultaneous trains of thought that can occur, of the differences between conscious and subconscious processing, and of what it means to focus upon one train of thought to the exclusion of others. What—who—does the focussing, what happens to those other trains of thought as they are excluded? (Some, I am certain continue silently, unheeded, as subconscious

processes that may later interrupt to again force conscious attention to themselves.) And what does it mean to have *conscious attention*? Can there be attention that is not conscious? What—who—experiences the result of conscious attentional processes?

Some of these issues seem to result directly from the properties of an animate cognitive organism. An animate organism can not afford the luxury of concentrating entirely upon a problem until it has been completed. Animate organisms must be multiple-minded, data-driven by environmental events, ever ready to capitalize on the accidents of the world, or to avoid the unexpected dangerous spots. The tasks we assign ourselves to do are often long and complex ones, things which we are incapable of completing at one sitting. We have finite cognitive resources and these must be deployed in some manner that is effective. We can't be entirely data-driven, else the steady flow of information from the sensory system would completely occupy our attention: we must be able to exclude the excludable, to concentrate upon that which is most important (or interesting) at the moment.

Subconscious processing seems essential to functioning. Whatever the special properties of consciousness, they are not needed by all mental processes. (Elsewhere I have argued that consciousness is important for the formation of intentions, the monitoring of their performance, and that it is needed only where the actions to be performed are not routine and well established.) More on this in the section on skills.

Glimpses into the role of conscious and subconscious processing can come from several sources. Hypnotic experiences offer one method, and they can be performed with some rigor in the experimental laboratory. Experiences of subconscious problem-solving or memory retrieval are often experienced and talked about, and there is some possibility that they too can be explored experimentally. Studies of attention are, of course, another possible route, one that has been under active exploration. And there are the errors that people make, slips of the tongue, slips of action, another source of information about subconscious processes and their relationship to conscious ones, to thoughts and motives and intentions. Experiential literature is relevant too, although it must be approached with caution, separating out the description of the experience from the interpretation of that experience, something the experiencer may not be able to do as well as an external observer.

Development

A child is not a small sized-adult, simply lacking in experience, in physical development, and in knowledge, waiting for its head to be filled with the mindstuff of an adult. As adults, we have a wide range of skills, enormous amounts of detailed, specialized knowledge, well established belief systems. We are not just more than children, we are different.

The study of development is well established, of course, hardly in need of suggestions or advice. (Although the studies concentrate upon the years just after birth, with little exploration of adolescence, adulthood, or aging). But in the study of adult cognition there seems to be the implicit assumption that once we come to understand adults, children will simply be seen to be at various stages along the pathway towards the adult. Perhaps. But perhaps also that the complexity and experience of the adult will forever mask some properties. Automatic behavior masks the underlying structure, pushing things beneath the conscious surface to the inaccessibility of subconscious processes. Well established belief and knowledge systems mask their content.

Much of cognitive behavior could be studied best through the developmental cycle, with the history of the development leading to better understanding of the adult. Animate organisms take very long times to reach adulthood: the human is learning new concepts throughout the entire life span. Language learning goes on through the late teenage years, and vocabulary learning never ceases. We are fundamentally organisms that learn, that develop over time. By ignoring this aspect of behavior and concentrating on the static phases we may miss the keys to understanding.

Emotion

And what is the role of emotion in the study of cognition? We leave it to the poet, the playwright, the novelist. As people, we delight in art and in music. We fight, get angered, have joy, grief, happiness. But as students of mental events, we are ignorant of why, how.

Emotion. Is it a leftover of a primitive alerting system, or is it a sophisticated set of states reaching its highest pinnacle in the human? Earlier I argued in the direction of the latter point of view. Now, I simply remind you of the issue. The study of emotion is an important field, with important findings and implications for the study of cognition. We cannot ignore our biological heritage, ignore our emotional states. Geschwind, in his paper, reminds us of the fundamental role that emotion plays in biological organisms, and of the close relationship between the neurological structures thought to be important for emotion and those thought to be important for memory. Indeed, there is some experimental evidence for state-induced memory retrieval, so that we remember best events whose emotional content matches our current state: sad events are best remembered while sad, happy ones while happy. Geschwind suggested that some neurological control structures have dual activations, one from below—from the emotions—one from above—the intellect. We smile, cry, and laugh from emotional signals: our attempt to mimic these acts from intellectual desires or upon receipt of a verbal command to do so recreates neither the true emotion, nor the same motor actions. An observer can often tell which behavior is real, which synthetic.

Interaction

Human beings are social organisms. Our intelligence does not operate in isolation, but rather in conjunction. We interact with others, we transmit knowledge through cultures.

We supplement our intelligence with social interaction, by our use of the environment, through the construction of artifacts: reading and writing (and paper, pen, printing press); machine transportation; communication methods that operate over distance (signalling devices, mail services, telegraph, phone, vision); machines for commerce, for other essentials of life; and machines for computation. The interactions that result become a fundamental aspect of our behavior. In some sense our intelligence has become partially externalized, contained in the artifacts as much as in our head. ("I don't need to know that," we say, "I just need to know . . ."—choose one: "who to ask," "where to look," "where to go to find out," "that it is known.")

My major concern here is social interactions, but the issues of interaction share properties, whether it be with person, society, machine. We need to have mental models of the people (and things) with which we interact, for communication depends strongly upon mutual use of shared knowledge, shared understandings. Without a good model of the digital-chronograph-stopwatch-calendar-timer-watch, remembering which buttons to push for what is a hopeless task. With a good mental model ("good" does not mean the "true" model, just a consistent one), the buttons make sense and the use is facilitated. Without a good mental model of our conversational partners, the conversation does not make progress. "Where is the empire state building?" The answer depends upon why the question was asked, in what part of the world it was asked, and how much the questioner needs to know.

Much of the study of cognitive processes has been the study of the isolated person. Much of the study of interactive groups has been of the dynamics of the situation, or of the behavioral aspects of the group. To my knowledge, little has been done to combine these efforts, to examine the individual cognitive processes as they are used within interactive settings. But, because the normal mode for the human is to interact, the studies of memory and language and problem solving and decision making in isolation address only one part of the mechanisms of human cognition.

The earlier section on "Cybernetics and Behavior" was intended to introduce the importance of the consideration of interaction, and so I pass on to the next issue.

Language and Perception

I include these two issues to remind us that they exist, to dispell any illusion that I have forgotten them. But I do not wish to discuss either language or perception, primarily because they are of such central importance that they have already

received sufficient emphasis. Actually, the emphasis is itself a problem. There is a tendency to identify the study of Cognitive Science with the study of these two topics (and within perception, with visual perception). I believe this to be mistaken, a view that is both wrong and unfortunate. Even language and perception themselves are complex topics, with many different aspects of cognition interwoven together. Like all of the issues within Cognitive Science, these different aspects support one another, enriching the performance of one domain through the knowledge and characteristics of the other domains. I do not believe we can solve the problems of interpretation of language and of perception until we have made substantive progress on the other 10 issues of cognitive science.

Learning

Learning. Recognized by many as a key issue. Still eluding us. In the early days of psychology and in the construction of artificial devices for intelligent behavior, learning was the core topic of study. Machines were constructed that were to learn through their interactions, perhaps to acquire broad, general intelligence as a result. Psychologists developed global theories of human and animal behavior, often built around such fundamental learning principles as “the law of effect” or “associative properties of learning and memory.” It all has come to nought. Today, the study of learning is not considered a central part of either psychology or artificial intelligence. Why? Perhaps because the understanding of learning requires knowing about problems of representation, of input (perception), of output (performance), and of thought and inference. It is only recently that we began to understand these issues with appropriate depth.

We spend much of our lifetimes learning: in a sense, we learn from everything we do. If learning is not yet understood, it is because there is more to it than the simple accumulation of knowledge. Accumulation is indeed one form of learning, but there are other things that must be done. One fundamental mode of learning is that of restructuring one’s knowledge, reformulating the very basis of understanding of some topic as a result of new concepts and new experiences. Then there is the tuning of behavior, the fine sharpening of adequate skills and understanding to that of the expert, smooth, efficient, effortless.

There has been remarkably little study of learning—real learning, the learning of complex topics, the learning that takes months, even years to accomplish. Elsewhere I have estimated that experts at a task may spend 5,000 hours acquiring their skills: that is not such a long time; it is 2½ years of full-time study, 40 hours a week, 50 weeks a year. Not much time to become a professional tennis player, or computer programmer, or linguist. What goes on during that time? Whatever it is, it is slow, continuous. No magic dose of knowledge in the form of pill or lecture. Just a lot of slow, continual exposure to the topic, probably accompanied by several bouts of restructuring of the underlying mental representations, reconceptualizations of the concepts, plus many hours of accumulation of large quantities of facts.

The relative importance of learning is well understood and often stated. We know how important learning is for the child, and how important the developmental sequence from child to adult. Most of us are professional educators. Surprise, then, that so little is known about learning (and so little about the complement, teaching). And in this case, the lack is, in part, from lack of trying. People talk fondly of computer programs that will start with some fundamentals and acquire all the knowledge needed by some natural sequence of learning, experiencing the environment in which it must function. Very little effort gets spent at studying what it would take to accomplish this, perhaps because there is implicit realization that the task is harder than it might seem. Perhaps the sober realization that a newborn infant takes 25 years to become a fledgling professional, perhaps 5,000 hours of practice and training after the fundamentals have been acquired. Who wants a computer program that can't perform well for the first 25 years of fulltime running (or even for the first 5,000 hours—try explaining that to the government funding agency or the University faculty committee, especially when the first few attempts don't even learn after those periods). And so the study and understanding of the learning process remains at a miniscule level. Pity.

Memory

Do not be impressed by all that is presumably known about the psychology of memory. Less is known than you might think.

Research on the properties of memory has several important functions, some obvious, some not so obvious. For one, it must be obvious that human memory is central to human cognition, and that in general, memory systems are central to cognitive systems (that PCS again). But the complexities of retrieval from a very large memory store are not well appreciated. In Computer Science, the real difficulties of memory retrieval have not yet been faced.

How does one find the information required to answer a question when the form of the question was not anticipated at the time of acquiring the information? Not possible with artificial systems today, a commonplace occurrence with people. And how is the desired information recognized once it is found if it wasn't known in the first place? If I seek the name of a long-lost colleague and retrieve the name Isaac Newton, how do I reject that as the name I seek when I do not know the sought-for name? This example provides its own clue to the solution, but the general case is not so simple. How do we remember stories, events, experiences? More to the point, how do we retrieve them when least we expect them?

Memory has some other puzzles. We recognize the meanings of words in tenths of seconds (as in reading), yet may take hours or days to retrieve one of those words when we seek it for use in a sentence. And what is it that keeps the memory search going for those hours or days, while conscious thought proceeds

in other directions, when the need for the word may have long passed? Current events bring to mind previous experiences, not always in any obvious fashion. It is a well accepted statement that memory is associative, that memory structures are organized into some form: networks, concepts, prototypes, basic levels, schemas, frames, units, scripts. How? We need to understand the representation of knowledge, including the process that operate upon the representation. What is motor memory like, or image, or spatial information?

Associations among memory concepts have the immediate suggestion that somehow there is the equivalent of wires interconnecting memory structures. A little thought indicates that the notion of wires (neurons) simply will not do. That implies much too much knowledge of the wire (or its biological equivalent) that is to snake its way among the already existing stuff to the spot some distance away that might correspond to the new stuff (hold with me for the moment the belief that memories are stored in places). Alternatives to wires are not easy to find, the major candidate being numbered, labelled places (don't worry about numbers: just realize that each place must have a unique name). Then, the association between two memory structures is done by giving each one the unique name of the other, trusting to the existence of some clever machinery that can get from one place to another if only it has this name. This problem—I call it the “address problem”—is fundamental to the organization of any large scale associative memory. Bobrow and I have suggested that memory access is by means of “descriptions” of the items sought, our attempt to overcome the address problem. In these meetings, Minsky proposes an alternative view, one that says there are indeed wires (nerves) between associated memory structures, and that the problems of physical interconnection thereby create severe practical constraints on the sorts of things that can get related to one another.

But wait a minute. Why is it that I assume that memories are stored in places. Can't they be distributed in space? (Remember the hologram.) They can. Essentially there are two different classes of memory structures: place memories (the sort I have just described) and additive memories, memory structures which superimpose particular memories on top of one another, relying on various schemes to extract the relevant information. Additive memories include holograms, so-called “associative memories,” and perceptrons (and its modern descendants). These memories offer, for free, content-addressable storage and retrieval, but pose their own host of problems. There has not yet been sufficient research on additive memory structures.

And finally, but of great importance, there are the functional properties of the memory system that have received some attention. Short-term (primary memory) working memory, activations in memory. Then there are various uses of memory: strategies for organizing, strategies for retrieval, rehearsal, the repeating over and over again of an item in temporary memory in order to maintain it while—while what?—while other operations can get done on it, I suppose. Is there one temporary memory? Many? Any? How is stuff represented in perma-

nent memory, in working memory? Images? Propositions? I stop. I could go on indefinitely, but these issues are well known.

Performance

Performance, too long neglected, now just starting to receive its due attention. The problem of output, of performance, of motor control. The human hand is a marvelously complex instrument, with 27 bones, controlled by over 40 muscles (most of the muscles being in the forearm, connected to the fingers through an intricate set of tendons). The high-speed typist or musician moves the fingers with intervals of less than 100 msec., fingers simultaneously moving in different directions for different targets, with different time schedules for their time of tapping the target key (or string). Interesting errors arise in these high-speed operations, errors indicative of control structures: the doubling error in typing in which the wrong letter of a word is doubled, as when "look" becomes "lokk"; the alternation error, similar in spirit to the double in which "these" might become "thses"; the transposition error, in which two neighboring letters exchange positions so that "music" gets typed as "muisc," almost always occurring across hands as if the difficulty resulted from a synchronization problem between the hands, hardly ever within hands. And once mis-synchronized, the hands can continue, smoothly, wrongly, as in my transformation of "artificial" into "aritifical" in the typing of the draft of this paper, each "i" coming one position early. The control process for going from perception of rough draft to the rapid movement of the fingers that produces the final copy is immense, involving synchronization of looking, perception, reading, motor programming, and feedback.

Consider handwriting, simple on the surface, complex in the details. A set of orthogonal muscle control systems, with intricate timing relationships (50 msec timing pulses, so some say). Handwriting can be thought of at many levels: organization of the ideas, determination of the words, physical organization of the words on the page, control of the letters, with individual motions of various sorts—micro motion to make the individual letters, macro motion to shift the palm across the page (the movement occurring only at orthographically determined locations), the global motion to place the hand on the page or move it when returning to a new line or adjusting the placement of words on the page. Each level controlled, perhaps, by different parts of the cognitive system, for the control of the precise timing signals that create the letter segments would seem to be a different problem than determining during what part of the word the palm may shift, which is in turn different from the backup required to dot the i's, cross the t's, or the large shift required when, say, deciding to set things off indented with a large gap from the preceding line.

The motor control programs are non-trivial in character, their set up being as much a cognitive function as is reading, or perceiving, or talking. They take

time—longer with longer or more complex motor sequences. They can be interfered with by simultaneous acts. They require long periods of training.

With all the muscles to control, with so many degrees of freedom possible because of the numerous joints and the flexibility of the body, the computation of the proper motion of each antagonist muscle pair seems beyond possibility. It probably is. Bernstein (1967), the Soviet investigator of motion, argued for complexes of motor control, systems in which one controls ratios and higher level parameters, the local computation available at the spinal cord and lower taking care of the local translations into muscle commands.

Huge hunks of the brain are devoted to motor control. The cerebellum, a marvelous device, seems dedicated to the function, as is the motor cortex. With so much of the brain dedicated to motor control, it seems unthinkable that this issue should be divorced from the study of higher mental processes. The sensory systems and the motor control systems are intimately linked, closely related neurologically. Probably closely linked psychologically.

The problems of performance are real, they require understanding of computational issues of considerable sophistication, and they interact with perceptual and thought processes in fundamental ways. It is possible to argue that much of our knowledge of the world resides in our knowledge of the procedures that interact with the world, that the perceptual-cognitive-motor schemas are unitary memory constructs, and the separation of one from the others destroys the whole.

Skill

Skill? Why is not skill the same as learning, or performance, or memory? Isn't skill simply expert performance?

Skill. A combination of learning and performance. But more than that, perhaps a fundamental aspect of human cognition. Suppose that our biological heritage developed by means of specialized subsystems for specialized behavior. Maybe skills are independent pockets of knowledge, with independent knowledge sources, computational resources, even independent brain and body structures. Maybe, maybe not. As usual, I suspect the truth is somewhere in-between: we are neither general purpose computational devices, all knowledge and abilities being treated alike, nor are we highly specialized subsystems, each independent of the rest. In fact, let me call separate skills "separable," as opposed to "independent." We cannot ignore the specialization of function of an evolving biological creature, and so the issue of whether we have separable skills is an important one, with major implications for theories of human cognition.

Skills, specialized subsystems of knowledge and of performance. The expert at a task performs differently than the non-expert: the statement is correct, but misses the essence of the difference. The expert performer is qualitatively and quantitatively different than the non-expert. Bartlett, in his book on thinking,

stated that a major difference between expert and non-expert performance was timing. Experts had lots of time. They did their tasks easily, smoothly, without apparent effort, and with plenty of excess time. The expert tennis player is there before the ball. The expert pilot flies "ahead of the plane." The difficult looks simple. The non-expert is always scurrying, barely able to cope, rushing from this to that. With the non-expert, the difficult task looks difficult.

There are other differences, differences in perspective. Consider what happens when you first learn to drive an automobile. The instructions you receive emphasize the actions and the mechanics: hold the steering wheel this way, synchronize foot (for clutch) and hand (for gearshift) that way. As you progress, the point of view changes. Now you are turning the wheel, not moving your hands clockwise. Then you are turning the car, later you are entering that driveway. Eventually, when a truly expert driver, you drive to the bank, go shopping. The differences in the qualitative feeling of the performance are great. At the expert level, you may no longer be aware of all the subsidiary operations that you perform: you look at the driveway, form the intention to enter, and the car obediently follows suit. Driving the car becomes as natural as walking, the car becoming as much a part of the body's controlled appendages as the limbs.

Thought

It is hardly necessary to state that the study of cognition should include the study of thought. The concern, though is not that thought should be included, it is with how the inclusion should go. You may have thought we know a lot about thought. I claim not: what we do know is important, but primarily restricted to that part of the thought process available to conscious awareness—and as long as we lack knowledge of consciousness, we will lack a complete understanding of the role of conscious thought.

A question to be debated seriously is how much thought can be studied in isolation, as if it were a pure, abstract activity, divorced from special knowledge or special mechanisms. The mathematics of thought does indeed have this character, and as that mathematics has been used for models of human thought, it has tended to yield the vision of the human as a general purpose computational device.

But what if we are not so general, if our thought processes are designed for world interaction, with mental models of experiences being the major reasoning method, with limited ability to hold formal constructions in mind while we perform abstract operations upon them. I believe that too much emphasis has been given to possible formal properties of human reasoning, not enough to informal, experiential based models of reasoning. Take care, though, with this argument. I agree with Newell that we must have some class of a general physical symbol system as a basis for much of cognition. We may be specialized, but we can also be general, learning new abilities, reasoning through novel

situations, planning. Still, the strategy used may be to model experiences, to use the properties of spatial arrays to aid our computations, perhaps by using wired-in, specialized spatial knowledge. So I agree with Johnson-Laird, too.

The environment plays an important role in thought. We solve some problems by imagining the environment, solve others by using the environment. Micronesian navigators evidently use the outrigger of their canoe as a sort of analog computer which, when coupled with star positions and rate of passage of water past the canoe, can be used to aid precise navigation for hundreds of miles, out of sight of land much of the way (Hutchins, 1979). We use external aids ourselves, such as pencils, papers, drawings, even the placement of objects on a table. The computer is, in some sense, an artificial extension of our intellect, invented by humans to extend human thought processes. Just as we no longer need to master the art of memory because of the ease of writing, and just as we may no longer need to master arithmetic because of the availability of the calculator, or calligraphy because of the typewriter, we may perhaps forego some forms of thought once small portable computers become commonplace. (Hopefully, thus freeing ourselves for higher levels of thought processes.) Here is not the place for social commentary on these changes, just notice of the heavy dependence our culture places on technological aids to thought processes.

AFTERTHOUGHTS

Is There a Thirteenth Issue? "You left out an issue," my readers rush to tell me, "why do you not have X?" The answer to some extent must be arbitrary. The division of Cognitive Science into 12 issues is idiosyncratic. My list is meant to cover the important principles and phenomena, to be those things that must be included in the study of cognition. The important point, therefore, is not whether my divisions are correct, but whether I have complete coverage. Have I left out anything? Among the various suggestions I have received, one stands out: motivation.

Motivation, the Thirteenth Issue? What makes something interesting? Why do I sometimes watch a television show when I pass by an active set, even when I do not wish to? ("Turn off that TV set," I tell my son, as I walk into the living room, sit down beside him, and watch for 30 minutes, muttering to myself all the while.) For years I studied learning, concerned about the proper way to present material to improve a student's understanding. I studied many things, including proper organizational structure of the material, various instructional strategies, the making of detailed models of teacher, of student, and of topic matter. Yet none of these variables seemed to be as powerful as the one I did not study: changing the motivation of the student. Why is it that we do some tasks easily, readily, while others, seemingly no different, repel us, requiring huge

amounts of self discipline to start, and then to finish? Interest and motivation seem intimately linked, the issues seemingly more complex than can be provided by simple analysis of missing knowledge structures or recourse to concepts such as the overall goals of a person. Note too that the desire to do something is not the same as being motivated to do that thing: I may want to do something, but find it difficult to force myself to do it. I may wish not to do something else, yet find it difficult to stop myself. I am reminded of the distinction Geschwind has made between laughing or crying to a verbal command or to internal signals, the one is difficult or impossible, the other natural and easy.

The arguments for motivation were pointed out to me by Craig Will, one of my graduate students who read an early draft of the paper. Motivation can make the difference between learning or not, decent performance or not, what one attends to, what acts one does. Once, it was a leading topic in psychology, although oftentimes linked to emotion: "Motivation and Emotion," one chapter of a textbook on Human Information Processing is called. Will was persuasive. Is motivation a thirteenth issue?

I think not. I believe motivation to result from a combination of things, from one's fundamental knowledge and goal structures, partially from emotional variables, and partially from decisions about the application of mental resources. Hence, the phenomena of motivation come from various aspects of several issues: Belief Systems, Emotion, and Consciousness. Moreover, and more important, I am not convinced that there is a single phenomenon of motivation (if there is, it should indeed be afforded the special status of an issue). Rather, I believe it to be a complex of things, some biological, some cultural, some emotional, some the result of conscious goals and intentions, others subconscious. Motivation is indeed important, worthy of serious study, and a major determiner of our behavior. I believe, however, that it is a derived issue, composed of different aspects of the others.

The Environmental System and Cultural Knowledge

One early reader of this paper, Michael Cole, suggested that I did not give proper consideration to the role of environment, especially in its role in development. It was his view that my treatment was not satisfactory, that I "... need badly to consider the ES, the Environmental System, consisting of physical and social parts, that is an equal partner in giving shape to the super-system comprised of RS and PCS" (The Regulatory System and the Pure Cognitive System). In part, Cole argued that I could make a much stronger case for:

... the relevance of evolutionary neuroscience, developmental psychology, and cognitive anthropology. Why? Because the mature system that encompasses RS and PCS must develop through a series of interactions between RS and ES (the Regulatory and the Environmental Systems). The PCS should be seen as an evolving adaptation. Where does culture enter? At some point it becomes a part of ES, one with an external

source of memory over generations to supplant the "memory" built in by evolution. The Cognitive System does some work on the Environmental System too, so the current environment is always the product of Regulatory—Environmental—Cognitive System interaction.

The discussion with Cole continued for several days. The behavior of people is shaped by their environment, we both agreed, and a good deal of one's cultural knowledge is the shared strategies that develop for the use of environment. We change the environment through our technological developments and our science, literature, and mathematics. The culture provides us with cultural transformers, and amplifiers, tools that expand our mental abilities. Cole argued that:

What culturally organized knowledge does is to carry a lot of information for us. An extreme way to talk about it is that the information is in the environment, not in the head, so a lot of the processing that experiments require to be done in the head can be, and is, short-circuited in real life. . . . One issue is how to describe cognition as an interaction between head and world where some of the thought power resides in each locus.

The comments are complete enough to be self-contained. Where they are not, where, for example, the manner in which we use environmental information is unclear, the field is unclear. A major task for Cognitive Science.

IMPLICATIONS FOR COGNITIVE SCIENCE

The fact that I can write such a paper, ask such questions, complain with some reasonable specificity, is a positive sign about the emergence of a new discipline. It is a sign of progress that things are sufficiently well understood that the list of non-understood topics can be prepared.

What are the implications for research? In part, my suggestions are going to be received with displeasure. Am I suggesting that everyone become everyone, each person an expert in all other disciplines, all issues? No. I do not believe that productive research, the sort that leads to solid advances in understanding, comes about when efforts are spread apart too thinly over too wide a range. I believe in depth-first research, in concentration upon the minute details of the problem.

But detailed, narrowly-based research should not take place within a vacuum. The choice of the area in which to make the detailed, deep probes must be selected with thought and care to the eventual product of the research. My argument is for goal directed, conceptually based research planning, leading to careful selection of topics, then plunging as deeply as possible into the tangled web of specific problems that exist within any area of concentration. Then, let the results drive the investigation, so that the studies become the driving force for

further research: ye olde standard combination of top-down and bottom-up processing, both conceptually driven and data-driven.

But the research efforts should cease now and then for pause and reflection. Where is the work leading to? Are the tangled problems worthy of further effort, or do they simply lead further from the goal? Without such broad reconsiderations there is a tendency in all domains of research to be captured by the problem of the year, by the race among competing research groups to untangle the theory problems that seem to be holding up current progress. The danger is that in the fun of the race, in the excitement of overcoming technical difficulties, one may forge off in directions of little concern to anyone. For me to urge the need to step back and resurvey the direction of research is not particularly novel (although perhaps it cannot be stated too often). But I am also urging a particular way in which to stand off, a particular set of issues to reconsider each time the overview is made.

I believe in the decomposition of difficult problems into smaller, nearly independent issues—what Simon has called the nearly decomposable problems. We make progress by picking the right size problem, the one we can handle with today's knowledge. I would be doing a disservice were I to convince too many of you to become generalists, biting off more than can be chewed, or swallowed, or digested, even if the bites were from the correct things. I urge the philosophy of nearly decomposable systems, but global considerations of the sort I discuss here are necessary in order to determine the right decomposition.

The major results of my concerns should probably be in the education of new researchers, education at the advanced levels. It is here that I think my points best made, for it is within the education of ourselves and our students that the wider implications and wider aspects of our field ought to be acknowledged, discussed, considered. I would certainly not want my 12 issues to become 12 examination questions or 12 reading lists. I wish Cognitive Science to be recognized as a complex interaction among different issues of concern, an interaction that will not be properly understood until all the parts are understood, with no part independent of the others, the whole requiring the parts, and the parts the whole.

REFERENCES

I avoided references in the text, for the purpose of the paper was to convey the spirit of the argument, not the technical details of each issue. Here I attempt to make up for that lack, citing key references for points discussed within the paper. On most of the issues of this paper I will not give references. Many of the areas are well known, well studied, and for me to select one or two references would do little good. I will present references to material that has heavily influenced me or to things I think are important but not well enough known. This set of references is not meant to be an exhaustive listing of the critical works in

Cognitive Science. It is a personal listing of works that I have found useful along with citations for key ideas discussed in this paper.

A view of the classroom relevant to my discussion—and in part the source of my ideas—comes from the work of Mehan and his colleagues: Mehan (1978, 1979). The report on the Tenerife crash comes from Roitsch, Babcock, and Edmunds (1979).

I restrict myself to two references on consciousness. I take this opportunity to refer to my own "Slips of the mind" (Norman, 1979), a paper that addresses some of the issues of conscious and subconscious control, and gives a reading list on these topics (including Freud, and the rather extensive literature on Slips of the Tongue and Spoonerisms). Second, I recommend Hilgard's (1977) treatment of hypnosis.

I am not expert enough in belief systems to give definitive references. Moreover, some of the work I do know of has only been reported to me verbally. But I recommend Berlin and Kay (1969) on color terms, Berlin (1978) on ethnobiological classification, Cole and Scribner (1974) on culture and thought, Cole, Sharp, and Lave (1976) on the cognitive consequences of education, D'Andrade (1976) on beliefs about illness, and Rosch (1978) on the principles of categorization. Abelson has long been studying belief systems from the point of view of social psychology, including the construction of simulation models of people's belief structures (see Abelson, 1973).

Emotion is a field with an extensive literature. The book that has had the major influence on me is *Mind and emotion*, by George Mandler (1975). This book gives an explicit treatment of the relationship between emotion and cognition from within the framework of human information processing. My discussion of the subservience of the cognitive system to the emotional and regulatory system had its origin here (and from my discussions with Mandler).

Memory has a large literature, and I will therefore avoid citing any of it. But several non-traditional sources seem important to mention. One new book that contains extensive discussions of neglected areas is Kihlstrom and Evan's (1979) *Functional disorders of memory*. Here are discussions of anomalies, amnesias, aging, dreaming, state-dependent memory, motivated forgetting, and repression. There are other treatments of most of these topics, but nowhere are they all so conveniently gathered together: an important collection of topics that ought also to be important, but that have been largely ignored in the cognitive psychology of memory.

Kohonen (1977) provides one treatment of associative (additive) memories, as does Anderson, Silverstein, Ritz, and Jones (1977). The book edited by Hinton and Anderson (in preparation) will treat many aspects of additive memories.

Performance and skill have a few simple references I can point you at. I have already mentioned the Soviet investigator Bernstein (1967), and his works provide an important foundation. Stelmach's (1978) collection provides a valu-

able place to start for the psychological literature. Gallistel (1980) gives an interesting and important review of the neurological work in this area. Other sources exist.

- Abelson, R. P. The structure of belief systems. In R. C. Schank & K. Colby (Eds.), *Computer models of thought and language*. San Francisco: Freeman, 1973.
- 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*, 1977, 84, 413-451.
- Berlin, B. Ethnobiological classification. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, New Jersey: Larry Erlbaum Associates, 1978.
- Berlin, B., & Kay, P. *Basic color terms: Their universality and evolution*. Berkeley: University of California Press, 1969.
- Bernstein, N. *The co-ordination and regulation of movements*. New York: Pergamon Press, 1967.
- Cole, M., & Scribner, S. *Culture and thought*. New York: Wiley, 1974.
- Cole, M., Sharp, D. W., & Lave, C. The cognitive consequences of education. *The Urban Review*, 1976, 9, 218-233.
- D'Andrade, R. G. A propositional analysis of U.S. American beliefs about illness. In K. H. Basso & H. A. Selby (Eds.), *Meaning in anthropology*. Albuquerque: University of New Mexico Press, 1976.
- Gallistel, R. *The organization of action*. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1980.
- Hilgard, E. R. *Divided consciousness: Multiple controls in human thought and action*. New York: Wiley, 1977.
- Hinton, G. E., & Anderson, J. A. (Eds.), *Parallel models of associative memory*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.
- Hutchins, E. Conceptual structures in pre-literate navigation. Unpublished manuscript. La Jolla, California: Program in Cognitive Science, University of California, San Diego, 1979.
- Kihlstrom, J. F. & Evans, F. J. (Eds.), *Functional disorders of memory*. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1979.
- Kohonen, T. *Associative memory: A system-theoretic approach*. Berlin: Springer-Verlag, 1977.
- James, W. *Principles of psychology*. New York: Holt, 1890. (Reprinted New York: Dover, 1950.)
- Mandler, G. *Mind and emotion*. New York: Wiley, 1975.
- Mehan, H. Structuring school structure. *Harvard Educational Review*, 1978, 48(1), 32-64.
- Mehan, H. *Learning lessons*. Cambridge, Mass.: Harvard University Press, 1979.
- Norman, D. A., *Slips of the mind and an outline for a theory of action*. (CHIP technical report 88). LaJolla, California: Center for Human Information Processing, the University of California, San Diego; La Jolla, CA. 92093, 1979.
- Norman, D. A. *Human learning and memory*. New York: Scientific American, 1980 (in press).
- Roitsch, P. A., Babcock, G. L., & Edmunds, W. W. *Human factors report on the Tenerife accident. Aircraft accident report: Pan American World Airways, Boeing 747, N737 AP: KLM Royal Dutch Airlines, Boeing 747. PH-BUF: Tenerife, Canary Islands, March 27, 1977*. Washington: Air Line Pilots Association, 1979. (Actual report is not dated.)
- Rosch, E. Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, New Jersey: Larry Erlbaum Associates, 1978.
- Stelmach, G. E. (Ed.), *Information processing in motor control and learning*. New York: Academic Press, 1978.

ACKNOWLEDGMENTS

This paper is the direct result of the interactions within our Cognitive Science Program. Two major seminars, each lasting for the entire academic year, served as the testing ground for many of these ideas, the source for others. One was our "Cognitive Science Seminar," a course held primarily for our visiting scholars in the Program, but including local faculty, graduate students, and visitors. (The major contributors were Dave Rumelhart, Bob Buhr, Larry Carleton, Geoff Hinton, Ed Hutchins, Ian Moar, Chris Riesbeck, and Len Talmy. Jim Anderson and Mike Maratsos joined us at the end of the year.) The other seminar was a faculty meeting, organized by Mike Cole, with fierce and useful interchanges among Aaron Cicourel, Mike Cole, Roy D'Andrade, George Mandler, Jean Mandler, Bud Mehan, and Dave Rumelhart. Thus, the section on the relationship between the regulatory and the pure cognitive systems originated in (and was modified by) the Cognitive Science Seminar, while the section on the cybernetic view of the classroom resulted from the faculty seminar's consideration of the studies of Mehan and Cole. The support for all of this was provided by a grant from the Alfred P. Sloan Foundation.

Several people have made useful comments on a draft of this paper, especially Craig Will and Mike Cole. I have incorporated their comments at relevant spots in the paper (Will on motivation, Cole on culture).

Although research results are not presented in this paper, various research projects over the years have helped develop these views. Research on learning (and memory) was supported by a contract from the Office of Naval Research and the Advanced Research Projects Agency, monitored by ONR under contract N00014-76-C-0628, NR 154-387. Research on motor performance was supported by ONR under contract N00014-79-C-0323, NR 157-437. Research on pilot errors and applied human information processing was supported by DARPA and monitored by ONR under contract N00014-79-C-0515, NR 157-434. Support was also provided by grant MH-15828 from the National Institutes of Mental Health to the Center for Human Information Processing.