p. 104 blank

CHAPTER 4

*The Role of Attention in Cognition**

Herbert A. Simon

In our progress toward understanding human thinking and information process, we will gradually, over the years, build bridges from the neurological level of explanation over to the levels of basic information processes with which Posner is mainly concerned, and of complex information processes with which I and some of my colleagues are mostly preoccupied. At the complex end of that bridge, of course, are the cognitive processes we encounter in school learning. A volume like this one is exceedingly useful in measuring, from time to time, the progress in our bridge-building efforts.

I. THE MEANING OF ATTENTION

It might be useful, although perhaps elementary and even redundant for the readers of this volume, to examine what we mean by attention and what its role is in controlling behavior and internal cognitive processes. The authors of the preceding papers provided their definitions—one of them reminded us of William James' (1890) well-known quotation: "Everyone knows what attention is." Nevertheless, we might pursue the definitional question a little further.

The concept of attention is of particular significance for an organism that, apart from its sensory and motor organs, must carry out its thinking serially, rather than in parallel. Although there is still some disagreement about the amount of parallel activity in human thinking, we might agree that those processes that require attention can only go on one at a time or a few at a time. We have a limited capacity for attention, which is best modeled as a serial system.

It is interesting to raise the evolutionary question of why higher or-

* This research was supported by Research Grant MH-07722 from the National Institute of Mental Health, and by a grant from Alfred P. Sloan Foundation.

ganisms generally evince this phenomenon of attention—why their thinking exhibits seriality and an attentional bottleneck. First—and this has been emphasized in the preceding chapters—such an organism needs a mechanism that will sustain its focus over some time on a particular problem context. Organisms are continually surrounded by a very rich and complex stimulus field, from which an enormous amount of information could be extracted each second. In the brains of higher organisms there is also stored a large amount of information that could be evoked at any moment, and thereby impact on, and influence, behavior. The organism needs a mechanism to guarantee that only a very small part of the potentially available stimulus information, and only a very small part of the information potentially available from long-term memory, is brought to bear on behavior during any short interval of time. In the absence of such a restriction, the organism would be buffeted by irrelevancies, and behavior would go off in all directions at once.

If I were to break out suddenly in Latin, that would be disconcerting both to the reader and to me. It is best that what little Latin remains stored in my brain not be evoked except in those rare instances when it becomes relevant to the problem that faces me. Attentional mechanisms enforce that relevance.

Second, on the positive side, a function of attention is to select out stimulus and memory elements that, though not active at the moment, *would* be relevant to the current problem context. Attention brings this information from the senses and from memory into the focus of attention, into the active processor.

A third, equally crucial, function of the attentional mechanisms is to allow a *shift* in focus and context: to break off the continuity of thought in order to respond to real-time requirements of the organism in adapting to its environment. Because bricks do fly through the air sometimes, it is good to be able to notice and dodge a brick even if you are not scanning the horizon for missiles when it comes flying. The attention mechanisms produce such interruptions when sudden, unexpected stimuli present themselves.

From a physiological standpoint, the eye can be regarded almost as two organs. The fovea is designed to transmit information that is in the focus of attention, while the periphery is designed to interrupt attention when urgent information comes from other directions. Thus, one important function of the visual system, and the same can be said of the auditory, is to provide a mechanism for interrupting the ongoing activity of an organism when what is currently being attended to can be postponed in order to deal with high-priority real-time needs.

This interrupting function of attention, which was not much discussed in the preceding chapters is also closely related to some of the emotions, and to the arousal mechanisms associated with emotion. Physiological psychology has paid a good deal of attention to this bridge, via arousal, between emotion and cognition. Lindsley provided a good account of it 40 years ago in the Stevens *Handbook* (1951). It is a theme that provides important linkage between the physiological and educational ends of the continuum represented in this volume.

Thus, both an evolutionary analysis of the functions that attention might perform in increasing the biological fitness of an organism, and the evidence that is provided by physiological research suggest that there are at least two clusters of mechanisms involved in attention: one producing attention shift through interruption, the other having to do with the filtering of information for relevance to the current context—what some researchers, like Anderson (1976), call "priming."

In their chapter, Posner and Freidrich also referred to two separate mechanisms in visual attention. They discussed covert attention and overt attention, although I do not believe they actually used the latter term. But I do not think the distinction they made is the same as the one that I have just developed. At least, I have not succeeded in providing for myself a mapping between the two.

The chapter by Posner and Friedrich focused primarily on selectivity, and they described attention as the system that enables us to select among competing data and to bias our recall from memory and our ongoing stream of thought toward some contexts rather than others. I am not sure whether they meant that description to encompass the interruption of ongoing attention and the selection of urgent messages that were irrelevant to the previous context.

II. LEVELS OF EXPLANATION

Both the chapter by Picton *et al.* and by Posner and Friedrich illustrate beautifully the complementarity of neural and behavioral studies, and the way in which they can provide, and now are beginning to provide, clues for each other. Perhaps the younger among us cannot remember a time when they had little to do with each other, but that time lies not very far in the past. A partial exception was the study of brain damage and other brain abnormalities, which have always had to be studied, in the human organism at least, mainly through their behavioral manifestations.

One of the exciting pieces of information in the preceding chapters is about the strong bridges that are being built between the physiological and the behavioral with new techniques made possible by the CAT scanner and measures of brain metabolism, and our new capabilities for interpreting EEGs, particularly EEG responses to specific stimulus events. The chapter by Posner and Friedrich and the one by Wittrock extend our notions of complementarity to the behaviors that are especially relevant to education.

Another way of looking at the links between the physiological and the behavioral is to look at the things that are linked; and that leads immediately to the topic of levels of explanation, that is, the idea that there is a physiological level of explanation of human cognitive behavior, a level of explanation in terms of basic or elementary information processes, and a level of explanation in terms of complex information processes.

As a characterization of the world of psychology today, this description in terms of levels is rather accurate. It is even reflected in the sociology of our profession. Students earn their degrees from different departments and study different topics and develop different skills if they are interested in research at one of these levels than if they are interested in another. It is thought to be a real mark of courage, or foolhardiness, for a researcher to try to operate at two levels. There are separate literatures which we do not feel obligated to know in detail—or perhaps not at all—if they are not at our own level.

So it is a fact about research in cognition that these three levels exist. The more important question is whether it is healthy for scientific progress that they exist. That, of course, raises the question of what would constitute a satisfactory theory of cognition, a satisfactory explanation of human thinking. Let's look at the matter for a moment from the viewpoint of the educational level downward (or upward, if that is the way your map is oriented). How detailed a theory of attention at one level will be helpful in building a theory at the next higher level? This is the "who-needs-it?" question. Do we in fact need a physiological theory in order to explain cognition at the level of elementary information processes; and do we in fact need a theory of elementary information processes in order to understand how children learn to read or to spell or to do arithmetic?

A theoretical case can be made for sealing off the levels from each other, or nearly sealing them off. We have many successful instances of this in other sciences. The whole field of biochemistry does not wobble and shake every time physicists get a new idea about quarks. It is fortunate that it does not, because there is as yet little stability in the world of quarks. Yet we believe that quarks lie at the basis of elementary particles, particles at the basis of atoms, atoms at the basis of molecules, and so on. We are aware of the relations among the levels, but we can provide scientific explanations for phenomena at one level without understanding the underlying levels in detail.

Science, in fact, is something that can be hung from skyhooks; it does not always have to have a foundation from below. Nineteenth-century chemistry built a powerful theory for understanding chemical reactions and predicting them long before there was any kind of an atomic model beyond the idea that molecules were made of little, mostly structureless, balls.

The other side of the coin is that when the atomic model did reach a certain point of development—let us say, the stage of quantum mechanics in 1926 or therabouts—there began to be major leakage, so to speak, between the levels, and the atomic level began to have the important impact on the chemical level that it continues to have today. The whole new discipline of physical chemistry developed to deal with the linkage between the two levels.

So, I think a correct statement of the situation is something like this: sciences do tend to form hierarchies, where level in the hierarchy has to do with the degree of minuteness of the phenomena and the amount of detail we are concerned with, how microscopic a theory of the phenomena we want. For many purposes it is possible to build a macrotheory at some level of aggregation without knowing the details of the underlying microtheory.

We have a very clear understanding of this phenomenon of layering today derived from our experience with computer programming languages. One can become an expert programmer in PASCAL, LISP, FORTRAN, or some other higher-level language without knowing anything about assembly language or machine language, and certainly without having the slightest hint of what a transistor or a chip is as a physical or functional device. In fact, as the past four generations of computers show—each new generation built with completely different physical devices from the previous one—the programming level of computer behavior can rest on the most diverse physical underpinnings. Obviously, there is an almost complete sealing off of the physical microlevel of tubes or transistors, as the case may be, from the macroscopic level of programming languages.

Although, as the chemistry example shows, sealing-off of levels is never hermetic, the question is what details from below, if we knew them, would affect in a major way the theory at the macroscopic level. That is the practical question of the relation of all of the physiological and information-processing research to education.

Now, the answer might be different if our interests were, say, in the education of normal children in schools having a 30 to 1 student: teacher ratio from what it would be if we were interested in special education or one-to-one tutoring. It might turn out that for the typical school situation we would want a very rough-grained theory, and would be little concerned with the more microscopic levels. However, if we were dealing with children who had learning deficits, and particularly if those deficits were known or strongly suspected to derive from specific physiological deficits, then it might be of great importance to have a fine-grained physiological theory.

Consequently, my hunch is that the physiological research on cognition, and to a lesser extent the research on elementary information processes, will have its first impact, and perhaps its largest impact, on our work with children who have specific physiological problems. It may be a longer time before it will have significant impact on our ways of educating normal children. Perhaps it will never have such an impact, but "never" is a long time. People who make predictions about "never" usually turn out to be wrong.

III. THE EDUCATIONAL LEVEL OF THEORY

Let me now write from the viewpoint of the educational level of theory. I am only half-qualified to do that, because I have taught only in universities, and university teachers do not have to demonstrate professional competence in education—gifted amateurism is fully acceptable. So I am going to put on my professor's hat, instead of my researcher's hat, and try to summarize some of the things we can learn from the experience of teaching.

Most of us would agree that there is a very close connection between attention and learning. Crudely put, the zero-order approximation is simply: No learning without attention. There has even been some research to prove that. We know, for example, that learning does not happen during sleep—there's no use putting on earphones when you go to bed. I do not know of any evidence to refute the proposition that attention to something is a prerequisite for learning about that something. That is the zero-order theory. (insert "t" ok?)

The next approximation, the first-order theory, is that a person who attends for 8 seconds to a chunk (I am using the term *chunk* in the same technical sense that George Miller did in his famous "magical number seven" paper [1956]) will learn that chunk and will have it stored in long-term memory. There is an enormous amount of support for that hypothesis from the rote-learning literature, going back all the way to Ebbinghaus (1885/1913). The actual time might be 9 seconds or 7 seconds, but 8 seconds is a good round number that summarizes well the

7-4

findings in the experimental literature. If you hold your attention on something for about 8 seconds, something is going to be stored in your brain, the kind of thing we now call a chunk, a familiar item.

Those are my zero-order and first-order theories of the relation between attention and learning. What they imply for education is that we should be focusing on how to gain and hold people's attention—how to get their attention applied to whatever it is we want them to learn. Teachers have believed that for quite some time. That is why we tell jokes in class, in order to get attention. (Sometimes there are better ways.)

A somewhat more sophisticated theory would claim that it is also important to consider what is learned. Most experienced teachers know that there is a fundamental difference, although it is still only partly understood, between rote learning and meaningful learning. If we knew how to secure differential attention, we would want to direct it in such a way that the resulting learning would be meaningful rather than rote (Katona, 1940).

Because *meaningful* and *rote* are very fuzzy words, long banned by behaviorism, let me give you at least an example of what I mean by them. There exist rather simple algorithms for solving single linear algebraic equations in one variable. Every child has to learn how to do that in about a week during his first year of high school. It's just the basic skill of being given $4x + 17 = 3x + 2$ and finding what x is.

The algorithm is not usually written out explicitly in the algebra textbook, but it could be written without any difficulty. But we all know that if it were written explicitly in the textbook, some students would suppose that what they ought to do is to memorize it. Having memorized it, they would be surprised that they could not solve any algebra equations.

We know that there is a difference between memorizing that algorithm and acquiring the ability to apply it when appropriate. By the distinction between *rote* and *meaningful* I mean just some such distinction between being able to store away sequences of words and to repeat them on demand and being able to deal with problem situations appropriately.

This leads to a next step. One thing we are learning today from research on complex tasks (e.g., how people solve physics problems or how they make chess moves), and particularly research on how experts do such tasks, is that when experts look at a problem situation in their domain of expertness, they immediately recognize familiar features in the situation, and these turn out to be the principal relevant features for correct handling of the situation (Larkin, McDermott, Simon, & Simon, 1980). I say "turn out to be." Of course, it is no accident; that is why they are experts. They have learned to recognize just those things that are, in fact, relevant to the task.

So the expert who looks at the equation $4x + 17 = 3x + 2$ says, "Oh, there's a 17 on the left-hand side before the equals sign. We have to get rid of that." Then he or she does something appropriate to get rid of it. But recognition of the relevant feature, the unwanted 17, is the first step.

One can argue convincingly from the research evidence that a large part of expert skill consists in recognizing these features of situations. This was the point of much of the research that William Chase and I did on chess experts: showing the experts' ability to recognize important features of chess positions, and to retrieve from memory appropriate responses stored in association with those recognitions (Simon & Chase, 1973). The chess grandmaster says, "Oh, there is an open file; I had better move my rook to it." Of course, he or she does not necessarily make that move, but she or he *always* notices that it is possible, and thinks about it as a possible action to take.

The hypothesis arises from this kind of research that a large component of expert skill resides in the ability to attend, upon seeing a stimulus in the domain of the skill, to the relevant parts of the stimulus; and, through that attention and the resultant recognition, to get access in long-term memory to the information that is required for executing the skill at that point.

Most of you know of the estimates that have been made of the number of patterns an expert might be able to recognize. The estimates always come out to be of the order of magnitude of 50,000, roughly the size of the natural language vocabulary of a college-trained person. A chess master can recognize 50,000 different little clusters of pieces on a chess board, and knows the sorts of things that ought to be done in response to the presence of these patterns (Simon & Chase, 1973).

From an educational point of view, we might argue that it is important for us to understand what the patterns are, in any discipline, that have to be discriminated and learned. And as educators, we have to understand how people can be induced to learn to attend better. If we look at current educational practice, and particularly at current textbooks, I think we will find that in most domains (I have looked mainly at science and mathematics textbooks) there is insufficient attention to the discriminative and recognition skills that are needed for expert task performance.

You have all experienced the professor who stands up at the beginning of the math class and starts writing a proof on the left side of the blackboard. The professor goes all the way across the board, and down in the lower right-hand corner finally writes "QED." You have followed every step and you know he or she has not cheated; every step follows by the laws of mathematics from the previous one. But at the end, you scratch your head and say, "What made her or him think of doing that?"

We are able to check the individual steps. But we lack the recognition capacity of noticing and attending at each step to that aspect of the stimulus that would tell us what step to take next. Attention and recognition play a key role in determining whether we understand what is going on, and are able to acquire the skill of constructing such proofs.

IV. MOTIVATION FOR ATTENTION

I have not mentioned at all yet the topic relating to attention that is most often discussed by educators: What is there in it for the student? What is the motivation for attention? How do you induce students to attend?

One might even better say, "How do you motivate students to be *able* to attend? It is not clear that attention is a wholly voluntary action. The research we have heard reported in this session certainly does not suggest that attention is always voluntary.

We have a good deal of common-sense knowledge about attention, gleaned from our experiences as teachers. We believe that there is advantage from an educational standpoint in intrinsic over extrinsic motivation. That is, we believe that learning is facilitated if the stimuli that have to be attended to are interesting to the learner. We have various reasons for thinking that such intrinsic motivation may provide a more satisfactory basis for education than extrinsic motivation—"If you get an A on your report card, your allowance will be increased." We all know some of the theory and research on which this preference for intrinsic motivation is based.

If it is correct that we need intrinsic motivation to get the kind of attention that facilitates learning, then we need a theory about what makes things interesting. I would simply remark—not by way of criticism, for space was limited—that this aspect of attention did not receive very much attention in this volume.

The largest body of work on the determinants of attention was that done by Berlyne (1965). Berlyne showed that people (and rats) will give their most persistent attention to stimuli that are neither too simple nor too complex. When a stimulus is so simple (relative to the intelligence and previous experience of the organism) that there is nothing new in it, it will be boring and not attended to. When the stimulus is so complex that the organism cannot detect any pattern in it, there is also nothing to attend to, and it will be boring. These pieces of the theory of attention, and the relation between complexity and interest, would add another dimension to the topic as we have developed it in this volume, a dimension that is highly relevant for educational practice.

How detailed a theory of attention do we need for improvement of education? We might decide that we really did not need a neurological theory, but that we did need an information-processing theory. Or we might decide that we need only a theory of the complex processes, and we do not really need the layers below.

My guess is that we are going to opt, eventually, for all of these. We are not going to decide that there is one level on which we should focus exclusively. I cite two examples from the chapters in this volume, one of which has immediate methodological implications.

From Posner's report, we see that there is covert attention or selectivity in the response to visual stimuli even in the absence of eye movements. We do not simply look at whatever it is the fovea is pointed at, but we may look at a highly selective part of the stimulus that is within foveal vision. The methodological implication is that we should not depend on eye movements as a sovereign instrument for saying everything that has to be said about attention, even at the macro level.

The second example is the demonstration that there are several mechanisms having to do with attention, some of which may be under obvious voluntary control, and some of them quite involuntary. These two parts, the voluntary and the involuntary, may play quite different roles in the learning process.

V. FINAL COMMENTS

In my comments, I have not mentioned the kinds of research that are closest to my own interest—computer modeling—and the relation of simulation research to the question of levels of theory. I would like to add just a footnote on that topic.

Suppose we are trying to model performance on school tasks, like solving physics problems (a principal current preoccupation of our research group [Larkin, McDermott, Simon, & Simon, 1980]) or spelling English words (on which my wife and I did some research a short time ago [Simon & Simon, 1973]). Suppose that we were trying to model these processes in a way that would be relevant to improving education in these subjects? How detailed would that model be? Would we need to construct a physiological model? Can we get away with a model that only goes down to the level of elementary information processes of the sort that Posner has discussed? Or can we even use a model that incor-

12-14

9K

porates rather gross and complex information processes without analyzing them in detail?

I do not think we need to accept a doctrinaire answer to this question. We can usefully understand such processes at all three of these levels. For example, in the spelling case, with a rather aggregate and coarse model, we were able to predict rather accurately what spelling errors would be made with a word like "responsible." Moreover, from our detailed predictions of the spelling errors, we were able to draw some more general conclusions as to the value of certain kinds of educational procedures in improving spelling. Hence, modeling human performance in school subjects can give us some clues as to what help we could get from a coarse theory and when we might need one carried to a more detailed level.

VI. REFERENCES

Anderson, J. R. (1976). *Language, memory, and thought.* Hillsdale, NJ: Erlbaum.

Berlyne, D. (1965). *Structure and direction in thinking.* New York: Wiley.

Ebbinghaus, H. (1913). *Memory* (H. A. Ruger & C. E. Bussenius, Trans.). New York: Teachers College. (Original work published 1885)

James, W. (1890). *The principles of psychology.* New York: Holt.

Katona, G. (1940). *Organizing and memorizing.* New York: Columbia University Press.

Larkin, J., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. *Science, 208,* 1335–1342.

Lindsley, D. B. (1951). Emotion. In S. S. Stevens (Ed.), *Handbook of experimental psychology.* New York: Wiley.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63,* 81–97.

Simon, H. A., & Chase, W. G. (1973). Skill in chess. *American Scientist, 61,* 394–403.

Simon, D. P., & Simon, H. A. (1973). Alternative uses of phonemic information in spelling. *Review of Educational Research, 43,* 115–137.