ATTENTION, PERCEPTION, AND MEMORY

Alexander Pollatsek, Caren M. Rotello
Department of Psychology, University of Massachusetts, USA

Keywords: Selective attention, divided attention, limited capacity, pattern recognition, visual "neglect," features, eye movements, short-term memory, long-term memory, implicit memory, explicit memory, procedural memory, eyewitness memory, context effects, amnesia

Contents

Summary

We begin this article by discussing Donald Broadbent's book, which was an important part of the "cognitive revolution" that led to an explosion of work on human cognition. This book, and our introduction, indicates the basic framework that people used (and still use, to a large extent) to think about the basic mental processes involved in attention, perception, and memory. However, much of what we discuss goes beyond this basic framework and indicates ways in which research has gone well beyond the model. For example, in our discussion of perception and attention, we indicate that the model said that people encode basic sensory information at a different level than they encode meaning, but the model said little about the details of how we do either. We've tried to describe how psychologists have learned a lot more about these processes, but that there are still issues about which we understand little (e.g. how individuals recognize an object). Similarly, in memory, Broadbent's basic box model pointed out the distinction between various types of memories and suggested, in a general way, how they should interact. However, the model said little about how long-term memory works. Moreover, it suggested that memories were simply filed (as in a library) and that remembering something depended more or less on whether it had decayed or not. Our discussion of memory indicates that memory is much more complex, and that whether someone remembers something depends a lot on the retrieval process—for example, how one

tests memory (e.g. a recognition, recall, or implicit memory test), and how the conditions at retrieval match those at test.

1. Introduction

In 1958, Donald Broadbent published a book, *Perception and Communication*, that attempted to provide a unified framework for discussing perception, attention, memory, learning, and various other cognitive activities. This framework has been adapted and extended since then, but it remains a central metaphor for how most cognitive psychologists think about many topics in cognition (especially perception, attention, and memory).

One of the central insights of Broadbent's book was that virtually any task we do, no matter how apparently simple, is quite cognitively complex, using a lot of machinery of the brain. Hence, although perception, attention, memory, and learning seem like quite disparate topics, laboratory tasks that are designed to study one of these topics almost always involve the other processes as well. This is perhaps clearest in thinking about a memory task, as most memory tasks involve perceiving some sort of cue that triggers a memory. Thus, one has to understand how the cue was perceived, which involves the degree of attention paid to the cue. Furthermore, whether the memory is successfully retrieved is clearly a function of whether the person tried to "memorize" the information in the first place, and, if so, how well and under what conditions it was learned or studied. In this brief space, clearly we can present only the barest outline of what has been learned on these topics since the end of the nineteenth century. By using Broadbent's model as a focus, we will be biasing our summary toward certain topics and (necessarily) leaving out other important areas of inquiry. We hope our loss of completeness will be compensated for by an increase in cohesion.

The focus of Broadbent's model was attention, especially *selective attention*. A practical situation that he was interested in was the air-traffic controller who is listening to a variety of messages from various pilots about to land and has to decide who might be having difficulty, whom to respond to, etc. This situation involves two task components: (1) a *divided attention* component, where the controller is attempting to process all the messages; and (2) a selective attention component, where the controller is attempting to process a single message carefully (perhaps a pilot in trouble) and trying to respond to the message as well as possible. In fact, the data from many experiments in both the field and laboratory had suggested to Broadbent and others that divided attention was quite difficult, and was often accomplished by rapidly switching back and forth between messages rather than by truly processing them all at once. In contrast, the data indicated that the task of selectively attending to one message and ignoring the others was quite simple, with one proviso: that the attended message be distinguished from the others by a low-level sensory cue, such as spatial location or type of voice. These observations were critical in formulating Broadbent's conception of the human information processor.

Before discussing Broadbent's model, it might be worth briefly discussing some of the methodology of the experiments that led to the model. As indicated above, the experiments were inspired by practical problems in listening, so they focused on selective attention in processing language. In the basic task (devised by Colin Cherry) several messages were played at once to subjects (usually over earphones) and the subjects were asked to "shadow" one of the messages (i.e. repeat it verbatim). The ability to shadow (with a lag of a couple of seconds between what comes in the ear and what goes out the mouth) takes a bit of practice, even with no competing messages, and people were tested after they had learned this skill. The central finding was that if the *selected* message (i.e. the one subjects were instructed to shadow) differs from the other message or messages by a physical cue such as which ear it comes in or if the selected message is in a man's voice and the other message is in a woman's voice, people shadow the message just as well as when there is no *irrelevant* message. In contrast, if the two messages do not differ by a physical cue, shadowing the selected message becomes very difficult. This is true even when the semantic content (e.g. topic) of the two messages is quite different.

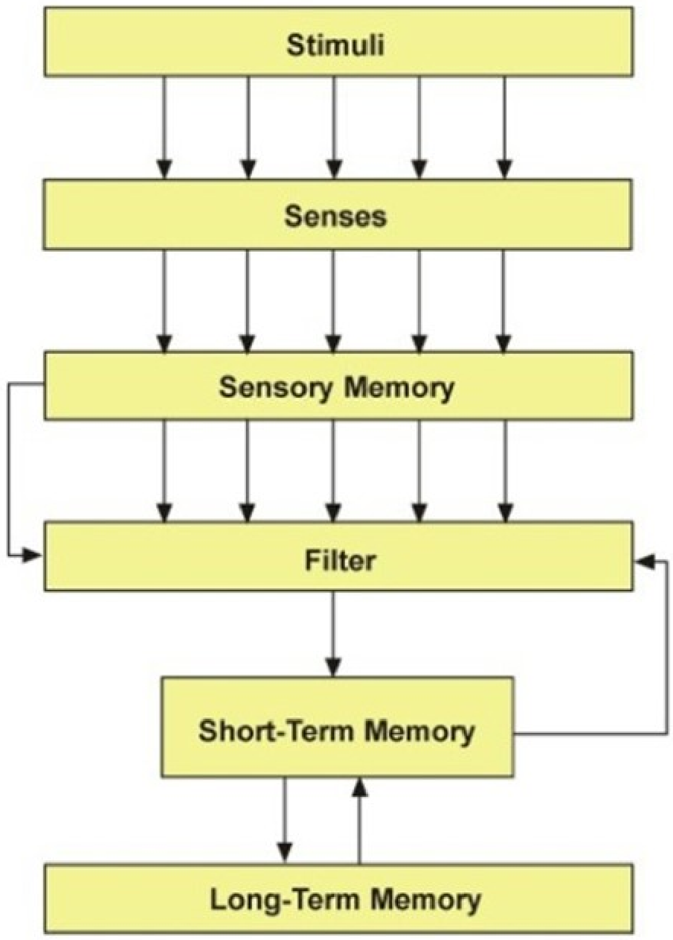

Figure 1. A schema representing the essence of Broadbent's model of cognition

This distinction between low-level sensory analysis and higher-level semantic analysis is key to Broadbent's model. In particular, the above type of finding suggested to him a) that they were performed by different systems; b) that the low-level sensory analysis was logically prior to the semantic analysis; and c) that efficient selective attention could be guided only by the sensory analysis. These verbal statements are captured by a significant portion of the model shown in Figure 1, which is a minor adaptation of

Broadbent's model. It is worth talking this figure through in some detail. Each box represents a brain system, and the model is quite modular in that these systems are viewed as largely independent of each other, except to the extent that they "talk" to each other (represented by the arrows between them).

Specifically, let's see how Broadbent represented the selective attention task with the model. The first box is called "senses" and wasn't specified particularly clearly in the model, but it presumably involves taking the raw physical data of the auditory message and analyzing it for spatial, location, pitch, and other such characteristics. These analyses are then passed on to the "sensory memory." The name sensory memory implies that these sensory characteristics are not only analyzed, they "hang around" for a while—automatically. (A key question that became a major topic of inquiry in the 1970s was the characteristics of these stores, including how long information remained.) This sensory information then heads to the "filter" which, if we are successfully doing a selective attention task, will allow only one message through. The key point of the model is that, until "the message" passes this filter, it has been analyzed only for these low-level sensory properties, and so the filter can filter only on the basis of them. After going through the filter, the message gets through to the central processor, where the meaning of the message is extracted. This central processor has two characteristics. First, unlike the sensory boxes before it, it is active and its behavior can be changed by the intentions of the person. Second, it has *limited capacity*, in that it can process only a limited amount of material at one time.

The way the model explains the selective listening data should thus be clear. The filter understands only physical cues, because that is all that is processed prior to the filter, so that the filter can only be set to select on the basis of these physical cues. Because the filter is all-or-none, it can effectively screen out unwanted messages on the basis of physical cues efficiently. In contrast, if the messages do not differ on the basis of physical cues, both of them go through to the central processor, which can sort out things on the basis of meaning, but because of its limited capacity it cannot do this sorting out very efficiently.

The fact that division of attention is difficult was explained in more than one way. One possible strategy people could have in a divided attention task (e.g. monitor a bunch of pilots' messages at once) is not to attempt any filtering and to let all of these messages come in and then to try to sort out whether anything needs to be attended to. This would be something like trying to find a potato in a thick stew and it would be hindered by the limited capacity of the central processor. The second strategy people could adopt is to attempt to attend selectively to each channel in turn rapidly (like going through the channels of a TV quickly). This would be hindered by getting only fragmentary information from each message, so that much of the information would need to be reconstructed, and could be slowed down if there was a non-trivial time needed to switch between information channels.

The rest of the structure in the model is not used heavily in the attention tasks, but is worth commenting on briefly here. First, the term "sensory memory" is perhaps a bit misleading and should be termed "sensory memories," as there are undoubtedly different stores for the different senses. The one in vision, commonly called "iconic" memory, lasts for only a fraction of a second and is probably of limited practical use. However, the one in audition, commonly called “echoic memory,” is undoubtedly a major part of the functioning of auditory processing. Estimates of auditory sensory memory range from fractions of a second to many seconds. In fact, there is thinking now that there are really two auditory sensory stores. One lasts a fraction of a second, and is merely an afterimage, but the other lasts several seconds and plays an important role in a lot of cognitive tasks. Indeed, auditory rehearsal of material undoubtedly uses such a system; if all auditory stimulation decayed instantaneously, it would be quite difficult to understand any spoken sentences because the first part of the sentence would be lost from memory before the second part was even spoken.

The central processor in Broadbent’s model is equivalent to what is usually termed short-term memory (STM). It is assumed that material stays in STM on the order of 10–30 seconds. Some interesting work done at about the time of Broadbent’s book dramatically illustrated a distinction between STM and long-term memory (LTM). There were patients suffering from what is known as Korsakov’s syndrome (usually the result of chronic alcoholism). These people seem to have a relatively intact STM. For example, they can remember about the standard size list of digits (6–8) in an STM task. More ecologically, when conversing with them, they appear to have little or no problem understanding conversations, which draw on STM (e.g. remembering the first part of a sentence to understand the second part of the sentence). They also seem to have relatively intact LTM for events that occurred before their problems began. For example, one of the authors met a Korsakov syndrome patient who conversed at length about his childhood on a farm and made subtle distinctions between various types of farm machinery. In contrast, these patients seem to have lost the ability to put down any newer long-term memories. They have little idea what they were doing half an hour ago and apparently can remember virtually nothing of what has happened to them since their problem began. For example, the patient described above thought he had been at the hospital for a day or two even though he had been there for 10 to 15 years. That is, it appears that the connection going from STM to LTM is severed (or at least severely damaged), preventing any new memories from being laid down. In contrast, the connection going from LTM to STM seems relatively intact, as they can retrieve old long-term memories. As we shall see later, this description of these patients (and the relation between LTM and STM) is a bit oversimplified.

The major remaining box is “long-term memory” (in contrast to the shorter-term memories in the sensory memory and STM boxes). The need for such a box to explain human behavior is perhaps obvious. What may be less obvious is that there are any qualitative distinctions between longer-term and shorter-term memories. This is still a matter of controversy, although we think most people would agree that there is some distinction. What may be still less obvious is that LTM is needed in the selective attention task, but it is. Presumably, in order for the instructions to the experiments to be understood, all sorts of long-term memories need to be activated (also for the meanings of the words in the messages to be understood). Memories in LTM were assumed to undergo some sort of decay or deterioration over time, but at a very slow rate.

We have presented this model in some detail, as it is a conceptual model that runs through virtually all of cognitive psychology. Even people who are critical of the model often invoke many of the terms as explanatory concepts. Needless to say, however, the model is hardly a complete explanation of human thinking or even of human attention, perception, memory, and learning, as it focuses on a few processes and ignores many other things. For example, it posits that the meanings of spoken words are perceived, but it doesn't say anything helpful about *how* this occurs. Moreover, people became suspicious of the simplicity of the model. For example, could selective attention really operate that simply, where the meaning of the unattended messages was completely blocked? Or, does processing of meaning really flow from the processing of "features" such as location, pitch, etc. In spite of these potential problems and shortcomings, however, the model (and its successors) led to many interesting experiments looking at the details of how information was processed and stored, especially testing the validity of assuming various "short-term memories" and examining the properties of each and has proved to be a valuable framework for thinking about these topics.

2. Attention

One of the pieces of data supporting Broadbent's notion that unattended messages were filtered out completely was that, in experiments by Cherry, some words appeared in the unattended message many times and yet people had no recognition memory of them. This appeared to be strong evidence for the assertion that the filter was all-or-none, as posited by Broadbent. Nonetheless, several experiments appeared to indicate that this assumption was wrong. In one, the subject's name appeared in the unattended channel, and the people almost always detected it. Thus, their name—because of its meaning—was getting through to consciousness. Other experiments showed, using physiological measures, that unattended words were being processed even though people were unaware of them. Thus, unfortunately, there are problems with the elegant simplicity of Broadbent's model. The first finding shows that, on some level, messages are breaking through to consciousness because of their meaning and the second finding shows that awareness and/or memory for things does not necessarily mean that things are not processed for meaning (i.e. that there is "subliminal perception").

A model proposed by Treisman, a student of Broadbent, expanded the model to explain some of these findings. She posited that there was a "dictionary" or *lexicon* that was intermediate between the sensory feature information and the filter. This means, for spoken discourse, that the meanings of words, in some sense, are "automatically" activated. She posited that there were thresholds for activation based on factors such as the frequency of a word or its importance. More frequent or more important stimuli have lower thresholds, and hence take less stimulus activation to lead to identification. She also posited that the filter only attenuated (i.e. weakened) unattended messages. Thus, even weakened messages (especially if they are important) could cross the threshold and make it into the central processor.

Although these changes may appear minor, they changed the model profoundly, as they imply that activation of meaning (for individual words) does not require attention, although attention may be instrumental in their being consciously perceived and used for further processing (e.g. understanding what a sentence means). In the visual domain, the implication was that "objects" were analogous units that could be processed automatically. Thus, this modification creates a different dichotomy of what can be processed without need of attention. In this modified model, individual words (either spoken or printed) or individual objects can be processed without attention (although attention might help in processing them), but attention is needed to process things at higher levels—either understanding phrases or sentences or perceiving "scenes" that consist of several spatially ordered objects. Although there is still some controversy about the limits of attention, there is probably a consensus that this dividing line is about right. There is some evidence, however, that some higher-order percepts may be processed as "units" given a lot of practice. That is, perhaps very familiar scenes or phrases may become just as much "units" as words or individual objects.

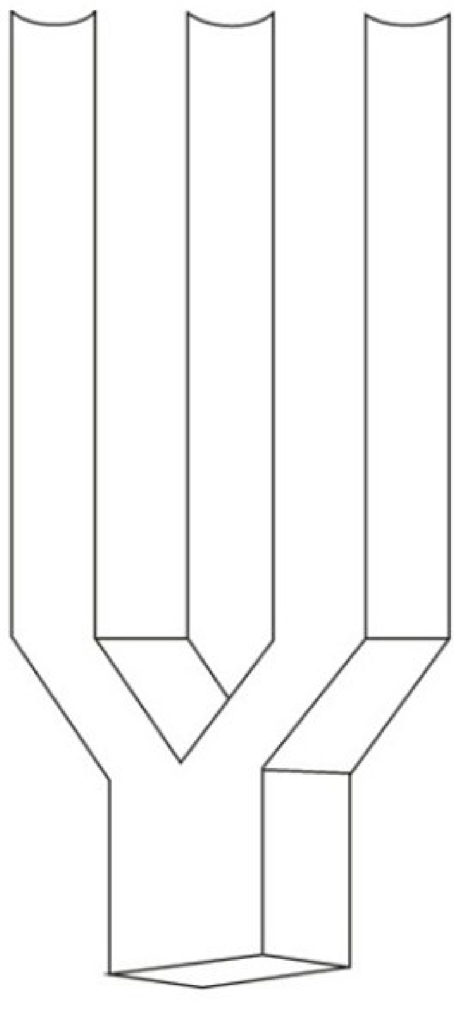

Figure 2. A drawing of an impossible figure

It is worth noting, however, that the impression that we can perceive a scene in an instant is largely an illusion. Perhaps the clearest illustration of this are the many optical illusions such as the tuning fork that starts out with two prongs and winds up with three (see Figure 2) or the endlessly ascending staircase in the well-known Escher drawing. If one could, in fact, process such visual displays all at once, their impossibility would be apparent; however, it often takes people minutes to find the absurdity in such a drawing, even when they are looking for it. This implies that only small regions of the display are attended to at any one instant, and if each small region "looks normal," absurdities are difficult to detect. Thus, finding the absurdity in these illusions requires a more abstract coding of each region and becomes more of a problem-solving task than a perceptual one.

Most of the early work on attention assumed that attention could easily be drawn to a stimulus, whether auditory or visual, by its location in space. Presumably such attention could be drawn either by instruction (e.g. "pay attention to the front door") or by an abrupt stimulus, such as motion of the front door or a loud sound coming from that direction. This leads to the question of whether people can direct their attention at will to any location or whether there are limits. One interesting finding in this regard is called "inhibition of return." What this means is that if you have just paid attention to location A and then switched your attention to location B, you have more difficulty returning your attention to location A than in switching your attention to location C. (This inhibition is relatively minor, as the inhibition slows things only a fraction of a second.) Besides this minor difficulty of attending that we all experience, there are patients who have difficulty attending to regions of space. This phenomenon is known as *neglect* or *extinction* (depending on its severity) and is associated with damage to visual areas in the *parietal lobe* of the visual cortex. Patients who have the severest deficits tend to act as though events in one half of their visual world (either the right or left half, depending on the location of the damage) do not exist. For the patients with less severe damage, the syndrome is something like an extreme form of inhibition of return. This is perhaps best illustrated in the clinical test given to such patients. (Let's say the patient is having difficulty perceiving things to his left.) The doctor tells the patient to look at his nose and wiggles a finger a bit to the right of her (the doctor's) head so that it is in the left visual field of the patient. If that's all that's happening, the patient will often see it. But if the doctor wiggles fingers on both sides of her head, the one in the patient's left visual field effectively disappears and he won't see it. That is, his attention is drawn to the other visual field and it can't be drawn away.

-
-
-

Bibliography

Broadbent D.E. (1958). *Perception and Communication*, 338 pp. London: Pergamon. [One of the key books in the "cognitive revolution" in psychology. Still an interesting book worth reading.]

Bruce V., Green P.R., and Georgeson M.A. (1996). *Visual Perception, Physiolology, Psychology, and Ecology*, 3rd edn. East Sussex: Psychology Press. [A good modern text on visual perception.]

Neisser U. (1967). *Cognitive Psychology*, 351 pp. Englewood Cliffs, N.J.: Prentice Hall. [The first textbook that attempted to integrate perception, attention, and memory from a modern perspective and the book that coined the term "cognitive psychology." Perhaps still the best text.]

Posner M.I. (1978). *Chronometric Explorations of Mind*, 271 pp. Hillsdale, N.J.: Lawrence Erlbaum. [A good introduction to attention.]

Rayner K. and Pollatsek A. (1989). *The Psychology of Reading*, 529 pp. Englewood Cliffs, N.J.: Prentice Hall. [An elementary text that covers the reading process and extensively deals with the perceptual aspects of reading, including a tutorial on eye movements. See, especially, Chapters 3 to 5.]

Tulving E. and Craik F.I.M., eds. (2000). *The Oxford Handbook of Memory*, 700 pp. New York: Oxford University Press. [This book provides an extensive introduction to memory in the form of short, easy-to-read chapters on a wide variety of topics. Of particular interest to readers of this chapter are the chapters on short-term memory, source monitoring, memory for life experiences, memory for emotional events, and memory systems.]

Biographical Sketches

Alexander Pollatsek is professor of psychology at the University of Massachusetts, Amherst, and has been at the University of Massachusetts since 1969. He received a B.S. in chemistry from the University of Michigan, an M.A. in chemistry from Harvard, and an M.A. in mathematics and a Ph.D. in psychology from the University of Michigan. His research interests are varied, although his current focus is on visual perception and reading. Much of his work on reading is contained in the book *The Psychology of Reading*, cited in the bibliography. He has also published papers on decision making, memory, and reasoning. A current secondary focus is on reasoning in science, mathematics, and statistics, with a focus on improving instruction in these areas. His work is supported by grants from the National Institutes of Health and the National Science Foundation.

Caren Rotello is an associate professor at the University of Massachusetts, Amherst, where she has worked since 1993. She received a B.S. in psychology and mathematics from the University of Michigan and a Ph.D. in psychology from Stanford University. Her research is focused on models of recognition memory and categorization and is supported by a grant from the National Institutes of Health.