COGNITIVE METHODS

and Their Application to Clinical Research

Edited by
Amy Wenzel
David C. Rubin

AMERICAN PSYCHOLOGICAL ASSOCIATION • WASHINGTON, DC

5

SELECTIVE ATTENTION TASKS IN COGNITIVE RESEARCH

NELSON COWAN

The goal of this chapter is to use research on selective attention within the field of cognitive psychology to provide a background that could guide researchers of clinical populations. The corpus of research on attention is huge but my policy nevertheless is to interpret the concept of attention broadly enough to make clear the rich interconnections between attention and other mental functions including perception, memory, and thought, and to point out the many options available if one is interested in studying attention.

Many essays on attention begin with a quote from the chapter on attention within William James' (1890) *Principles of Psychology*, in which he asserts, "Every one knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought." Although this statement is true, it seems equally telling that people do not often find similar quotes about how everyone knows what habit is, what memory is, what reasoning is, and so on. Perhaps more than most other aspects of the mind, there is a paradoxical difference between our ready folk knowledge of the concept of attention and

This chapter was written with support from National Institute of Health Grant R01 HD-21338.

the difficulty of understanding it, or even recognizing it, in clear and definite terms. For example, it is not as if, by looking at another person, one can consistently tell whether he or she is paying attention. In an episode of the cartoon strip *Dave* (by David Miller), our tongue-tied hero sits with his girlfriend, who is torturing him with a conversation about their relationship and is gazing intently at him while he silently muses, "How do I know that there is a god . . . and he's a guy? Because listening intently and zoning out look exactly the same!" On the whole, attention is known as a concept that is difficult to pin down and operationalize. I suspect that many a researcher has started a talk on some aspect of attention, only to be interrupted by a member of the audience who asks for a definition of what is meant by attention, perhaps mischievously expecting that the speaker will not be able to answer in a satisfactory manner. But life, and lectures, must go on despite the difficulty of defining attention.

The present chapter differs from most others in this book in that the topic is broader; attention is not a specific type of method as is, say, the Stroop effect, and it probably covers even more ground than do other relevant phenomena such as, say, autobiographical memory. Indeed, the deployment of attention affects almost everything humans do, and numerous methods have been developed to study it. The method must be shaped by the aspect of attention that is of interest and the starting assumptions that one is willing to make about how it operates. The floor plan for the present chapter consequently is as follows. First, I take stock of reasons why clinicians might be interested in seriously studying attention. Second, because a chapter cannot do justice to the broad topic of research on attention, some of the available book references are discussed. Third, a taxonomy of tasks that have been used to study attention are offered briefly, with a few examples and illustrations. Fourth, a key substantive question that can be addressed on the basis of such tasks—namely the question of what factors appear to cause attention or inattention—is discussed. Fifth, and finally, a more in-depth description of one line of attentional research, from the author's laboratory and collaborations, is presented to illustrate the methodological issues that must be taken into account in order to carry out a study of an aspect of attention.

WHY A CLINICIAN MIGHT STUDY ATTENTION

Why might a clinician or clinical researcher be interested in attention despite such difficulties? Perhaps because it is a key concept in understanding the human mind. It has been clear at least since James (1890), and almost certainly earlier, that there is a limit to how much information a human being can deal with at once, or within a limited period. I can think of at least three ways in which this human predicament is likely to be clinically relevant.

1. In abnormal individuals, there might be an abnormal attentional profile of diagnostic relevance. People with psychopathologies often attend to stimuli related to certain domains of clinically relevant interest (e.g., sex, violence, or emotional dependence) more than do normal individuals. Yiend and Mathews (see chap. 6, this volume) discuss this type of mechanism in detail. Conversely, it is theoretically possible that there could sometimes be a tendency to avoid such stimuli and therefore to focus attention on competing stimuli, or perhaps on no stimuli at all. In either case, an abnormal profile of attention might provide information about the topics of special concern to the patient, about the degree of severity of abnormality, and perhaps about the correct diagnosis. Yiend and Mathews offer numerous examples of different clinical populations (e.g., depressed versus anxious patients) responding differently in the same attention tasks.
2. An obsession with particular thoughts or types of thoughts can prevent attention from operating flexibly and normally. This kind of attentional abnormality can be very debilitating in its own right and warrants treatment. Of course, knowledge of research on attention is necessary for an effective treatment to be developed.
3. Some types of psychopathology (e.g., schizophrenia or attention-deficit/hyperactivity disorder) may physiologically cause cognitive impairment, including a debilitating deficit in the control of attention. Here, the treatments might differ from the case in which attentional deficits result indirectly, as a secondary consequence of obsessive thoughts.

EXTENSIVE SUMMARIES OF ATTENTION RESEARCH

In an important recent review of research in attention, Luck and Vecera (2002) stated, "The term *attention* has been used in the title or abstract of over 40,000 journal articles, books, and book chapters in the past 30 years. This greatly exceeds the 8,300 works that have used the term *emotion*, and it almost equals the 48,000 works that have used the term *memory*." Given the breadth of the topic of attention, it is best to start by describing a few resources that would allow the in-depth investigator to find out about a wider variety of experimental procedures and results than can be discussed in this chapter. Styles (1997) has written a short, general text on attention that is fairly accessible and Pashler (1998a) has written a longer but still-accessible one. A number of special, edited volumes on attention research also are worthwhile (e.g., Baddeley & Weiskrantz, 1994; Pashler, 1998b; Shapiro, 2001).

The book reporting proceedings from the regular *Attention and Performance* conference series has many useful entries, and the *Annual Review of Psychology* (Pashler, Johnston, & Ruthruff, 2000; Rensink, 2002) includes two recent summaries that are of relevance. Cowan (1995) summarized evidence of the interrelatedness of attention and memory. For historical insight regarding the development of the attention concept, one could read the chapter on attention by James (1890) and then the seminal book by Broadbent (1958). Shiffrin (1988) gives a comprehensive overview of aspects of attention research, including research on automatization of processes, that greatly increased the sophistication of researchers' understanding. Näätänen (1992) gives an in-depth summary of procedures that were used to examine the brain representation of attention, and that field has been growing rapidly (e.g., Braun, Koch, & Davis, 2001; Humphreys, Duncan, & Treisman, 1999; Parasuraman, 1998). Davies and Parasuraman (1982) and Hancock and Desmond (2001) give in-depth coverage of sustained attention, or vigilance, and practical aspects of it. There is enough overlap among these sources that it is not necessary to read all of them to gain a useful, working knowledge of research on attention. They do, however, reveal a range of theoretical perspectives and emphases as well as a great deal of empirical research.

MAJOR VARIETIES OF ATTENTION TASKS

There are various ways in which authors have classified the many different tasks that typically have been used to examine attention. Luck and Vecera (2002) provided a useful classification that distinguished between cuing paradigms, search paradigms, filtering paradigms, and dual-task paradigms and that classification will be followed here. Before that, however, a word on dependent measures is in order.

Dependent Measures in Attention Tasks

Within each category of task, some procedures measure the proportion of trials in which a correct response is given, or the difference that the direction of attention makes for the correctness of a response. Such procedures can involve very rapid presentations in order to make the task sufficiently difficult to discriminate between conditions or between individuals. Other procedures focus upon the reaction time to make a correct response. Ideally, both the reaction times and accuracy of responses must be taken into consideration in some way because there can be tradeoffs between the two. Specifically, some participants might answer more slowly so as to be sure to answer correctly, whereas others might take a greater risk of making incorrect or inaccurate responses in order to respond more quickly. This emphasis on accuracy versus speed falls along a continuum and can be modified with task instructions.

As a further complication, one can imagine tasks in which accuracy itself falls along a continuum. Such tasks are not very common in the laboratory but might be critical in real life, as for soldiers or police officers in a gun battle. Whereas either too slow or too inaccurate a firing upon a threatening attacker may mean death for the officer as a result of return fire, hitting the attacker just slightly off-target, that is, in an unintended anatomical location, still can have big consequences for either party but is not as risky for the officer as is missing the target completely. In a less life-threatening but still important example, a clinician has to respond to a statement by a patient within a relatively short time window but too hasty a response can throw the therapy onto the wrong track. People probably differ in their typical, overall tendency to be impulsive or reflective in responding in a particular situation.

Finally, in some tasks it is not just the correctness of responding that is at issue, but the particular response that is given. Sometimes the nature of erroneous responses is at issue (in error analysis) and other times there may be no correct answer or multiple possible correct answers, each of which has a different significance to the experimenter. All of these possibilities can be found in cuing, search, filtering, and dual-task procedures, although the emphasis is most often upon the pattern of proportion correct and reaction times across different conditions.

Cuing Paradigms

In cuing paradigms, as Luck and Vecera (2002, p. 240) put it, "a stimulus or an instruction is used to lead subjects to expect a specific source of inputs, and then the processing of this source of inputs is compared with the processing of other inputs." A well-known example is the procedure of Posner (1980) in which a participant is informed with a visual cue as to whether a target will appear on the left or right of visual fixation. The task is to respond with a buttonpress indicating which side of the screen the target is on, as quickly as possible. When a directional cue appears the cue usually (80% of the time) points to the correct location but on other trials (20% of the time) it is misleading, pointing to the incorrect location. On still other, control trials, there is no directional cue. For example, in some experiments the cue was a central arrow pointing left or right or a plus sign. This type of procedure produces an advantage for a correctly cued location and a disadvantage when attention is sent to the side opposite from the target. Investigators have drawn a distinction between *peripheral* or *exogenous* cues that appear at the location where the target probably will appear, which recruit attention to their locations rather involuntarily, and *central* or *endogenous* cues, such as the central arrow pointing to one side, which draw attention to the center but then allow the participant to move attention voluntarily in the direction in which the target is to be expected.

The value of cuing paradigms is that one can examine factors that recruit attention or allow attention to be shifted voluntarily, the speed and timing of these attention-shifting mechanisms, the difficulty of returning attention to a region that was recently attended, and other aspects of attention-shifting. Attention-shifting to a peripheral cue occurs automatically and, in fact, takes effort to suppress. Attention-shifting away from a central arrow toward the direction in which it points takes effort because the automatic tendency to direct attention to the arrow itself must be overcome. In this type of procedure one can usefully measure the response reaction time for trials with valid, invalid, and neutral cues and also eye movements to the cues. Given that the real-world application of attention often involves a shifting of attention to understand rapidly emerging threats and crises, cuing procedures can be quite informative. In clinical work, one could make a distinction between, say, an automatic aversion to threatening stimuli versus difficulty in using effortful processing to direct attention to threatening stimuli (see chap. 6, this volume). These mechanisms have different psychological implications. It is perhaps easier to combat a tendency that occurs with effort and voluntary processing than one that occurs automatically because the effortful processing is more visible to consciousness (e.g., Shiffrin, 1988).

Search Paradigms

In search paradigms, one item or several from a target set are to be found as quickly as possible within an array of items presented in the visual field (or, in principle at least, within a field of stimuli presented in any sensory modality or even across modalities). In a *memory search* procedure, there is just a single probe item to be examined but one searches through a memory representation to determine whether that probe was present within a previously learned or encountered set of items (Sternberg, 1966). In contrast, in a simple *visual search* procedure, there is only one item in memory, and it must be compared with a larger array of items in the visual field. Memory search and visual search also can be combined in the same experiment by having participants search for multiple possible targets, held in memory, within a multi-item array (e.g., Schneider & Shiffrin, 1977).

One benefit of search paradigms is that one can examine what factors influence the ability to use attention to find things in the environment or within a mental representation of a stimulus set. One also can examine the pattern of search times or accuracy as a function of how many items there are to be searched. For very rapid searches or searches in which items are processed in parallel, little time should be added to the reaction time for each additional item in the set to be searched. When searches must be made in a method in which only a small portion of the processing can be completed at once, say one item at a time, then each additional item in the set to be searched increases the reaction time substantially.

A key example is the visual search procedure used by Treisman and Gelade (1980). They distinguished between feature searches and conjunction searches. In an example of a feature search, a screen of dots might be presented, and one would have to press a button indicating if a red dot is present (among a field of blue dots). If a single red item is detected anywhere in the array, one can immediately ascertain that the target is present. The feature search might also involve, say, finding a triangle among circles. The reaction time for such a search does not increase much as a function of the number of objects in the display. In a conjunction search, one might be looking for a red circle among a set of items that includes at least two types of distractors: red triangles and blue circles. It is not enough to determine if the target features are present; one must carefully search to determine if they are present in the right combination. In this case, search time generally increases linearly as a function of the number of items in the display. Moreover, the slope of increase is twice as high when the target that one is searching for is absent from the display than when it is present. That is because one must search all items in a display to determine that a target is absent whereas, when a target is present, on average one finds it after searching through half of the items in the display. This type of procedure could be used to determine what the units of perception are. Clinically, we could speculate that certain objects of one's obsessions or concerns tend to be treated as features rather than conjunctions although to my knowledge there is no such research as yet.

Filtering Paradigms

In filtering paradigms, instructions given before a stimulus field indicate which part of the stimuli are to be attended and which others are to be ignored. One of the oldest examples of this type of paradigm is the selective-listening procedure in which competing spoken messages are presented to the two ears (Broadbent, 1958; Cherry, 1953). In Cherry's procedure, for example, one hears a message in one ear and a different message in the other ear and must quickly repeat (or *shadow*) everything that is said in one ear. It is clear that people can do a reasonably good job of processing only one coherent speech message at a time and that very little of the other message(s) is subsequently recalled, presumably because it could not be processed.

Filtering paradigms are useful because it is possible to learn the conditions under which distractions can or cannot be excluded from processing and the ways in which processing is impaired as a result. Learning all of this has considerable real-world application inasmuch as we live in a world with multiple concurrent stimuli and, much of the time, must exert effort to stay on task. Yiend and Mathews (see chap. 6, this volume) briefly mention early clinical work in which an emotional item is embedded within the ignored message in selective listening. To the extent that the subject's processing system is tuned to the issue of the emotional item, the item may tend to be noticed, resulting in a disruption of attention and revealing some of the subject's mental properties. An even older filtering procedure is the task by Stroop (1935/1992) discussed in previous chapters within this volume. When a color word is printed in a conflicting color of ink, people cannot name the color of the ink without massive interference from the printed color word, provided that they know how to read it. Also commonly used is the flankers task, in which the participant is to respond to a central target but not to distractors placed on the left and right of the target (see the review by Eriksen, 1995).

We have thus examined both auditory and visual filtering procedures. Although visual filtering studies traditionally have been easier to conduct with simple equipment, it is not clear that this is still the case with modern computers. In one way at least, auditory tests are especially useful: They allow ignored and attended stimuli to be presented at the same level of sensory acuity (because humans cannot close or redirect their ears). In contrast, ignored visual information tends to be perceived out of focus or peripherally, with less sensory acuity than the attended items, which can be focused with high acuity on the fovea and the closely surrounding regions of the retina. The uniform acuity found in audition makes it easier to avoid confounding sensory and attentional processes.

There are two ways in which a distractor can interfere with behavior. First, the control of the focus of attention may not be fine-grained enough to include the target and still exclude the distractors. Doing so depends on the precision of attention, the similarity of the features among target and distractors, and their spatial proximity. Second, some aspects of distractors may be processed automatically, without the need for attention. Nevertheless, this automatic processing may disrupt performance in one way or another. It may recruit attention away from processing of the target (as when a loud noise occurs in the unattended channel in selective listening) or it may contaminate the response-planning process (as in the Stroop effect). Variations on the filtering procedure are quite helpful in investigating the means whereby processing is disrupted.

A key type of evidence that ongoing processing has been disrupted is the orienting response (Sokolov, 1963), in which a shift of attention from one stimulus to another is accompanied by motor slowing and physiological signs such as heart-rate slowing, accompanied by privileged processing of the item that caused the orienting. Novel, abrupt stimuli tend to cause orienting, whereas repeated stimuli result in a habituation of orienting (see Cowan, 1995, for a review). Sokolov's theory stated that one builds a neural model of stimuli that are repeated, and there is evidence that this neural model can help a person to ignore the repeated stimuli and pay attention to something else (e.g., Cowan, 1995; Elliott & Cowan, 2001; Waters, McDonald, & Koresko, 1977). Thus, the switching of attention from one event to another is partly, but not entirely, under one's own voluntary control. If a very loud noise is unexpectedly heard, for example, attention will invariably shift from the ongoing activities to the noise. One can be internally conflicted, as when a student needs to pay attention to a boring lecture for the sake of a grade but finds it nearly impossible to do so, and attention often is a struggle between voluntary and involuntary processes. Of particular interest for clinical research, stimuli of special significance to the participant can cause continued orienting even when repeated (see Cowan, 1995).

Multiple-Task Paradigms

The final category of attentional paradigm described by Luck and Vecera (2002) was the dual-task paradigm, but we can talk more generally about multiple-task paradigms. When two or more tasks are to be accomplished at the same time, one can ask whether they all draw upon the same attentional processes or *resources*. To the extent that they do, improved performance in one task can be accomplished only at the expense of poorer performance in the other task or tasks. My own specialty in this regard is getting lost if driving to a rarely visited location while holding a discussion with a passenger, as several of my colleagues can attest (e.g., Monica Fabiani, personal communication, March 10, 2000). Broadbent (1957) described a prototypical dual-task procedure. Two different lists of three items were spoken simultaneously, with one list presented to the left ear and the other to the right ear, and then all six items were to be recalled. Recall was superior when the required order of recall was first one ear, and then the other, as opposed to an order in which the first, second, and third items in each ear were recalled in temporal order by alternating between left- and right-ear stimuli. The presumed reason is that sensory memory was used to hold the items that had not yet been recalled, and it was easiest to access one ear's representation at a time, without switching attention back and forth.

A particular dual task may have been groundbreaking in fostering the realization that human perception is fallible. As Boring (1957) has described, in Greenwich in 1796, an astronomer (Maskelyne) fired his assistant (Kinnebrook) because the two of them consistently produced results that were very different from one another. The task was one in which a metronome had to be used to time the movement of a star across hairlines in the telescope. Another forward-thinking astronomer, Bessel, from Königsberg, realized that the assistant may not have been at fault. Visiting various astronomers across Europe, he found that each one gave a different measurement but with an impressive amount of stability in judgment within an individual. Boring attributed the individual differences to the *law of prior entry*, which states that when two sensory impressions arrive concurrently, it is possible to allocate attention more to one or the other and that the attended sensory impression will seem to have arrived sooner than the initially unattended one. Differences in the style of attending to the telescope versus the metronome could have caused the astronomers to differ in their judgments.

Attention may have to be divided across time, across space, or both. It must be divided across time in the currently popular attentional blink procedure (Raymond, Shapiro, & Arnell, 1992; based on earlier, similar findings not using that nomenclature, by Broadbent & Broadbent, 1987; Reeves & Sperling, 1986; Weichselgartner & Sperling, 1987). In the attentional blink procedure, a stream of items is presented rapidly on the computer screen (e.g., 10 items per second), with one item replacing the next at the same spatial location. The participant must search for two targets, sometimes on different bases, and make two responses after the stream of characters ends. For example, the participant might be asked to search for the occurrence of two digits among letters, and then report the digits. The finding is that there is a period after the first target during which recognition of the second target is poor. If there are no items between the two targets, this *attentional blink* is not found. However, there is a period in which an occurrence of the second target will tend not to be noticed, which produces a scallop in the performance function for that target. Performance may be 50% down from its peak when there are 100 to 200 ms between targets and it returns to normal when there are about 500 to 600 ms between them.

One might worry that an attentional blink could occur because of an actual eye blink after the first target, but the data seem to rule out such an interpretation. For example, event-related brain potential recordings show that a second target that is not reported because of the attentional blink still elicits electrical components corresponding to sensory processing (P1 and N1 components) and semantic analysis (N400) but that a component corresponding to the updating of working memory (P3) is missing (Vogel, Luck, & Shapiro, 1998).

Another popular procedure in which attention is divided across time is the psychological refractory period procedure (e.g., see Ruthruff & Pashler, 2001). In that procedure, two stimuli requiring responses again are placed near one another with a variable interstimulus interval. However, here the dependent measure is not the accuracy of stimulus perception but, rather, the reaction times to the two stimuli. Again it is supposed that the allocation of attention to the first stimulus uses up some resource that may not yet be available again when the second stimulus arrives. The benefit of procedures in which attention is divided across time is that one can gain powerful indices of the time course of various types of processing that cannot be carried out for two stimuli at once.

In other procedures, attention must be divided over space instead of time. For example, in a dichotic listening experiment, such as the one by Broadbent (1957) described above, different left- and right-ear stimuli are