Chapter 2

Selective Attention

Contents

2.1 Definition and Domain

Selective attention is the use of one source of information rather than another. To refine this general definition we will distinguish terms that refer to phenomena and terms that refer to theoretical concepts. Consider phenomena. An *effect of selective attention* refers to the consequences of manipulating the relevance of one source of information rather than another. For an example from vision, we might instruct you to read either the first word of the chapter title or the second word of the chapter title. You can easily follow this instruction and read either "selective" or "attention".

Now consider the corresponding theoretical terms. A *mechanism of selective attention* refers to the internal process by which one source of information is used rather than another. To distinguish the theoretical concept, we will use the term *selection* to refer to the mechanism of selective attention. Consider again the reading example in which you must respond to a cued word and not another. One account is that while both words are processed early in the visual system, at some point the representation of one word is selected for further processing and determines the response.

How is the representation of one word selected over another? One part of the answer to that question is that we can point our eyes at one word of the text, rather than another. But that is not the whole answer. Even when multiple words are within our line of sight and clearly visible, we can read a particular word while avoiding others. Another question is what consequences are there in visual processing, for selecting (or not selecting) one word over another? In the extreme, the selected word is recognized (i.e., it is read), while the unselected word is not. But what exactly changed in visual processing to lead to that final outcome and how exactly was that word selected?

We begin our discussion of selective attention by considering the simple case of attending to spatial locations. In addition, this initial foray is limited to a consideration of relatively simple visual tasks such as the detection and discrimination of simple stimuli. A much wider range of tasks and stimuli, as well as selection on the basis of visual dimensions other than space (e.g., color and time) will be considered as the book progresses. The goal of this early chapter is to consider a small number of phenomena that reveal the effects of selective attention. We will use the simplicity of task and stimulus to advantage when asking basic questions about how attention works and how it changes visual processing. Complexity will be added in each following chapter.

2.1.1 The Role of Space in Vision

Before asking about the role of space in selective attention, consider the role of space within visual processing more generally. Vision is fundamentally a spatial modality as contrasted, for example, with audition which is fundamentally a temporal modality. Vision begins when an image is formed on the retina from light that is reflected from surfaces in the world. That image reflects spatial relationships among stimuli within the world. This *spatiotopic representation* is maintained as information is transformed from initial retinal responses to organized representations of objects within the scene in cortical processing areas of the brain. As these spatial relationships are maintained across transformations of the retinal image, "channels" of information flow that are defined on the basis of space are established. These spatial channels constitute a medium through which spatial attention can function.

Another aspect in which space is fundamental to vision is that the information at the retina at any given moment necessarily derives from only part of the surrounding world. This is because our eyes can point in only one direction at a time. Humans (and many other animals) move their eyes in quick point-to-point movements called *saccades* at a rate of approximately 4 times per second. It is as if the visual world is being sampled as a sequence of spatial windows defined by changing eye fixations.

Where one fixates has consequences. First, it determines which part of the world is visually represented at all; the parts of the world behind our heads is not represented. In addition, visual information processing is not homogenous across the retina. Information from locations near central fixation are represented in greater detail than information from peripheral locations. Consider the demonstration of Figure 2.1. This figure contains two Landolt-Cs which are C-like figures that can have a gap pointed in any direction (e.g. left, right, up or down). To control your eye position, fixate the central cross. With your eyes at this position, the Landolt-C on the left has an eccentricity that is about 1/4 of the Landolt-C on the right. For typical viewing conditions, you can easily identify

Figure 2.1: An illustration of the effects of eccentricity. Please fixate the central cross. Keeping your eyes on the cross, you can clearly see the Landolt-C on the left but not the more eccentric Landolt-C on the right.

the nearby figure on the left but not the more eccentric figure on the right. This is an effect of eccentricity.

To quantify eccentricity, one must consider the nature of the visual image. The image at the eye is two-dimensional. It does not have a direct representation of the distance from the eye to the relevant objects. As a result, the location in the 2-dimensional image is measured using degrees of visual angle within the image rather than physical locations as in the the 3-dimensional world. The geometry and calculation of visual angle is illustrated in Figure 2.2. Given the page is viewed at an arm's length of 60 cm, the two eccentricities in the demonstration have visual angles of about 1.5° and 6°. Many of the experiments in the opening chapters of this book are conducted in peripheral vision at eccentricities of about 6°.

Eccentricity effects are due to several factors including a greater concentration of photoreceptors at the central part of the retina (i.e., the *fovea*) and greater connectivity of foveal photoreceptors to higher-order cells within the visual system. Eccentricity effects constitute an important effect of space on visual processing. They are not, however, attention effects. They reflect more-or-less fixed properties of the visual system that cannot be altered based on one's goals except insofar as one controls where one's eyes are pointed. Eccentricity effects must, therefore, be considered and controlled for when asking about effects of spatial selective attention.

In order to simplify the study of spatial selective attention in light of known eccentricity effects, researchers often limit their inquiry to processing of visual information that is available within a single fixation. One way of approximating this to to present stimuli briefly (50 to 250 ms duration). This duration is too short for observers to move their eyes. Thus what is available during that time is about what is available within a single eye fixation.

In summary, in this chapter we focus on selective attention to spatial locations. Furthermore, we consider this question for the domain of single eye fixations. In practice, this is archived by instructing observers where to fixate and presenting brief displays.

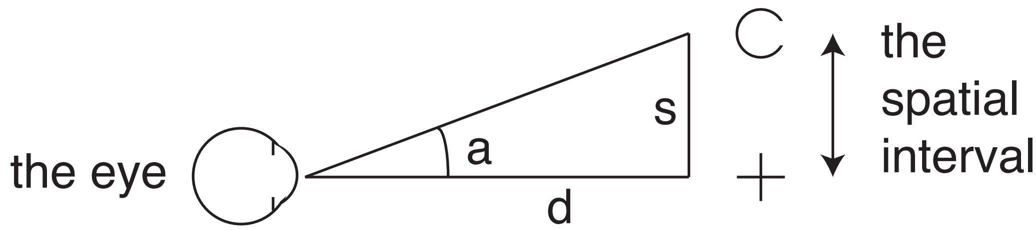

Figure 2.2: An illustration of how visual angle depends on size and distance. Suppose you want to calculate the visual angle between the fixation cross and the Landolt C of the demonstration in the preceding figure. This figure illustrates the triangle made by the two objects and the eye. To calculate the visual angle a you needs to know the spatial interval s and its distance d from the eye. The relation between these variables is given by the trigonometric relation: $tan(a) = s/d$.

2.1.2 Simple Stimuli and Tasks

In addition to limiting our initial scope of inquiry to the processing of information that is available within a single fixation, we are also going to limit it to simple stimuli and simple tasks. Again, complexity will be built up as the book progresses. But for our first steps simplicity offers the advantage of allowing us to ask whether basic visual functions are altered by selective attention.

What we consider simple stimuli are easy to describe. Typically they are small spots of light displayed at different location relative to fixation. Some experiments will use other favorite stimuli in visual science such as sinusoidal gratings (see Figure 2.3).

Regarding tasks, perhaps the simplest visual task is *detection*. In a detection task, observers are asked to view a display and report whether or not it includes some stimulus (e.g., a point of light, a disk of light, or some variation in light intensity of the entire field). Discrimination is another commonly used simple task. Here the task is for the observer to view a display and report which of multiple stimuli (e.g., a black disk or a white disk) is present.

By carefully manipulating the stimuli that are used in these simple tasks, inferences can be drawn about basic sensitivities of the visual system. One can ask, for example, how intense a light flash has to be for an observer to reliably detect in the periphery, compared to it at fixation. Or, as illustrated in Figure 2.1, how large must be the gap in a Landolt-C for an observer to discriminate its orientation? These sensitivities reflect basic visual processes without regard to selective attention. By measuring sensitivity using simple tasks and stimuli, one can determine how sensitivity depends on selective attention. This is our reason for starting simple.

Even using a detection task to measure sensitivity can be challenging. The simplest approach is to ask: do you see it "yes" or "no"? We will sometimes use this *yes-no method*. Unfortunately, it is prone to bias and distinguishing effects of bias from effects of sensitivity isn't always easy. Observers dislike saying that they see something that isn't present. For that reason, when observers

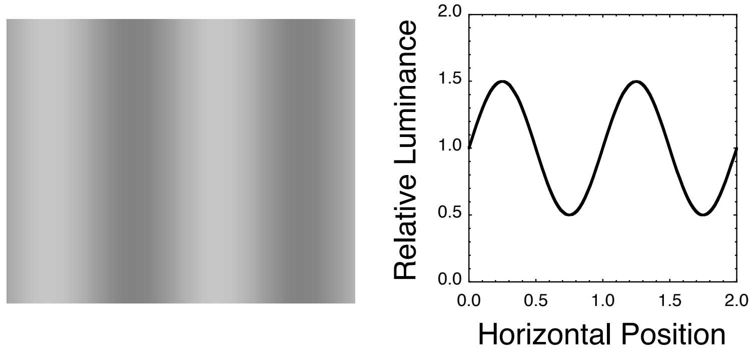

Figure 2.3: An illustration of a sinusoidal grating. On the left is an image of a grating with light and dark vertical bars. On the right is the corresponding plot of relative luminance versus horizontal position in which luminance varies sinusoidally with position. Gratings have parameters for both spatial and intensive properties. The primary spatial parameter is the period of the grating which is the distance required for one complete pattern of light and dark. In this example, the period is 1 unit of horizontal position. Equivalent to the period is the spatial frequency which is the number of periods per unit space. The intensive parameters are the mean luminance and the contrast. In this example, the mean is 1 unit of relative luminance and the contrast is 0.5 (or 50%). For periodic gratings, the contrast is defined as the difference between the highest and lowest luminance relative to the sum of the highest and lowest luminance. In summary, this is an example of a sinusoidal grating with 2 periods and a contrast of 50%.

are uncertain they tend to say "no" rather than "yes". This bias has been measured and much studied (Green & Swets, 1966) for its own sake. But for the purpose of measuring sensitivity, we favor *forced-choice methods* that minimize bias. For example, a stimulus can be presented to the left or right side of fixation and the observer can response on which side was the stimulus. Now there is always a hard-to-see stimulus present and observers must do their best to determine on which side it is present. Compared to yes-no, forced-choice methods have little response bias. Other forced-choice methods use time intervals or other stimulus features to provide a choice for the observer.

In summary, in this chapter we focus on simple visual stimuli such as spots of light and the simplest tasks such as detection and discrimination. Furthermore, we will use forced-choice methods whenever possible to avoid the need to distinguish bias and sensitivity.

2.2 Spatial Cueing

The very first question that we consider is whether selective attention to space even exists. That is, beyond orienting your head or eyes toward a given location, can you selectively process visual information from one location over another? The basic idea goes back to Helmhotz (1894/1968). He briefly illuminated a scene to prevent eye movements and described selectively reporting different parts of the scene. This idea was formalized by the partial report studies of Sperling (1960) which we will consider in Chapter xxx on memory.

Figure 2.4: An illustration of the stimuli used in Davis et al. (1983). The three possible displays differ in the spatial position of a grating. This creates spatial uncertainty for the observer.

2.2.1 The spatial cueing paradigm

The idea in spatial cueing is to manipulate observers' knowledge about the location of stimuli and ask whether this influences how those stimuli are processed. Suppose an observer must detect a simple stimulus that could be presented at one of several possible locations and that the stimulus is low contrast to make it hard to see. The idea of spatial cueing is to manipulate what the observer knows about the possible location of the stimulus. In particular, a cue precedes the display to indicate the location of the stimulus. If such knowledge improves the detection of the stimulus relative to when the same stimulus had to be detected with no knowledge about location, then the knowledge about location somehow changed the way information was processed across locations. This is an effect of selective attention.

2.2.2 A sample spatial cueing experiment using a contrast detection task

Consider a spatial cueing experiment reported by Davis, Kramer and Graham (1983). They used a contrast detection task and stimuli that were presented at more than one possible location. The stimuli were gratings, which are patches of repeating patterns of light and dark across space. A particularly common type of grating is a sinusoidal grating in which the light level varies sinusoidally across space. Sinusoidal gratings are widely used in studies of optical systems including human vision. Figure 2.3 shows an example of such a grating. As described in the figure caption, gratings have several parameters that can be usefully manipulated to study the responsiveness of the visual system.

Davis and colleagues used low-contrast gratings that were just visible under their conditions. They then asked whether sensitivity to the gratings was determined entirely by the stimulus that was presented, or whether it is influenced by the observer's knowledge of the spatial position in which the stimulus might appear. That is, they asked whether contrast sensitivity is subject to influence from spatial selective attention.

To manipulate observers' knowledge about the location of the stimulus, they manipulated the *spatial uncertainty* of the grating through auditory cues that were presented prior to the visual displays. In one condition the cue provided no information about where the grating would be presented. Following this type of cue observers knew that the grating was equally likely to appear

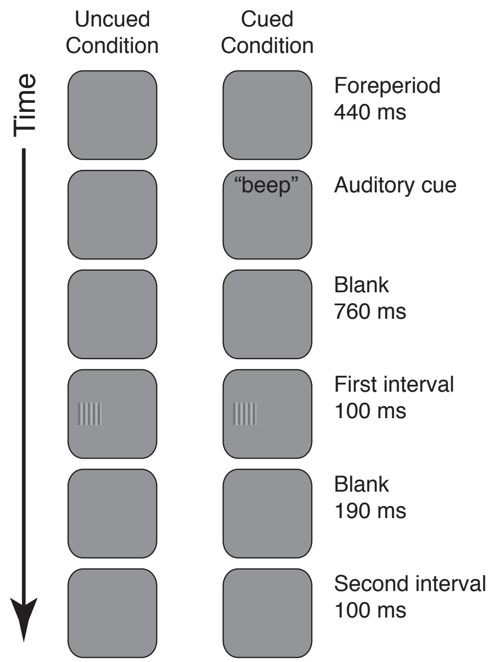

Figure 2.5: An illustration of the procedure used in Davis et al. (1983). The displays in a trial are shown in sequence going down the illustration. On the left is the sequence for an uncued trial and on the right is the sequence for a cued trial. They differ only in the presence of an auditory cue specifying the spatial position of the stimulus. The cue changes the number of relevant locations from three to one.

in any of three different locations (left, center, right; see Figure 2.4). In another condition, the cue indicated one of the three locations as the location where the stimulus would be presented (one, two, or three tone bursts indicated left, center or right location). Following the cue, the observer knew that the stimulus would appear in the indicated location. In short, the cue reduced spatial uncertainty.

It is important that nothing differed in terms of the stimuli themselves across the different cueing conditions. The only thing that differed was the knowledge provided by the cue regarding where the stimulus would be presented. In this way, the manipulation was designed to assess effects that are specific to selective attention. Any differences across cueing conditions must be attributed to the differences in knowledge about spatial location because there were no stimulus differences.

To measure detection, Davis and colleagues used what is known as a *two-interval forced choice* task. Two temporal intervals were defined within each trial, the beginnings of which were indicated by tones. A grating was presented on every trial during one of the two intervals, and observers reported which of the two intervals contained the grating. This forced-choice task provided a measure of contrast sensitivity that minimized the concerns that the cue might affect response bias rather than sensitivity.

Details of the procedure are illustrated in Figure 2.5. It shows the sequence of displays for one trial of the uncued and cued conditions. Trials were initiated by the observer using a keypress.

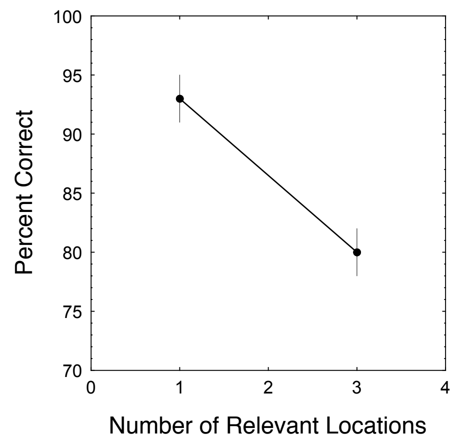

Figure 2.6: The results of Davis et al. (1983). Percent correct is plotted as a function of the number of relevant stimuli. The error bars are the standard error of the mean. There is a reliable effect of the cue.

The initial part of the trial either included an auditory cue (cued) or did not (uncued). This was followed by two, 100-ms stimulus intervals separated by 190 ms. A single grating appeared at one of three locations in either the first or second interval. These stimulus intervals were marked by tones to make it clear when the grating could appear. The observer's task was to indicate whether the grating was in the first or second interval using a keypress. Observers were instructed to fixate the middle of the display at the beginning of the trial and maintain fixation throughout the trial.

The results are shown in Figure 2.6. The mean percent correct for 2 observers is plotted as a function of the number of relevant stimuli. A cue that indicated a single location improved performance by more than 10% relative to one that provided no information about the location of the stimulus. The aggregate effect of the cue over observers and conditions (including some conditions not discussed here) was approximately 13%. In addition, when the effect was broken down for each of the three positions separately, the advantage occurred for each one. Clearly the cue affected performance.

The effect of the spatial cue on performance in this experiment provides an answer to our initial question. Yes, selective attention to space exists. This effect is attentional because processing of the gratings was altered by knowledge of the location in which the stimulus was presented. Performance was determined not just by the stimulus, but by knowledge of the location of the stimulus.

2.3 Partially-Valid Cueing

A useful variation on spatial cueing is known as *partially-valid cueing*. It has long been known that varying the probability of a stimulus affects the accuracy and response time to detect that stimulus (Hyman, 1935). Indeed, early treatments of stimulus probability related stimulus probability to stimulus uncertainty using information theory. It has also been clear that stimulus probability is

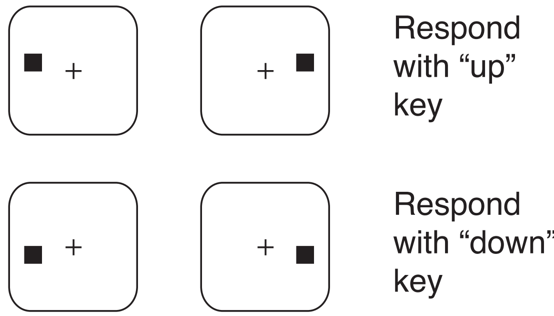

Figure 2.7: An illustration of the stimuli used in Experiment 3 of Posner et al. (1980). The stimuli were small squares in four possible locations: 7° to the left or right of fixation and just above or below the horizontal midline. The observers task was to judge if the squares were above or below the horizontal midline.

usually confounded with the response probability and its effect is probably mediated by decision and response processes as well as perceptual processes (LaBerge, Legrand & Hobbie, 1969). What is needed is a way to separate the perceptual effects from the other effects. The first step is to create an attentional manipulation aimed at the stimulus that does not affect the overall stimulus and response probabilities.

In partially-valid cueing, the cue indicates where a target stimulus is *most likely* to appear, but it is less than 100% valid. For example, a cue may indicate that the probability of a target appearing in the cued location is .8. That means that 80% of the time the target will appear in that location, but the remaining 20% of the time it will appear in some uncued location. Trials in which the target appears in the cued location are referred to as *valid* trials, whereas trials in which it appears in an uncued location are referred to as *invalid* trials. Now one can ask whether providing such probabilistic information about the relevant spatial locations has an effect on behavior.

There are two key innovations in this paradigm relative to simply manipulating the stimulus probability. First, the conditional probability of the stimulus given the cue is manipulated while the stimulus probability is held constant. Second, the task is chosen to be independent of the cued dimension. For example, one might be cued to location but judge color. The location cue says nothing about the correct response to a color task. Broadbent called this distinction the difference between stimulus set and response set (1970).

An early and influential example of the partially-valid cueing paradigm is due to Posner, Snyder and Davidson (1980). Here we will consider in detail their Experiment 3. The possible stimuli are shown schematically in Figure 2.7. The stimuli were 0.5° small squares presented 7° to the left or right of fixation. In addition, these squares were presented just above or below the horizontal midline defined by a central fixation cross. The observers task was to judge the relative vertical position of the stimuli.

The procedure is illustrated in Figure 2.8. Each trial began with a central arrow cue indicating which side of the display was more likely to have the stimulus. The cue correctly indicated the

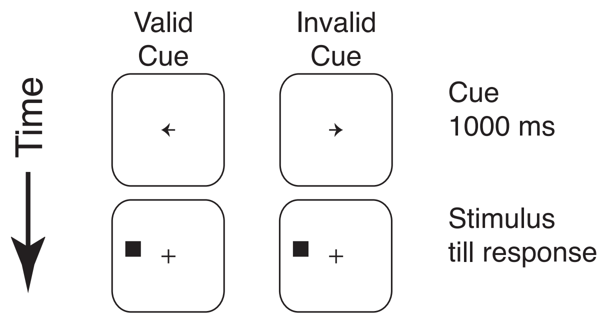

Figure 2.8: An illustration of the procedure used in Experiment 3 of Posner et al. (1980). The displays in a trial are shown in sequence. Trials begin with a cue and then the stimulus appears until response. On the left is the sequence for a trial with a valid cue and on the right is the sequence for an invalid cue. They differ only in the cue.

horizontal position of the stimuli on 80% of the trials. A second later the stimulus was presented and the observers were to respond whether the stimulus was above or below the horizontal midline. The observers were to respond as quickly as possible without making undue errors. Thus this is a choice response time experiment. It is also a spatial forced-choice experiment. They task was to report the location of the stimulus, not whether there was a stimulus present. Furthermore, the judged vertical location was independent of the cued horizontal location. This last point is crucial to minimize the effects of the cue on response bias and preparation. The cue tells one nothing about what response to prepare.

The results of this experiment are shown in Figure 2.9. On the left is mean correct response time as a function of cue validity and on the right is mean percent errors as a function of cue validity. There was a reliable effect of cue validity on response time of about 30 ms. There was also a small but unreliable effect on errors in the same direction of about 1%. Thus, giving a probabilistic location cue does improve performance.

Is this an effect of attention? As with the spatial cuing paradigm, the answer to that question is clearly yes. The stimuli being judged were identical for the two levels of cue validity. The only difference between a valid and invalid cue conditions is the direction of the central cue presented well before the stimulus. Thus, performance must be influenced by the state of the observer and not the stimulus alone.

Stepping back from the details of this experiment, consider the differences from the spatial cueing paradigm. The key difference is in the pair of conditions being compared. For partially-valid cueing, one compares trials in which the cue validly specified the stimulus location to the trials with otherwise identical conditions in which the cue did not specify the stimulus location. In short, the stimulus uncertainty is manipulated by the validity of the cue (e.g. .8 vs. .2). In contrast, for the spatial cueing paradigm, one compares a 100% validly cued condition with a uncued condition. Here the spatial uncertainty is manipulated by the number of relevant locations in the cued vs. uncued trials (e.g. 1 vs 3).

There are many other details of these examples that are not critical. One example used unspeeded accuracy while the other used a response time task. Both paradigms have been done both

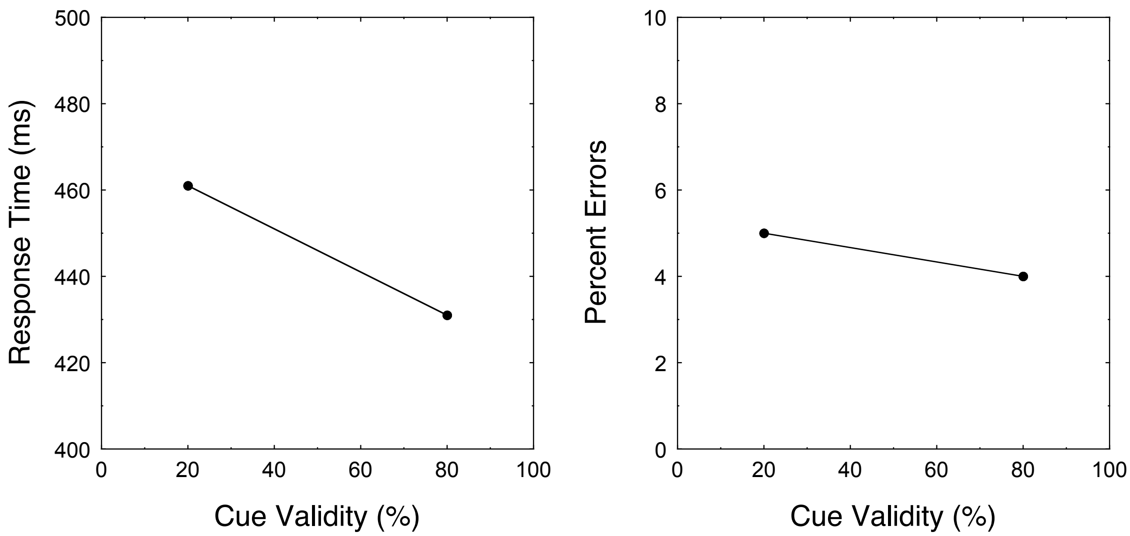

Figure 2.9: The results of Experiment 3 in Posner et al. (1980). On the left, the mean response time of 8 observers is plotted as a function of cue validity. On the right, the mean percent errors is plotted as a function of cue validity. Response time and errors decrease with increasing cue validity.

ways and yielded effects of selective attention. One example used temporal forced choice and the other used spatial forced choice. Again, both methods have been used and share the critical feature of the judged dimension is independent of the cued dimension.

2.4 Spatial Filtering

Consider next another variation in selective attention paradigms called *spatial filtering*. As in the *spatial cueing* paradigm the effect of selective attention is measured by manipulating the spatial uncertainty of stimuli using spatial cues and asking what effect, if any, it has on performance in a visual task. The information carried in the cue can be helpful if observers are able to use it. The spatial filtering paradigm is different in that observers *must* use the cue in order to do the task. That is, a location is cued, but in this case, the cue actually defines the task. It indicates the cued location as relevant and other locations as irrelevant. Stimuli are presented in both cued and uncued locations, and the stimuli are drawn from the same set. Therefore, the only thing that defines a stimulus as being the relevant one is whether or not it appears in the cued location. In summary, in filtering the cue is necessary to correctly perform the task while in spatial cueing the cue is merely helpful.

Consider a filtering task in audition from the early literature on attention. Broadbent (1958), and others, used dichotic listening to study selective attention. One auditory stream (e.g., a voice reading some text) was played into one ear, while a second auditory stream was played into the other ear. One ear was cued as the relevant ear, and observers were asked to repeat back the message in that ear while ignoring the message in the other ear. It was common to use a filter

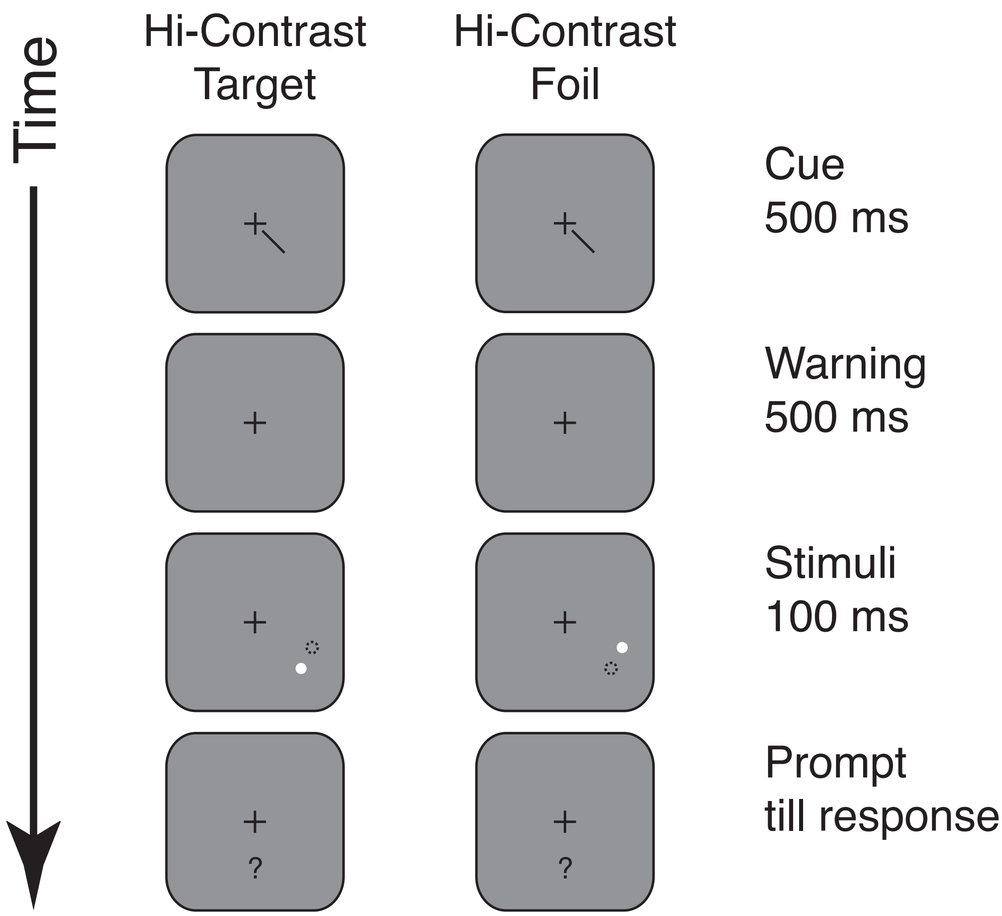

Figure 2.10: This figure shows the sequence of displays used in Yigit-Elliott, et al. (2011). Time goes from top to bottom. The cue is presented and then followed by a brief target display. The left column gives an example with a high-contrast target and a low-contrast foil. The right column gives an example of a trial with a high-contrast foil and a low-contrast target. The low-contrast stimuli are represented by a dashed black ring because low-contrast stimuli cannot be reliably reproduced. With eccentric viewing, the low-contrast stimuli were just barely visible. Thus, on most trials, observers reported seeing only a single stimulus.

metaphor in regard to this task, and in fact Broadbent's (1958) theory of attention is known as *Filter Theory*. The idea is that to do the task, one must filter out the information in the uncued stream and respond only to the information in the cued stream. Since this beginning, the term *filtering* has been adapted to refer to the paradigm and the term *selection* has been adapted to refer to the theoretical mechanism.

Now consider a visual analog of dichotic listening used by Neisser and Becklan (1975; see also Simons & Chabris, 1999). Two video sequences were presented with one superimposed spatially onto the other. One of the two videos was cued as being relevant to the task and the other irrelevant. For example, one video showed a ball game in which players pass the ball among themselves, while the other video showed a hand-clapping game in which participants interact in a series of hand-clapping gestures. For the ball game, the task was to count the number of times the ball is passed; for hand clapping, the task was to count the claps. Sometimes the ball-game video was cued as relevant and other times the hand-clapping video was cued as relevant. To the extent that observers can do this task at all, it implies the use of selective visual attention. Nothing differed in the stimulus between conditions in which one video was relevant or the other was relevant. Yet if observers did the task, then it follows that the stimulus was processed differently depending on which video was relevant. This is a compelling study, and we will return to some of the issues that it raises in later chapters when we discuss selection on the basis of dimensions other than space (e.g., objects or semantics). We leave it for the moment to return to a much simpler example of spatial filtering.

2.4.1 A sample spatial filtering experiment

Yigit-Elliott, Palmer and Moore (2011) reported a spatial filtering experiment that used contrast discrimination and brief displays. Figure 2.10 illustrates the basic procedure. On each trial, observers were cued to a peripheral location. A disk, which was either lighter than or darker than the background, was briefly presented at the cued location. This disk was the *target* because it was presented in the cued location. The task was to report the contrast polarity (i.e., lighter or darker than the background) of the target. In addition to the target, a second disk was presented in a nearby location. This disk was the *foil*. Observers were instructed to ignore the foil. The target and foil were drawn from the same stimulus set. Therefore the only thing that defined a given stimulus as a target or a foil was its location. In addition, the contrast polarity of the two disks varied independently across trials so that they both could be light, both dark, or one dark with other light. In other words, the polarity of the foil predicted nothing about the polarity of the target.

Both the target and foil also varied in the degree of contrast. Notice that *contrast* is different from *contrast polarity*. A "darker" target can be much darker than the background (high contrast) or only a little bit darker than the background (low contrast), and similarly for "lighter" targets. The full experiment had a wide range of contrasts. But for the moment consider only two contrast conditions: one in which the target had high contrast and the foil had low contrast, and one in which the foil had high contrast and the target had low contrast. Figure 2.11 illustrates some sample displays that are defined when these two contrast conditions are factorially combined with whether the contrast polarity of the target and foil are the same or different. The target is always

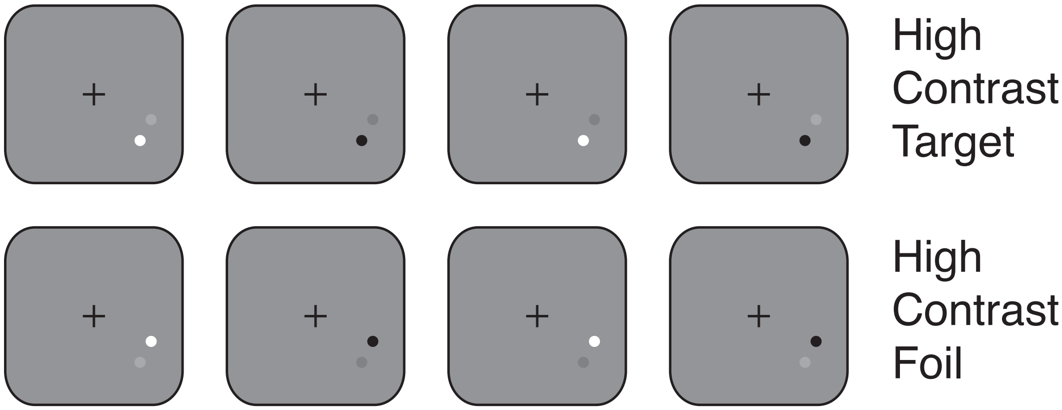

Figure 2.11: This figure shows the variety of stimuli used in Yigit-Elliott, et al. (2011). The top row shows the cases in which the target had a high contrast (100%) and the foil had a low contrast (5%). The bottom row shows the cases in which the foil had a high contrast (100%) and the target had a low contrast (5%). The two displays on the left side had targets and contrasts with the same polarity and the two displays on the right side had targets and foils with the opposite polarity. Depending of the reproduction, the contrast shown for the low-contrast condition is probably not representative. With eccentric viewing, these stimuli were just barely visible.

at the 45° location relative to fixation and in these examples, the foil appears in a location just above it along an isoeccentric arc. In other displays (not shown), the foil would appear below. The top row shows stimuli in which the target has high contrast; the bottom row shows stimuli in which the foil has high contrast. Within each row, the target and foil vary in polarity.

Each trial began with a cue that indicated the relevant location (see Figure 2.10). For a given session of the experiment, the target location was always the same. In the example that location is to the lower right of fixation at an 8° eccentricity. Because this was always the target location for that run, the cue was technically unnecessary, but it was presented nonetheless to reinforce and remind observers of the relevant location. One second following the cue, the stimulus display was presented for 100 ms. In this display, the target was always presented at the cued location. The foil was presented to one or the other side of the target along an isoeccentric arc, separated from the target by 0.6°, 1.2°, or 2.4°.

The heart of this experiment was to compare responses to a high-contrast foil relative to a high-contrast target. One way of measuring performance is to report the percent of trials on which the response corresponded to the value of the target. This is the familiar measure of percent correct. Another way of measuring performance, however, is to report the percent of trials on which the response corresponded to the value of the foil. Because the target value and foil value were independently varied, if observers are completely successful in responding only to the target, then this second measure—*percent "foil" response*—will be 50% (i.e., chance). If this value is greater than 50%, then it indicates that selection failed in some way and that the response was influenced by the foil instead of just the target.

Figure 2.12 shows the results averaged across 6 observers. In the left panel, we show performance in terms of percent correct as a function of separation for the two contrast conditions. Not surprisingly, observers were very good at reporting the polarity of a high-contrast target in the presence of a low-contrast foil, regardless of the separation between them. In contrast, for the low contrast target that accompanied the high-contrast foil, observers were near chance (50%).

In the right panel, performance is shown in terms of the percent responses with the same polarity as the foil. For short, this measure is labeled as the *percent "foil" responses*. These responses are also shown as a function of separation for the two contrast conditions. For the trials with high-contrast targets, the foil responses were near chance. This is also unsurprising because the foils had a low contrast. The critical condition was the one with high-contrast foils (and low-contrast targets). For these trials, the target was difficult to see. The percent of "foil" responses was high for the small separation (0.6°) and fell to near chance for the large separation (2.4°). This pattern indicates that observers were able to select the target from the foil at large separations, but for the small separation, they were not able to successfully select the target. Because target and foil could not be selectively processed, responses were be driven by the high-contrast stimulus, which in this condition was the foil.

2.4.2 Is this an attention effect?

Let us now consider the same two questions about the results of this spatial filtering experiment as we did about results from spatial cueing experiments. The first is the question of whether the effect observed here reflects selective attention. The fact that observers could do that task at all (at least at larger separations), suggests that they selectively processed information on the basis of space because the target was defined by its location. But could the effect involve a subtle stimulus-driven effect? Did the target differ from the foil in any way other than that the target location having been cued and the foil location not? In the experiment of Yigit-Elliott and colleagues, the target and foil were always presented at the same eccentricity. They differed in their exact location: the target was always at 45°, while the foil appeared above or below this location. An ideal experiment would have alternated the locations of the target and foil. In fact, this better match between target and foil was used by Palmer and Moore (2009, Experiment 2) and yielded similar results. Therefore, setting this detail aside, there is little doubt that the filtering effect is an effect of the task goals of the observer (attention), rather than some purely stimulus-driven effect.

How striking is this finding? On the one hand, it seems intuitively obvious that one would be able to do this task if the spatial separation is large enough. But here is the point. There are tasks in which selection by location fails even when the separations are large. We will take up this important comparison in our discussion of the Stroop effect in Chapter xxx.

A further point to consider is the effect of separation that was revealed in these experiments. When the separation is small, observers act as if they cannot distinguish the target and foil; when the separation is large, observers act as as if they can perfectly distinguish the target and foil. In other terms, selecting the relevant information goes from no success with a small separation to nearly perfect with a large separation. Thus, these results provide something akin to a spatial tuning function for the spatial extent of visual selective attention. For these 8° eccentric stimuli, selection is half way between chance and perfect at a little over 1°. Additional measures of the

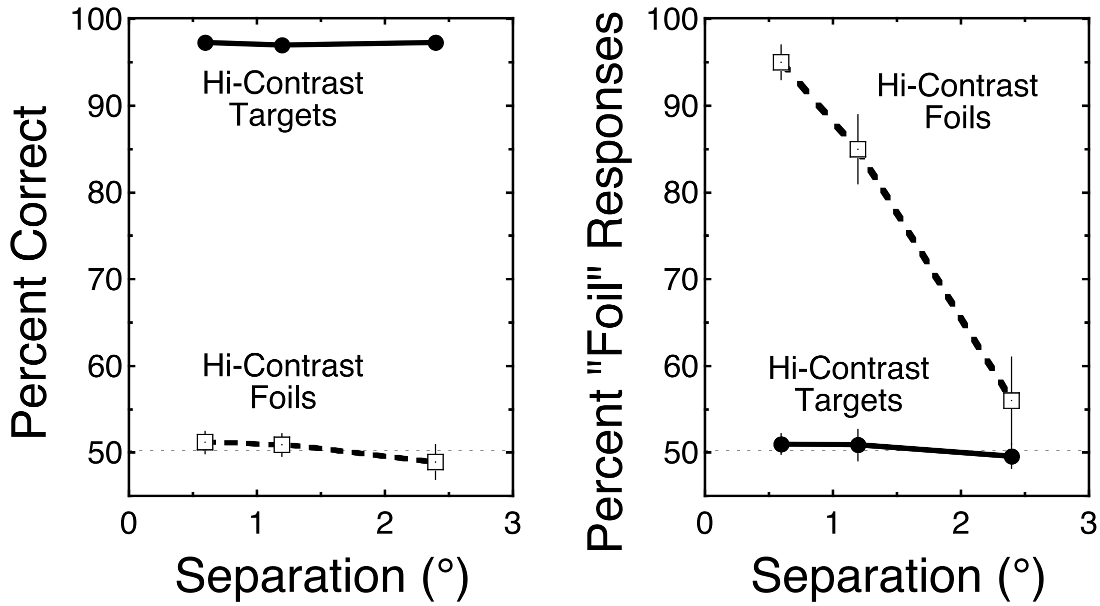

Figure 2.12: Results are shown for the experiment in Yigit-Elliott, et al. (2011). For the left panel, the percent correct is shown as a function of separation for two conditions: one with high-contrast targets (and low-contrast foils), and another with high-contrast foils (and low contrast targets). Observers are accurate with the high-contrast targets but not the low-contrast targets in the second condition. For the right panel, the percent “foil” response is shown as a function of separation for the same two conditions. These are the responses in which the response polarity was the same as the foil polarity. Observers were not affected by the foil in the high-contrast target condition (with low-contrast foils). Of key interest, are the results for the high-contrast foil condition. The foil did have an effect at small separations when it had a high contrast and the target had a low contrast.

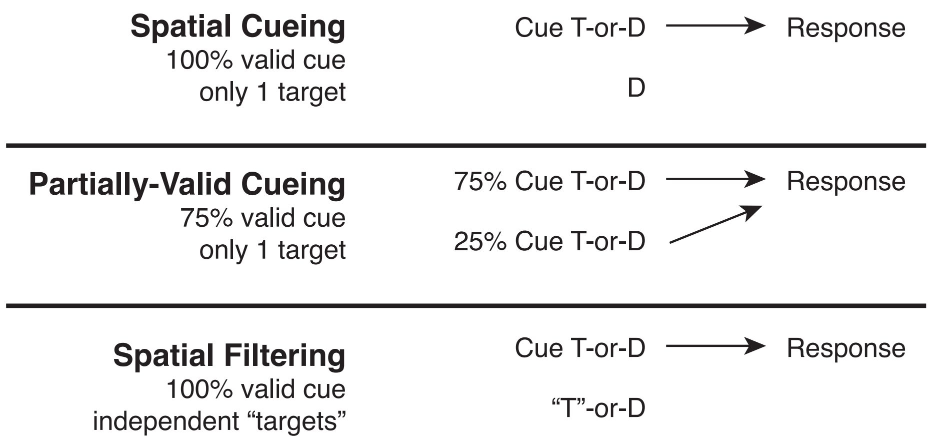

Figure 2.13: The three paradigms are compared by a simple task analysis. The top panel illustrates spatial cueing; the middle panel illustrates partially-valid cueing; and, the bottom panel illustrates spatial filtering. The possible stimuli for a two location experiment are shown with their mapping to response. For spatial cueing, a target or distractor is presented at the cued location while a distractor is presented at the uncued location. Only the cued location is mapped to a response. For partially-valid cueing, both targets and distractors are presented at both locations and both are mapped to responses. For spatial filtering, both "targets" and distractors are presented at both locations but only the cued location is mapped to response.

spatial extent of attention can be found in Palmer and Moore (2009). We will consider the spatial resolution of attention quantitatively in Chapter xxx.

2.5 Comparing the Three Paradigms

We have now introduced three selective attention paradigms. These paradigms obviously have many similarities. They also have some critical differences that make each of them useful in different ways. Consider the task analysis in Figure 2.13 to illustrate the important differences.

For purposes of this comparison, imagine the task is to detect targets T among distractors D and respond either "yes" or "no". In the examples given thus far, the distractors were actually a blank screen but there are similar visual search experiments that we will discuss in the next chapter with explicit distractors. The left column of the figure labels each panel and highlights key properties of each paradigm. For example, two paradigms use 100% valid cueing and two paradigms present only a single target while the other paradigm presents stimuli independently at the different locations.

The key part of the figure is the middle column illustrating the possible stimuli at the two locations and the arrows in the right column illustrating the mapping to response. For the cued condition of spatial filtering, targets are presented at only one location and only stimuli at that one location are mapped to a response. For partially-valid cueing, targets and distractors can occur at

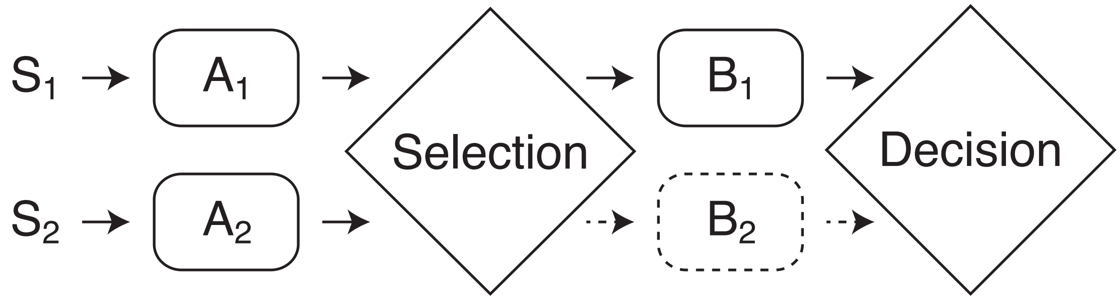

Figure 2.14: An illustration of selection between two stimuli using a diagram of the flow of stimulus information. The first stimulus is selected over the second stimulus. This results in reduced processing of the second stimulus after selection (B_2).

both locations and both are mapped to responses. For spatial filtering, both targets and distractors can occur at both locations, but only stimuli at the cued location are mapped to a response. Beware that the term "target" is not quite right for spatial filtering and we have used the term "foil" for targets at an task-irrelevant location.

Several comments highlight the differences. Spatial cueing is very simple but it lacks a natural comparison condition. It uses an uncued condition for comparison. This simple manipulation is used in many situations in the following chapter on divided attention.

Partially-valid cueing has a built-in comparison condition. On any trial, there is always a validly cued location and an invalidly cued location. But here, the constraint is that both locations are always relevant to some extent. One cannot consider 100% valid cues without loosing the needed invalid cue comparison condition.

Spatial filtering also has a built-in comparison condition. The target (T) and the foil ("T") can occur on any trial. Thus one can measure the effect on the response of either the relevant target or the irrelevant foil. In addition, spatial filtering uses 100% valid cues to provide the strongest attention instructions possible. Indeed, the limitation of this paradigm is that partially-valid cues cannot be used without making the "foils" into targets and thus breaking the paradigm.

As we proceed in this book, each of these paradigms will find multiple applications. And when we reach the chapter on quantitative models, we will explicitly consider their pros and cons in further detail. All three paradigms have a place.

2.6 The Locus of Selection

2.6.1 The concept of selection

We now turn to theoretical accounts of selective attention phenomena. The central concept is an internal process that is often called *selection*. Early authors have also used the term "filter" (e.g. Broadbent, 1958). For perceptual tasks, selection is the preferential processing of information from one stimulus over another. This idea is illustrated for two stimuli in Figure 2.14. Suppose the stimuli are distinct enough that they initially receive separate and independent processing as represented by A_1 and A_2. At some point in processing, a selective process affects the processing of the two stimuli. In this example, S_1 is given priority over S_2. This might be done by blocking all processing of S_2 (Broadbent, 1958) or by attenuating the processing of S_2 (Treisman, 1960). The figure shows this by the dotted lines and boxes for the B_2 process. Ultimately, information from these processes is used to make a decision and response. Later in this book we will consider the question of blocking versus attenuation in detail (Chapter xxx). At this point, we ask a different question: where in the processing sequence does selection occur? We next lay out two broad alternative hypotheses of the locus of selection.

2.6.2 Selective perception

One way that selection might affect visual processing is to change the development of a perceptual representation. We will refer to this hypothesis as *selective perception* because it focuses on how selection affects perception. The metaphor of a "spotlight" is often used to illustrate selective perception. The idea is that spatial selection is like a spotlight shining on a particular location, thereby illuminating stimuli at that location and not stimuli at other locations which remain in relative darkness. Thinking more formally, this hypothesis is sometimes referred to as the *signal enhancement* hypotheses. By this account, selection serves to enhance the signal so that it stands out better from the noise.

The defining feature of selective perception is that selection acts directly on the perceptual representations of individual stimuli. This is illustrated in the top panel of Figure 2.15. For the brief displays considered here, processing is feedforward and selection precedes perception. The figure also allows for some sensory processing occurring before selection. The point is that a part of perceptual processing is modified by selection. In this sense, selective perception is like the classic "early selection theories" in which selection occurs before the identification of individual stimuli. For the moment, we emphasize that that selection changes the way information at the selected location is processed.

2.6.3 Selective decision

Next consider what we call *selective decision.* According to this hypotheses, selection only affects decision processes and not perception. This is illustrated in the bottom panel of Figure 2.15. For brief displays, selection follows perception and affects only decision.

More specifically, a spatial cue allows information at uncued locations to be discounted, and thereby prevents irrelevant stimuli from interfering with the perceptual decision. Keep in mind that all of these experiments typically limit visibility to cause errors. Therefore, on those trials in which no single location was cued, visual noise in the information sampled from uncued locations might be erroneously attributed to there being a target stimulus. In contrast, when a location is cued, information sampled from uncued locations can be discounted thereby reducing the noise that

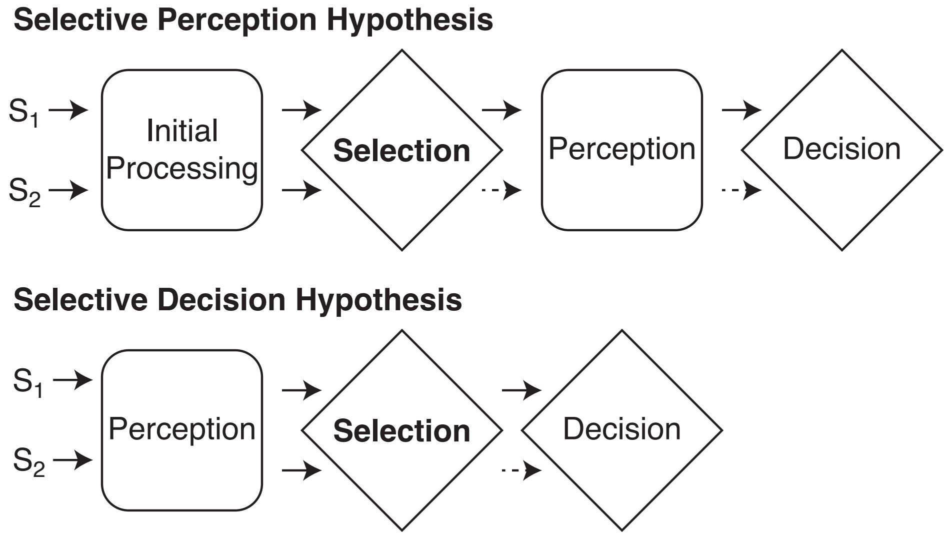

Figure 2.15: An illustration of two possible loci for selection in the processing of briefly presented stimuli. The top panel shows the selective perception hypothesis. In it, selection modifies perception as well as decision. The bottom panel shows the selective decision hypothesis. In it, selection modifies only decision.

would otherwise be fed into the decision process. These ideas are part of the classic "late selection theory" in which selection occurs after the identification of stimuli. For the moment, we emphasize that selection affects how information at multiple locations is combined to make a decision

The selective decision hypotheses has its roots in Tanner (1961) who argued that all visual tasks, including simple ones like detection, involve some element of decision making. The role of decision has already been introduced in separating sensitivity from response bias. While it is clear that spatial cueing affects sensitivity, it is possible that the effect of the cue is nevertheless mediated by decision. The possibility arises because the spatial cues change the decision problem as well as the perceptual problem. In the cued condition, only one location is relevant and need contribute to the decision process. By comparison, in the uncued condition, several locations are relevant and to be accurate all must contribute to decision. If the evidence for a stimulus is noisy at each location, then any decision will be less accurate as one considers more locations. This is because the noise at each locations gives a new chance for an error in decision. This hypothesis is often referred to as *noise reduction* because spatial cueing improves performance by reducing the contribution of noise from the irrelevant locations. In other words, selecting information from the relevant location avoids the noise at irrelevant locations.

2.6.4 Summary

In summary, we have identified two general hypotheses regarding how spatial selective attention affects visual processing. The *selective perception* or *signal-enhancement* hypothesis holds that the effect is on the representation of the stimulus at the cued location. The *selective decision* or *noise reduction* hypothesis holds that the effect occurs through the reduction of noise from uncued locations that would otherwise be included in decision processes about whether or not a grating was present. These hypotheses are not mutually exclusive. Selective attention could involve selection affecting both perception and decision. These ideas will be formalized in Chapter xxx on qualitative models.

2.7 Are Attention Effects always Perceptual?

The locus of attention effects is considered throughout this book. At this early point, we take up a particularly simple version of the locus question. Are the effects of selective attention always perceptual? In other words, does selective perception hold for all tasks and stimuli. This strong version of the selective perception hypothesis provides a good starting question. This view has been championed by Posner (1980) and continues to be influential.

2.7.1 Shiu and Pashler's critical experiment

Shiu and Pashler (1994) devised a test between the selective perception and selective decision hypotheses based on a version of the partially-valid cueing paradigm. We will first describe their version of partially-valid cueing and then consider their critical test of the hypotheses.

Figure 2.16 illustrates the procedure. Stimuli and cues were presented in one of four possible locations in the separate quadrants around fixation. Each trial began with a fixation cross and after a blank interval a cue (small square) was presented near one of the four possible stimulus locations. After a short blank interval, the stimulus was presented for 50 ms followed by a mask to reduce its visibility. In this experiment, the stimuli were the digits 4, 5, 6, and 7; and the observer's task was to identify the stimuli by pressing a corresponding key. The dependent measure was accuracy and the observers could take as much time as they wished. The main manipulation was whether on a particular trial the cue predicted the target location (valid cue) or did not (invalid cue) as shown in the two columns of the figure. The probability of the target appearing in the cued location was .75, whereas the probability of it appearing in any given uncued location was .083 (.25 divided by 3). In other words, the cue validity varied from 75% to 8%. Critically, the displays in these two conditions were identical except for the location of the cue.

Results from Shiu and Pashler's experiment are shown in Figure 2.17. To begin, consider only the multiple-mask condition. The mean percent correct of 12 observers is plotted as a function of the cue validity. Cue validity had an effect. Specifically, percent correct identification with valid cues was 11% higher than with invalid cues. This effect of cue validity must be attentional because the stimuli were identical in the two conditions.

Consider now whether this cueing effect reflects a change in perception or decision? According to the selective perception hypothesis, the effect of the cue is on processing information at the cued location. It influences the establishment of a perceptual representation of information at the cued location. So according to this hypothesis, valid cues lead to an enhanced perceptual representation of the target. Invalid cues do not. Instead, invalid cues should result in enhanced processing of information at the wrong location, which was empty initially and then contained a mask.

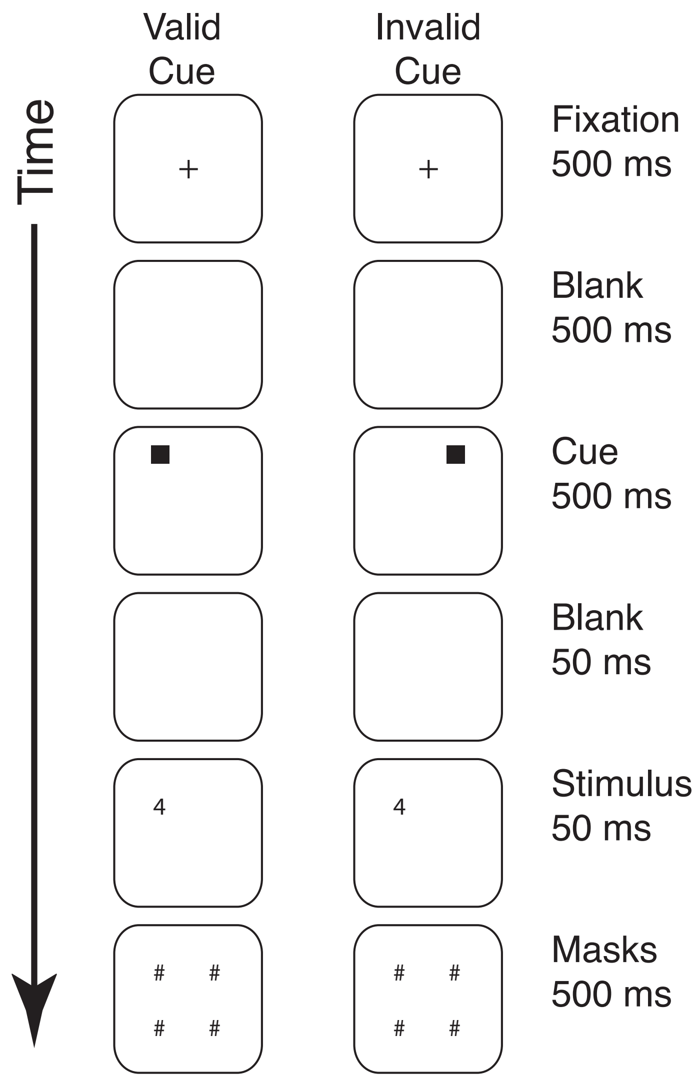

Figure 2.16: An illustration of the procedure used in Experiment 1 from Shiu and Pashler (1994). On the left is the sequence for a trial with a valid cued and on the right is the sequence for an invalid cued. They differ only in the location of the cue specifying the spatial position of the stimulus. The cue predicts the location on 75% of the trials and on the remaining trials is distributed over the other three possible positions. This illustration is for their multiple-mask condition. The single-mask condition differed in having a mask at only the location of the target.

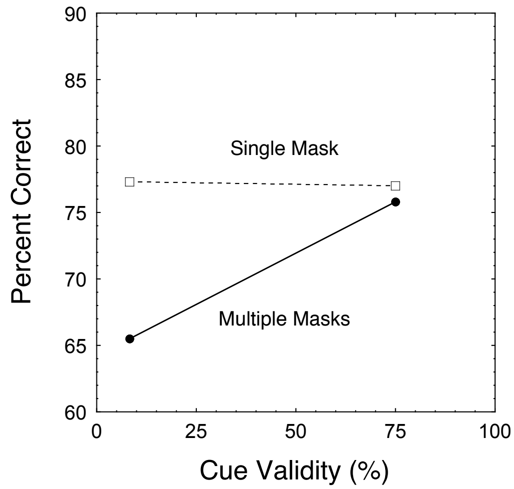

Figure 2.17: Results from Experiment 1 in Shiu and Pashler (1994). The mean percent correct is plotted as a function of the cue validity both single- and multiple-mask conditions. There is a reliable effect of the cue for the multiple-mask condition but not for the single-mask condition. This is evidence against selective perception.

Now consider the selective decision hypothesis. According to this hypothesis, the perceptual representations of information at all locations (cued and not cued) are the same regardless of the validity of the cue. What the cue does is allow the observer to discount information from uncued locations and limit input to decision processes to information from the cued location. On valid trials, such discounting is a benefit. Information from the target location is not discounted, while information from other locations is discounted. On invalid trials, however, such discounting does not help. Information from the target location is discounted while information from the wrong location is unaffected. Thus, the selective decision hypotheses predicts an effect of cueing.

A critical distinction between these two hypotheses concerns the role of information in uncued locations. For selective perception, the effect is determined by information at the cued location only. That information is enhanced by the cue. This improves performance when it is relevant information (valid trials) and impairs performance when it is irrelevant (invalid trials). The information at the uncued location has no role in this hypothesis. For selective decision, information at the uncued locations is integral to the effect. The cue reduces noise in the decision process by discounting information from uncued locations. For this hypothesis, the information at the cued location has no role in cueing effect.

Shiu and Pashler took advantage of this distinction to test between the two hypotheses. They manipulated the amount of noise at uncued locations. Selective decision predicts that more noise at uncued locations will yield larger spatial cueing effects because there is more advantage to be had for discounting the information from uncued locations. In contrast, selective perception predicts that the amount of noise at the uncued location will have little or no impact on the spatial-cueing effects. This is because the cue has its effect by acting on the processing of information at the cued location only.

Shiu and Pashler manipulated the amount of noise at uncued locations by manipulating the masks that followed the target displays. Specifically, in the multiple-mask condition that we have already described, four masks followed the four possible locations of the target display. In contrast, the single-mask condition had a single mask that was presented at the target location only. Under selective decision, the cueing effect should be smaller in the single-mask condition than in the multiple-mask condition. With high-contrast stimuli such as those used in this study, there will be little or no confusion between the high-contrast stimuli and an empty location. Without the masks, there is not much noise to be discounted. So cues will have little or no effect. Under the perceptual hypothesis, by contrast, the cueing effect should be unaffected by masks that follow the non-targets. If the spatial cue influences performance by enhancing processing at the cued location, then what appears at uncued location is irrelevant. So the effect of cues will be unchanged.

The results for this new condition are also shown in Figure 2.17. In the single-mask condition, cue validity has little if any effect ($\sim 1\%$). Consistent with selective decision, adding a mask to reduce visibility at non-target locations had a powerful effect on performance.

The results of this study are consistent with the selective decision hypothesis. But the two hypotheses are not mutually exclusive. If contributions from the decision processes are minimized, is there any remaining cueing effect? That is, controlling for an effect of decision, is there any evidence of an effect that could be due to perceptual enhancement of information at the cued location? The data from this particular experiment indicates that the answer to this question is no. In the single-mask condition, there was essentially no benefit to be had for reducing noise from uncued locations and there was no difference in performance between the valid and invalid conditions. Thus there is no evidence of selective perception under these conditions. The entirety of the spatial cueing effect in the Shiu and Pashler study can be accounted for by selective decision. This does not rule out the possibility that spatial cueing might play a role under other conditions, such as with more complex stimuli and tasks.

2.8 The Generality of Selective Attention Effects

We next turn to a brief review of the literature to put our example studies in context. At this point in our story, we are considering only simple stimuli such as luminance increments and simple tasks such as detection. Our goal in this review is to establish the generality of selective attention to space over these kinds of simple stimuli and tasks.

2.8.1 Spatial cueing

The Davis et al. example experiment is typical of one part of the spatial cueing literature. It demonstrated spatial cueing for two-interval, forced-choice detection and sinusoidal gratings. Other studies have demonstrated cueing effects for related stimuli and tasks: Cohen and Lashly (1974) using lights and two-interval, forced-choice detection; Foley and Schwartz (1998) using spatial two-alternative, forced-choice and gratings, and Grindley and Townsend (1968) using lights and a "yes-no-maybe" task. This last example is perhaps the earliest relevant example but used a task that is prone to intermix bias and sensitivity effects. In addition, there have been a variety of spatial cueing experiments using response time such as Eriksen and Hoffman (1972). Finally, in this chapter we emphasize a particular kind of cueing experiment with informative cues presented away from the relevant stimuli (sometimes called central cues or endogenous cues). Other kinds of cues are discussed in Chapter xxx on attentional control. In summary, these and related studies establish that knowledge about the location of stimuli can influence performance in visual tasks, beyond any effects of the stimulus itself. That is to say one can selectively attend to space and it has behavioral consequences.

2.8.2 Partially-valid cueing

The example partially-valid cueing experiment from Posner et al. (1980) used light detection and a spatial choice response time task. It is one of several early examples of response time studies using partially valid cues (Posner, 1980; Posner, Nissen & Ogden, 1978). Several of the other examples from that time use a simple response time task but we emphasize the choice response time task because it is easier to separate effects of bias and sensitivity. There have been many followup experiments using this paradigm to address a variety of questions (reviewed in Wright & Ward, 2008). This paradigm has also been used with accuracy measures. We have already described the study of Shiu and Pashler (1994) that used letters and a 4-alternative identification task. Similar examples include Bashinski and Bacharach (1980) using light detection and a yes-no task with confidence ratings, Van der Heijden, Schreuder and Wolters (1985) using letters and a identification task, and Shiu and Pashler (1995) using vernier stimuli and a two-choice discrimination task. All of these studies showed cueing effects but differed in the author's interpretation. There have also been partially-valid cueing experiments that added postcues to reduce the role of decision (Downing, 1988; Hawkins, et al, 1990). These studies also showed cueing effects and the authors argued for a perceptual rather than attentional interpretation. We defer considering such postcue methods until we consider postcues as part of explicit memory experiments in Chapter xxx. Finally, in this review, we have ignored the presence of neutral cues and focused on the difference between valid and invalid cues. This is because of the difficulty in establishing what is appropriate neutral cue (Jonides & Mack, 1984). In summary, for these partially-valid cues as well as for 100% valid cues, there is little doubt of the existence of cueing effects for space. The arguments are about the interpretation of these cueing effects.

2.8.3 Spatial filtering

Spatial filtering tasks designed for simple stimuli and detection tasks are a relatively recent development. The Yigit-Elliott et al. example used lights and a light-dark polarity discrimination. Similar studies by Palmer and Moore (2009) used lights and a spatial 2AFC task. They also have conducted unpublished experiments with lights and a 2IFC task. Early applications of filtering can be found in Eriksen and Hoffman (1973) using letters and a response time task. The best known such example is the *flanker paradigm* introduced by Eriksen and Eriksen (1974) that typically uses letters and a response time task. This paradigm also usually includes a many-to-one mapping of stimuli to responses that is aimed to address the locus of selection. There are many variations on this task including those that use accuracy (Whur & Musseler, 2005) and simpler stimuli such as color (Cohen & Shoup, 1997; Mordkoff, 1998). Almost all flanker experiments show some degree of successful selection particularly with large separations between the target and flanker. We defer a discussion of the interpretation of flanker effects until Chapter xxx on letters and words.

2.8.4 Is the locus of selection always perceptual?

In the last part of this chapter we began to address how to interpret the effects of selective attention. Is it due to selective perception, selective decision, or both? Studies such as Shiu and Pashler (1994) allow one to address the simpler question of whether these effects are entirely perceptual. Shiu and Pashler argue that they must be due to decision under their conditions. Several studies (e.g. Cheal & Gregory, 1997; Luck, Hillyard, Mouloua & Hawkins, 1996) have replicated the effect of single vs multiple masks on performance. But unlike Shiu and Pashler, they show some cueing effect in the single-mask condition. This invites the interpretation that there are contributions of both selective perception and selective decision. In more recent studies, some authors have accepted the presence of selective decision and tested for evidence of selective perception under conditions with minimal spatial uncertainty. For example, Smith (2000) suggests that there is some contribution of selective perception for experiments with masking but not for experiments without masking. We will pursue such locus questions throughout this book. As already suggested, we will argue that the locus depends on the stimuli and the task. For now, we summarize the evidence as ruling out the idea that spatial cueing effects are always perceptual. Selective decision makes a contribution for the paradigms introduced in this chapter.

2.9 Chapter Summary

2.9.1 Phenomena and paradigms

In this chapter, we have introduced the phenomena of selective attention to space. This phenomena is revealed by manipulating the observer's intentions about the stimuli relevant to the current task. These phenomena are captured in our examples by the use of cues to give observers information about the stimuli relevant to a given task. The chapter focussed on three paradigms that manipulate selective attention in different ways: spatial cueing, partially-valid cueing, and spatial filtering. These three paradigms will be used throughout this book.

In spatial cueing, the observer has to monitor multiple locations on some trials and on other trials can monitor a single cued location. The cue effect is measured by comparing a condition in which the cue minimizes spatial uncertainty to a condition with spatial uncertainty. This paradigm shows off the advantage of a to-be-attended location relative to a neutral control.

In partially-valid cueing, the observer is given a probabilistic cue that makes one location more likely to have the target than another. The cue effect is measured by comparing a condition in which the target occurs at a likely location to when it occurs in an unlikely location. This shows off the advantage of different degrees of prior information about location on the detection of a target.

In spatial filtering, the observer is cued that one location is relevant and another location is irrelevant. The same set of stimuli appears at both locations so now one has to ignore foils at the irrelevant location. The filtering effect is measured by comparing responses to the foil at the to-be-ignored location, to responses of an otherwise identical target at the to-be-attended location. This paradigm shows off the fate of stimuli at a to-be-ignored location.

2.9.2 Two theoretical questions

In our analysis of these paradigms, we raised two questions that will reappear in various forms throughout this book. The first can be termed the *attention question*: Is a given effect due to attention modulating stimulus processing, due to purely stimulus-driven processes, or due to a combination of both kinds of processes? Each of the three paradigm isolates attention-dependent processing from purely stimulus-dependent processing.

The second question can be called the *locus question*. Is a given attentional effect due to perceptual processing or some other process such as decision? For these effects of selective attention, these hypotheses were called selective perception and selective decision. To begin addressing this question we provided a counterexample to selective attention effects always being mediated by perceptual processes. As the book progresses, we will consider an variety of processes (decision, memory, etc) that might mediate attention effects under various conditions. To preview the answers, much depends on the particular tasks and stimuli.

2.9.3 One theoretical mechanism

Throughout this book, we will take pains for distingusih terms for emperical phenomena from terms for theoretical concepts. To that end, we introduced the theoretical concept of *selection* as an account for the phenomena of selective attention. For perception, selection is an internal process that favors one source of stimulus information over another. Alternative hypotheses about the nature of selection will be considered as this book progresses.