

Involuntary Covert Orienting Is Contingent on Attentional Control Settings

Charles L. Folk
Villanova University

Roger W. Remington and James C. Johnston
NASA—Ames Research Center

Four experiments tested a new hypothesis that involuntary attention shifts are contingent on the relationship between the properties of the eliciting event and the properties required for task performance. In a variant of the spatial cuing paradigm, the relation between cue property and the property useful in locating the target was systematically manipulated. In Experiment 1, invalid abrupt-onset precues produced costs for targets characterized by an abrupt onset but not for targets characterized by a discontinuity in color. In Experiment 2, invalid color precues produced greater costs for color targets than for abrupt-onset targets. Experiment 3 provided converging evidence for this pattern. Experiment 4 investigated the boundary conditions and time course for attention shifts elicited by color discontinuities. The results of these experiments suggest that attention capture is contingent on attentional control settings induced by task demands.

Early visual information processing is often characterized as having two functionally distinct modes. One mode is *preattentive*; information is processed in a spatially parallel manner, without the need for attentional resources. The other mode involves the allocation of attentional resources to specific locations or objects for more complex analysis. Preattentive processes are assumed to perform basic analyses that segment or isolate areas of the visual array, creating functional perceptual units to which attention can then be allocated for more sophisticated, resource-demanding processing (Broadbent, 1958; Egeth, 1977; Hoffman, 1979; Julesz & Bergen, 1983; Neisser, 1967; Treisman & Gelade, 1980; Ullman, 1984).

Research concerning the specific manner in which attentional resources are selectively allocated or "oriented" to preattentively defined units or locations in the visual array has yielded evidence for another functional dichotomy. Attentional orienting can be voluntary, controlled by central strategic goals, or involuntary, driven by particular stimulus events (Jonides, 1980; Müller & Rabbitt, 1989; Posner, 1980; Remington, Johnston, & Yantis, 1992; Yantis & Jonides, 1984). These two modes of orienting attention in space have been referred to as *endogenous* and *exogenous* control, respectively (Posner, 1980; Posner & Cohen, 1984).

The present article investigates the characteristics of exogenous attention allocation. Specifically, we explore (a) the conditions under which exogenous attention orientation occurs and (b) the stimulus properties to which the exogenous

attention allocation system is sensitive. Our primary thesis is that the occurrence of involuntary shifts of attention is critically dependent on the relationship between stimulus properties and task demands and is not, as suggested by others, solely a function of the presence of specific stimulus properties (Jonides & Yantis, 1988; Müller & Rabbitt, 1989; Remington et al., 1992; Yantis & Jonides, 1984). We first review research on the characteristics of exogenous attention orientation. We then propose a new hypothesis concerning the mechanisms underlying exogenous shifts of attention. Finally, we present four experiments that test this new hypothesis.

Exogenous Attention Orientation

Exogenous attention orientation refers to shifts of spatial attention that are elicited by salient external events and that occur involuntarily (i.e., without conscious intent). A growing number of studies suggest that one stimulus property in particular—abrupt change in luminance over time—can generate involuntary, exogenous attention shifts (Jonides, 1981; Jonides & Yantis, 1988; Müller, 1989; Müller & Findlay, 1987; Müller & Rabbitt, 1989; Posner, 1980; Posner & Cohen, 1984; Yantis & Jonides, 1984). Jonides (1981), for example, presented subjects with "cues," consisting of a peripheral, abruptly onset bar marker, occurring near the probable location of a subsequent target character. These cues produced significant costs in response time to identify the target when they occurred at a nontarget location and benefits when at the target location. More important, the magnitude of costs and benefits was unaffected by manipulations of concurrent processing load, degree of cue validity, and task expectancy. Thus, Jonides argued that such cues produced involuntary, "automatic" shifts of attention.¹

Portions of this research were conducted while Charles L. Folk held a National Research Council—NASA ARC Research Associateship. Experiment 4 was conducted at Syracuse University. This research was supported in part by National Institute of Mental Health Grant 1 RO3 MH 45008 to Charles L. Folk.

We would like to thank Steven Yantis, Ray Klein, and two anonymous reviewers for comments on earlier versions of the article.

Correspondence concerning this article should be addressed to Charles L. Folk, Department of Psychology, Villanova University, Villanova, Pennsylvania 19085.

¹ The extent to which a process is involuntary or remains unaffected by intention is only one of the criteria for establishing "automaticity" (Hasher & Zacks, 1979; Posner & Snyder, 1975; Schneider & Shiffrin, 1977; Yantis & Jonides, 1990). The present experiments, however,

Similarly, Remington et al. (1992) found that response times for a two-choice letter discrimination task were elevated when the target letter was preceded by an abrupt flash of light at a location other than the target location. This effect occurred even when subjects were fully aware that the target would never occur at the location at which the flash occurred. Furthermore, Müller and Rabbitt (1989) found that even probabilistic foreknowledge of a target's location could not eliminate the distracting effect of an abrupt, peripheral luminance increment. More recently, however, Yantis and Jonides (1990) and Theeuwes (1991) have shown that when the location of an impending target is known with absolute certainty, the distracting effect of peripheral luminance changes is eliminated. These results suggest that when spatial attention is highly focused at one location (i.e., when spatial uncertainty is completely eliminated), involuntary shifts to changes in luminance do not, in fact, occur.

Evidence for the ability of dynamic luminance changes to elicit involuntary shifts of attention, at least under conditions of spatial uncertainty, is not limited to spatial cuing tasks. Yantis and Jonides (1984), for example, presented subjects with an abruptly onset character among a variable number of characters created by the removal of camouflaging line segments. They found that when the target was the abruptly onset character, response time did not increase with display size. The authors argued that attention was involuntarily shifted first to the location of the abruptly onset character, obviating any need for further search when that character was the target. Miller (1989) and Theeuwes (1991) have shown that abrupt offsets, in addition to onsets, can produce similar, albeit weaker, results.

There is also evidence to suggest that dynamic luminance changes are unique in their ability to elicit involuntary shifts of attention. Jonides and Yantis (1988) compared the effect of display size on response time when targets differed from distractors by abruptness of onset, by color, or by intensity. Only the abruptly onset targets reduced the display size effect. These results suggest that only *dynamic discontinuities*, such as abrupt onsets, are capable of generating involuntary attention shifts. Discontinuities defined by changes in feature properties across space that do not change over time (e.g., an "odd-man-out" in color), which we will refer to as *static discontinuities*, appear not to elicit such shifts.

In summary, this line of research has led to the conclusion that exogenous attention allocation is uniquely sensitive to dynamic discontinuities in visual information and that, under conditions of spatial uncertainty, the shifts elicited by such discontinuities are involuntary.

The Contingent Involuntary Orienting Hypothesis

Contrary to both of these conclusions, we propose that dynamic discontinuities do not, in fact, play a unique role in exogenous orienting of attention. We argue that the research

supporting this claim has not addressed the potentially critical relationship between properties exhibited by a cue (or distractor) and the properties involved in locating and processing targets. In the spatial cuing studies just discussed, the occurrence and location of both targets and cues were signaled by the same stimulus property (i.e., abrupt luminance change; Jonides, 1980; Müller & Findlay, 1987; Müller & Rabbitt, 1989; Remington et al., 1992). The finding that abrupt luminance changes yielded involuntary shifts of attention could have resulted from subjects' use of luminance changes to locate the target. Thus, previous findings leave open the possibility that irrelevant, abruptly onset cues only produce an unwanted shift of attention because subjects are "set" to attend to the very same onset properties of targets. A similar case might be made for the visual search studies conducted by Yantis and Jonides (1984; Jonides & Yantis, 1988), as well as Miller (1989); both targets and distractors were signaled by luminance changes. Thus, existing evidence for involuntary attentional orienting to luminance changes might only hold when the occurrence and location of both targets and cues (distractors) are signaled by luminance changes.

The lack of evidence for involuntary attention shifts to static discontinuities in color and luminance is similarly problematic. In the static discontinuity conditions of the visual search study by Jonides and Yantis (1988), in which an entire search set was simultaneously presented and one element of that set appeared in a unique color, locating the target did not require monitoring for a static discontinuity; on the majority of trials the target was not the character unique in color. Thus, unlike the luminance change conditions of Jonides and Yantis, the static discontinuity conditions did not require subjects to monitor the display for a property that specified the unique element.

There is some indirect evidence that when locating targets does require monitoring for a static discontinuity, other, irrelevant, static discontinuities can involuntarily attract attention. Pashler (1988, Experiment 6) had subjects indicate which side of a textured display contained a "form singleton," a texture element that differed from the others in shape. In addition to this form singleton, an irrelevant "color singleton" (i.e., an element different in color from the rest) appeared simultaneously with the target. Pashler found that the color singleton significantly impaired performance on the task, a result consistent with the notion that when the location of a target is signaled by a static discontinuity, an irrelevant static discontinuity can attract attention.²

² It should be noted that the results of Pashler (1988) regarding the effects of static discontinuities on attention allocation are consistent with the conclusions drawn from a substantial body of previous research. Many studies have shown that the detection of static discontinuities in "simple" visual feature dimensions (e.g., color, orientation, size, and brightness) occurs preattentively (e.g., Bergen & Julesz, 1983; Connor, 1972; Donderi & Zelnicker, 1969; Egeth, Jonides, & Wall, 1972; Folk & Egeth, 1989; Treisman & Gelade, 1980; Treisman & Gormican, 1988). Moreover, many authors have proposed that the detection of static discontinuities in simple features underlies preattentive grouping processes (Banks & Prinzmetal, 1976; Beck, 1972, 1974, 1982; Kahneman & Henik, 1977, 1981) and serves

focus solely on the "involuntariness" of shifts of spatial attention. We therefore refrain from addressing whether such shifts are truly automatic (see Yantis & Jonides, 1990, for a full discussion of the topic).

In sum, we propose that under conditions of spatial uncertainty, involuntary shifts of attention to a given stimulus event (e.g., a dynamic or a static discontinuity) will be contingent on whether that event shares a feature property that is critical to the performance of the task at hand. Specifically, involuntary orienting of attention will occur if the event shares the critical property and will not occur if it does not. We refer to this proposal as the *contingent involuntary orienting hypothesis*.

This hypothesis has two important empirical consequences: (a) If target detection does not require monitoring for an isolated, dynamic discontinuity (e.g., an abrupt onset), there should be no evidence of involuntary shifts to a cue (distractor) containing a similar, isolated, dynamic discontinuity, and (b) if processing of a target requires monitoring for a stimulus property other than an isolated, dynamic discontinuity, such as a static discontinuity in color, then a cue (distractor) containing that property will elicit an involuntary shift of attention.

Focus of Present Research

The following experiments were designed to provide a direct test of these predictions of the contingent involuntary orienting hypothesis. The experimental procedure was a variant of the spatial cuing paradigm, adapted from Remington et al. (1992). A peripheral precue was presented at one of four potential target locations. To provide a strong measure of the occurrence of involuntary shifts of attention, the spatial compatibility of the cue and target was varied across blocks of trials (i.e., the cue was either 100% valid or 100% invalid with respect to predicting the exact target location). Because, in the 100% invalid condition, subjects were fully aware that the cue location never predicted target location, they had a strong incentive to withhold any shift of attention to the cue if they could. Thus, cost in this condition, relative to control conditions, is assumed to reflect the inability of subjects to withhold an attentional response to the cue.³

To address the notion of contingent involuntary orienting of spatial attention, the stimulus properties signaling the location of the target and cue were independently manipulated. Thus, for a given property signaling target location we could observe the differential effects of cues signaled either by the same relevant stimulus property or a different, task-irrelevant property.

Experiment 1

A strong empirical prediction of our hypothesis is that abruptly onset cues should capture attention only when an abrupt onset is used to locate the target, not when some other property is used for target location. The first experiment was conducted to test this prediction. For one group of subjects,

the target was a single, abruptly onset character (an "X" or an "=") appearing in one of four possible locations. For another group, all four locations contained characters and the target was defined as the character different in color from the rest. In both conditions, the cues were abruptly onset flashes appearing around one of the four locations. Subjects made a forced-choice response with regard to target identity.

According to the contingent involuntary orienting hypothesis, the effect of invalid abruptly onset cues, relative to controls, should vary depending on the target-property condition. The strong form of the hypothesis predicts that 100% invalid onset cues will produce costs for onset targets and no costs for color targets. According to the extant hypothesis that involuntary shifts of attention are always and only elicited by the presence of dynamic changes in visual information (Jonides & Yantis, 1988), invalid abruptly onset cues should produce significant costs regardless of the property used to locate the target.

Method

Subjects. Forty subjects participated in this experiment. Twenty-four were paid volunteers, recruited from the NASA-Ames subject pool. The other 16 were Villanova University undergraduates who participated to partially fulfill a requirement for a general psychology course. Twelve subjects in each between-subjects condition were run at NASA—Ames Research Center and 8 in each condition were run at Villanova. The subjects ranged in age from 18 to 35 and all reported normal or corrected-to-normal visual acuity and color vision.

Apparatus. Stimuli were presented on a Princeton Graphics SR-12 monitor controlled by an IBM PC microcomputer equipped with a Sigma Design, Color 400 EGA high-resolution graphics board.

Stimuli. Displays were of three basic types: fixation display, cue display, and target display. Examples of each of these displays, along with their sequence of presentation, are shown in Figure 1A. The fixation display consisted of five boxes measuring 1.15° visual angle from a viewing distance of approximately 40 cm. Four boxes were located at the vertices of an imaginary diamond centered on the fifth box with a center-to-vertex distance of 4.7° visual angle. The boxes were light gray (IBM color designation 8) against the black cathode-ray tube (CRT) screen.

The cue displays consisted of the fixation display with the addition of four small circles, each subtending 0.36° visual angle, around the four sides of one of the five boxes. The circles were placed such that each was centered approximately 0.3° visual angle peripheral to its respective side of the box. The circles were high-contrast white (IBM color designation 15) against the black CRT screen.

Target displays consisted of the fixation display with the addition of a variable number of characters, depending on the between-subjects condition. A given character was either an X or an = and subtended 0.57° visual angle.

³ It has been argued that in the spatial cuing paradigm, blocking cue validity can lead to strategic differences in the treatment of precues, which can confound the interpretation of costs and benefits (Jonides & Mack, 1984). In the present experiments, however, inducing strategy differences across conditions was an explicit goal. Specifically, the 100% invalid condition was used to induce a strategy of completely ignoring irrelevant precues, allowing a strong test of whether shifts of attention are unaffected by the intentions of the subject (i.e., involuntary).

as the basis for resource allocation by systematically drawing attention to those groups (Bergen & Julesz, 1983; Cave & Wolfe, 1990; Hoffman, 1979; Kahneman & Henik, 1977, 1981; Koch & Ullman, 1985; Ullman, 1984; Wolfe, Cave, & Franzel, 1989).

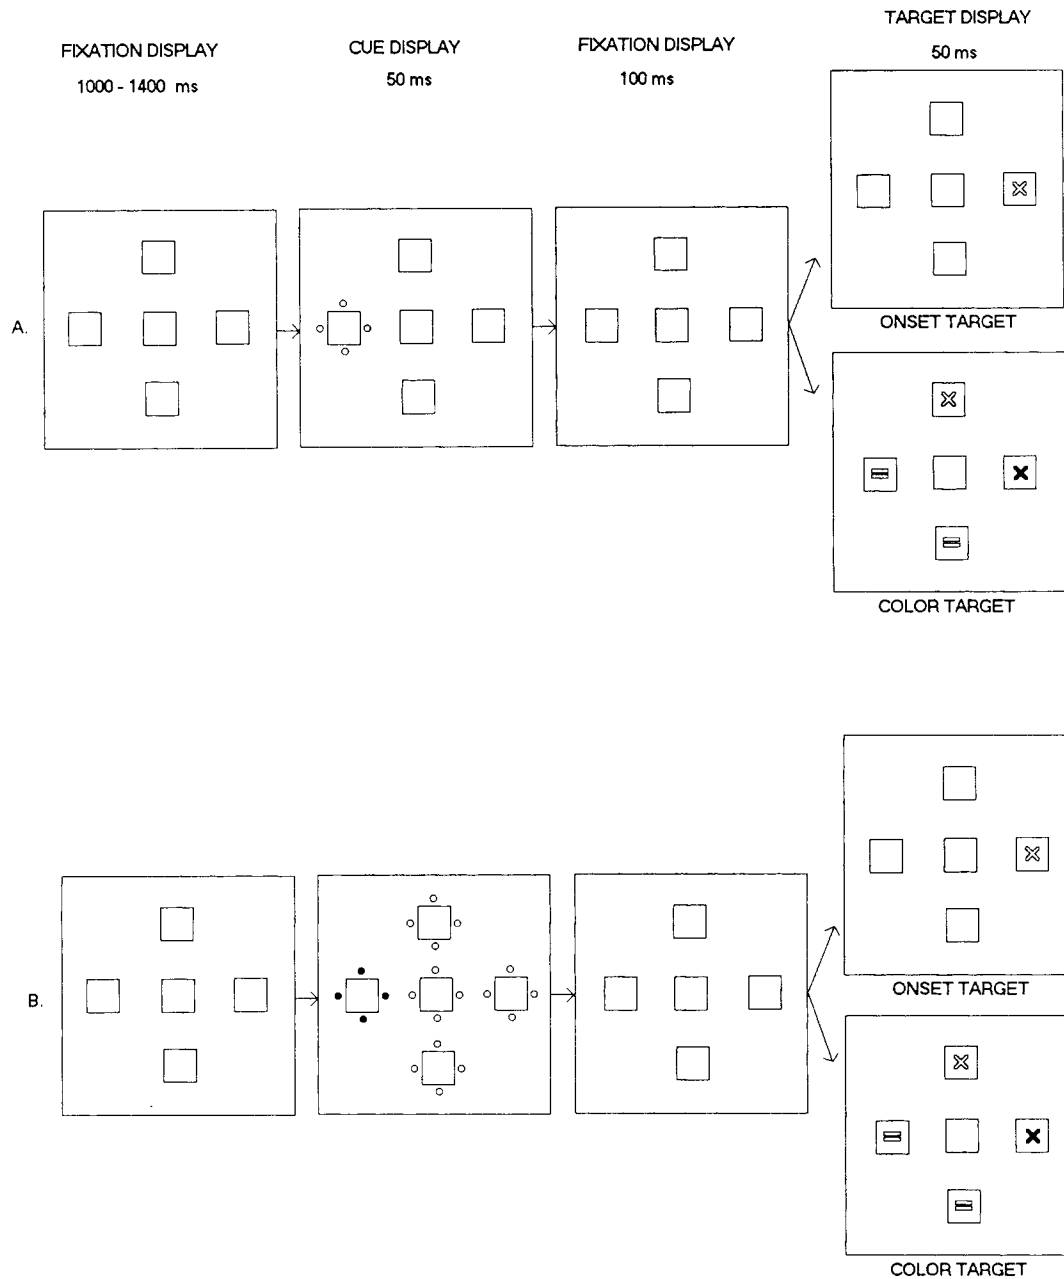


Figure 1. Examples of displays and sequence of events in Experiment 1 (A) and Experiment 2 (B). (Both examples represent invalid-cue trials. Outlined characters were actually white on the computer monitor; black characters were red on the monitor.)

In the onset-cue-onset-target condition, a single X or = (the target) appeared in one of the four outer boxes. This target character was high-contrast white (IBM color designation 15) against the black CRT screen.

In the onset-cue-color-target condition, each of the four outer boxes contained an X or an = character. Three of the characters were high-contrast white (IBM color designation 15) and one (the target) was high-contrast red (IBM color designation 12).

Design. As described earlier, two between-subjects conditions were created by varying the stimulus property signaling target location. In the onset-target condition, the target was a single, abruptly

onset character appearing in one of the four outer boxes. In the color-target condition, each box contained a character and the target was defined as the one red character among the white characters. Subjects were told to respond to this red character. Examples of each type of target display are shown in Figure 1A.

For each of the between-subjects conditions, four within-subjects cue conditions were created by varying the spatial relationship between the cue and the target. In the valid-cue condition, the target always appeared in the box marked by the cue. In the invalid-cue condition, the target always appeared in one of the three boxes not marked by the cue. In the center-cue condition, the cue always

appeared around the center box and the target appeared in one of the four outer boxes. Finally, in the no-cue condition, no cue was presented around any of the boxes in the cue display. These no-cue trials served as the control for the other three conditions.

Each of the four cue conditions was presented to each subject in separate blocks of trials. Therefore, in the invalid- and center-cue conditions, the cue was 100% invalid with respect to the target position; in the valid-cue condition, it was 100% valid.⁴ In other words, prior to each trial, subjects were fully aware of the spatial relationship between the cue and target.

Each cue condition consisted of two contiguous blocks of 48 trials. Cue condition order was counterbalanced across subjects according to a Latin square. Subjects were randomly assigned to orders.

Within each block of 48 trials, each of the two possible targets (X or =) appeared in each of the four possible locations equally often. In the color-target condition, the white, nontarget characters appearing in the three nontarget locations were chosen randomly from among X and =. Within each invalid- and valid-cue block, the cue appeared in each of the four possible locations equally often. Cue positions in the invalid-cue blocks were chosen with the additional constraint of appearing equally often in each of the three possible nontarget locations for each possible target location.

Ten practice trials were presented at the beginning of each block. In addition, after any trial on which an error was made, a "buffer" trial, chosen randomly from the set of possible trial parameters for that condition, was inserted. Responses to practice and buffer trials were not included in the data analysis.

Procedure. Each subject participated in a single 1-hr session. Subjects were given written and oral descriptions of the stimuli and task. They were fully informed of the blocked structure of the experiment and were encouraged to take advantage of the knowledge of cue-target positions if possible. Subjects were instructed to respond "as quickly as you can while making as few errors as possible." Maintaining fixation on the center box in the display was heavily stressed. Subjects were told that because the trial events would be occurring very rapidly, attempting to make eye movements would impair performance. The experimenter answered any questions the subject had, and then led the subject to a dimly lit testing room and familiarized him or her with the experimental equipment. The subject was seated at a distance of approximately 40 cm from the computer CRT screen.

Each block began with the presentation of a screen indicating which of the four cue conditions (no cue, valid cue, central cue, or invalid cue) and which block of that condition was about to begin. Subjects then pressed a key to begin the sequence of trials in that condition. A message appeared at the end of each block instructing the subject to rest before beginning the next block. Subjects then pressed a key when they were ready to continue.

The sequence of events on a given trial began with a 500-ms presentation of the fixation display. The center box then blinked off for 100 ms, which served as a warning signal to subjects that the trial was beginning. The fixation display then remained on for a foreperiod interval randomly chosen from the set 1,000, 1,100, 1,200, 1,300, or 1,400 ms. After the foreperiod, the cue display was presented for 50 ms. The target display then appeared for 50 ms after an interstimulus interval of 100 ms. Phenomenally, the five display boxes appeared to remain on the screen throughout each trial (except for the warning blink of the center box). The total interval between cue onset and target offset was 200 ms; at this short interval saccadic eye movements to the cue or other locations were unlikely (see Colgate, Hoffman, & Eriksen, 1973), allowing any effect of cue validity to be attributed to attentional, not acuity, factors.

Subjects made a speeded, forced-choice response to the target display by pressing one of two labeled keys on the keyboard. If the target was an =, subjects were instructed to press the left key with

their left index finger. If the target was an X, subjects were to press the right key with their right index finger. A 500-ms, 1000-Hz tone was sounded by the computer for an incorrect response. If a response was not made within 1,500 ms, an error was scored and the trial sequence continued. Following the subject's response, an intertrial interval of 500 ms elapsed before the center box blinked to indicate initiation of the next trial.

Response time and error rates for each trial were measured and recorded by the computer. Response time was measured from the onset of the target display until a response was made. Error trials and trials on which response time was less than 200 ms were excluded from the data analysis.

Results and Discussion

Average mean response times, error rates, and standard errors for each cue condition within each target property, pooled across target identity and block, are shown in Table 1. Preliminary analyses established that target identity and block did not interact in any way with cue condition.⁵ Thus, for the following analyses, response times were pooled across these two factors.

According to the contingent involuntary orienting hypothesis, the effect of abrupt-onset cues should depend on the stimulus property used to locate the target. To test this hypothesis, two sets of planned comparisons were performed on the data. The first was conducted to assess the effects of abrupt-onset cues relative to no-cue controls within each target-property condition. The second was conducted to determine the degree to which these effects varied across target-property conditions.

Within-condition analysis. Separate within-subjects analyses of variance (ANOVAs) were conducted on the response time and error data for the two target-property conditions with cue condition as the single repeated measures variable. Analysis of mean response times for the onset-target condition yielded a significant main effect of cue condition, $F(3, 57) = 8.95, p < .001$. A Newman-Keuls post hoc analysis conducted on the means ($\alpha = .05$) revealed that invalid cues produced a significant cost in response time relative to no-cue trials (as well as to valid- and center-cue trials). Valid cues, on the other hand, yielded a significant 36-ms benefit in response time relative to any of the other three cue conditions. Center-cue and no-cue response times did not differ significantly.

⁴ The terms *valid* and *invalid* are typically used to refer to expectancy information provided by a cue. In the present case, however, these terms refer strictly to the relationship between the spatial positions of the cue and target, not to the expectancy information they may provide to the subject. Technically, according to the latter usage, the cues in all three of these conditions are valid, in that subjects know with 100% certainty whether the target will appear in the cued position or not. Nonetheless, we use these terms for ease of exposition.

⁵ The only significant effects involving these factors occurred in the color-target condition. Response time in the second block was, overall, 10 ms faster than in the first, $F(1, 16) = 8.58, p < .01$. In addition, X responses were 25 ms faster overall than = responses, $F(1, 16) = 23.72, p < .001$. The most likely reason for this effect is that the X response was mapped onto the right key and no attempt was made to control for handedness.

Table 1
Average Mean Response Times (RT) and Error Rates (ER)
as a Function of Cue Condition and Target-Property
Condition: Experiment 1

Cue condition	Cue-property–target-property condition			
	Onset-cue– onset-target		Onset-cue– color-target	
	RT	ER	RT	ER
No cue	502	.068	520	.036
Valid cue	466	.078	476	.050
Center cue	498	.086	497	.050
Invalid cue	528	.080	516	.052

Note. Response times are in milliseconds.

This pattern replicates the essential findings of Remington et al. (in press).

For the color-target condition, analysis of response times also yielded a significant main effect of cue condition, $F(3, 57) = 9.08, p < .001$. However, unlike the onset-target condition, a post hoc Newman-Keuls analysis revealed no significant cost for invalid-cue trials relative to no-cue trials. On the other hand, as in the onset-target condition, valid cues produced a significant benefit in response time relative to the no-cue trials, as well as to center- and invalid-cue trials. Center cues also generated significant benefits relative to no-cue controls, as well as to invalid-cue trials.

The results of these within-condition analyses are consistent with our hypothesis. The exact same abrupt-onset cues had differential effects on performance depending on the property signaling the location of the target. Specifically, when target location was signaled by an isolated abrupt onset, there was significant cost for invalid, abrupt-onset cues. This is consistent with a model in which attention was drawn to the location of the abrupt-onset cue. However, when target location was signaled by color, there was no cost for invalid, abrupt-onset cues and, thus, no evidence that attention was drawn to the abrupt-onset cue in this condition.

It is important to note that although the costs associated with invalid cues varied depending on the target-property condition, significant benefits were obtained for 100% valid cues relative to the no-cue trials, regardless of target property. In the color-target condition, the abrupt-onset cues could be used to shift attention when desired, even though they did not appear to cause involuntary shifts when the cue was invalid.

One puzzling aspect of the present results concerns performance with center cues. As was the case with invalid cues, center cues always indicated a location where the target would not occur (i.e., in the center box). Thus, by our logic, these cues should have produced significant costs in the onset-target condition and no costs in the color-target condition. Surprisingly, these cues produced no costs in the onset condition and significant benefits in the color condition. We have no sound explanation for this pattern. One possibility is that the variations in center-cue response times relative to the other cuing conditions represent random error, perhaps indicating a large

natural variability for this condition. On the other hand, they could reflect differential effects of the alerting properties of the center cue on the color-target and onset-target tasks. We will return to this point in the discussion of Experiments 1 and 2.

Error rates averaged 7.8% and 4.7% for the onset- and color-target conditions, respectively. Repeated measures ANOVAs on the error data yielded no significant effect of cue condition (both $F_s < 1$). Thus, there was no evidence that subjects in either condition traded speed for accuracy across cue conditions.

Across-condition analysis. In an effort to provide a direct statistical test of the differential effects of invalid abrupt-onset cues across the two target-property conditions, a 2×2 mixed ANOVA was performed with a subset of the cue conditions (invalid cue vs. no cue) as the within-subjects variable and target property (onset vs. color) as the between-subjects variable. Neither of the main effects were significant, $F(1, 38) = .01, p > .05$, for cue condition and $F(1, 38) = 2.10, p > .05$, for target property. The interaction between these two variables, however, was significant, $F(1, 38) = 4.40, p < .05$, confirming the results of the previous post hoc tests showing that significant costs were present for onset targets but not for color targets.

The overall pattern of results in this experiment is consistent with the notion of contingent involuntary attention shifts; costs for invalid cues, which are assumed to index involuntary shifts of attention, occur only when the distracting cue contains the stimulus property used to locate the target. The results are clearly inconsistent with the view that under conditions of spatial uncertainty, dynamic discontinuities (i.e., abrupt onsets) produce involuntary shifts of attention regardless of the stimulus properties that are relevant to the task.

Experiment 2

A second empirical prediction of the contingent orienting hypothesis is that stimulus properties other than dynamic discontinuities, such as static discontinuities in color, will elicit involuntary shifts of attention if locating the target requires monitoring for that property. Experiment 2 was conducted to test this prediction. As in Experiment 1, subjects searched for and identified either an abruptly onset target or a color target. Instead of abruptly onset cues, however, the location of the cue in this experiment was signaled by a static discontinuity in color. All boxes in the cue display were surrounded by white circles except for one box (i.e., the cued box) that was surrounded by red circles. Cue validity was varied as in Experiment 1. According to our hypothesis, significant costs should occur for invalid color cues when paired with a color target but not when paired with an isolated onset target.

Method

Subjects. Forty new subjects participated in this experiment. Twenty-four were paid volunteers recruited from the NASA—Ames subject pool. The other 16 were Villanova University undergraduates who participated to partially fulfill a requirement for a general psy-

chology course. Twelve subjects in each between-subjects condition were run at NASA—Ames Research Center and 8 in each condition were run at Villanova. One subject was replaced as a result of equipment failure. The subjects ranged in age from 18 to 33 and all reported normal or corrected-to-normal visual acuity and color vision.

Apparatus and stimuli. The apparatus was exactly the same as that used in Experiment 1. The stimuli were also the same with the exception of the cue displays. Cue displays consisted of the fixation display with the addition of cue circles (see description of cues in Experiment 1) around all five boxes. The cue circles around one of the five boxes were red (IBM color designation 12). The cue circles around all other boxes were white (IBM color designation 15). The box surrounded by the red circles was defined as the cued location. An example of a cue display for an invalid trial is shown in Figure 1B. In the no-cue condition, there were no red circles; all boxes were surrounded by white circles. Note that although the cues still had an abrupt onset, this property was not unique: all of the circles around the boxes had an abrupt onset. The cued location was determined by a difference in color from the circles at other locations.

Design and procedure. With the exception of the nature of the cue displays (described earlier), the design and procedure for this experiment were identical to those used in Experiment 1.

Results and Discussion

The average mean response times and error rates for each of the four cue conditions within each target-property condition, pooled across target identity and block, are shown in Table 2.⁶ As in Experiment 1, two sets of planned comparisons were performed to determine the effect of the color cue within each target-property condition as well as across conditions.

Within-condition analyses. Separate within-subjects ANOVAs were conducted on mean response times and error rates for each target property with cue condition as the single variable. For the color-target condition, the overall effect of cue condition on response time was significant, $F(3, 57) = 24.98, p < .001$. A Newman-Keuls post hoc analysis ($\alpha = .05$) revealed a significant cost for invalid-cue relative to no-cue trials. Valid cues, on the other hand, produced a significant benefit relative to no-cue trials, as well as to the other three conditions. Mean response time for center-cue trials was significantly greater than for no-cue or valid-cue trials but did not differ significantly from invalid-cue trials. This pattern of response times is consistent with our hypothesis; static discontinuities elicited involuntary attention shifts when the target location was signaled by the same static discontinuity. Invalid cues produced significant costs relative to the no-cue trials even though subjects knew with 100% certainty that the color cues would be in a location that was incompatible with respect to the target location.

The ANOVA for the onset-target condition also yielded a significant effect of cue condition on response time, $F(3, 57) = 2.99, p < .05$. As in the color-target condition, a Newman-Keuls test revealed a significant cost for invalid color cues relative to no-cue trials. This cost for invalid cues, however, was accompanied by significant costs for both valid and center cues relative to the no-cue condition. In fact, mean response times for invalid cues did not differ significantly from the mean for center-cue or for valid-cue trials. Thus, because response time was unrelated to variation in the spatial rela-

Table 2
Average Mean Response Times (RT) and Error Rates as a Function of Cue Condition and Target-Property Condition: Experiment 2

Cue condition	Cue-property-target-property condition			
	Color-cue-color-target		Color-cue-onset-target	
	RT	ER	RT	ER
No cue	510	.032	479	.067
Valid cue	489	.047	492	.063
Center cue	542	.032	494	.058
Invalid cue	552	.033	500	.071

Note. Response times are in milliseconds.

tionship of cue to target (i.e., there were no differences in response time as the validity of the cue was manipulated), it is unlikely that the cost obtained for invalid color cues was due to shifts in spatial attention.

Unlike Experiment 1, the center-cue conditions in the present experiment behaved in a fashion consistent with our hypothesis; they produced significant costs in the color-target condition but not in the onset-target condition. The latter result suggests that center cues signaled by a color discontinuity did not capture attention when the target location was signaled by an abrupt onset. The former result, however, is qualified by the fact that costs were found for all cue types in the color-target condition. As pointed out earlier, this suggests that any costs obtained cannot be attributed to spatial shifts of attention.

Error rates for the color-target and onset-target conditions averaged 6.5% and 3.6%, respectively. ANOVAs yielded no significant effects of cue condition on error rate for either target-property condition (both $F_s < 1$).

Across-condition analysis. As in Experiment 1, to assess whether the effect of invalid color cues varied significantly as a function of target-property condition, the no-cue and invalid-cue conditions for each target-property condition were entered into a 2×2 mixed ANOVA with cue condition as the within-subjects variable and target property as the between-subjects variable. Both main effects were significant; response times for color targets were greater than for onset targets, $F(1, 38) = 4.21, p < .05$, and response times for invalid cues were greater than for no-cue trials, $F(1, 38) = 39.60, p < .05$. More important, the interaction between these variables was significant, $F(1, 38) = 4.36, p < .05$. Simple effects analyses revealed that response time for the no-cue condition did not differ for color and onset targets, $F(1, 38) = 2.21, p > .05$. For invalid color cues, however, response times were significantly greater in the color-target than in the onset-target condition, $F(1, 38) = 6.26, p < .05$. Thus, the

⁶ Preliminary analyses once again confirmed that target identity and block did not interact significantly with cue condition. The only significant effect associated with these variables was a 6-ms advantage for the second block over the first block of the color-target condition, $F(1, 16) = 7.51, p < .05$.

exact same color cues had differential effects on performance depending on whether subjects were monitoring for a color target or an onset target.

The results of this experiment support the second prediction of the contingent involuntary orienting hypothesis. When a static discontinuity in color was used to locate the target, invalid cues consisting of a similar discontinuity in color produced response time costs, suggesting that attention was involuntarily shifted to the location of the cue. This provides direct evidence that stimulus properties other than dynamic discontinuities can elicit exogenous attention shifts.

Discussion: Experiments 1 and 2

The pattern of results across the first two experiments is consistent with the hypothesis that involuntary shifts of spatial attention to a given stimulus event are contingent on whether that event shares a feature property for which the attentional system is set to respond. Table 3 summarizes the costs obtained with invalid cues for each of the four possible cue-target combinations. For a given cue type, the magnitude of costs varied significantly as a function of target property. Specifically, costs were larger when cue and target location were both signaled by the same stimulus property than when signaled by different properties.

There are, however, a number of aspects of the first two experiments that could potentially weaken the conclusions drawn above. First, as is clear in Table 3, the strong version of the contingent involuntary orienting hypothesis is violated in the color-cue-onset-target condition of Experiment 2, in that a significant cost was obtained. As discussed earlier, this cost may represent a systematic interference effect associated with color cues that is independent of the effects of spatial attention and that adds a constant increase in response time relative to control. Consistent with this interpretation, notice that in addition to the 21-ms costs in the color-cue-onset-target condition, the magnitude of costs in the color-cue-color-target condition is approximately 20 ms larger than in the onset-cue-onset-target condition. On the other hand, we cannot absolutely rule out the possibility that the cost found in the color-cue-onset-target condition is related to involuntary shifts of attention.

Second, our conclusions are based on the assumption that the no-cue conditions are appropriate baselines for assessing costs for 100% invalid cues. Our no-cue conditions may not, in fact, provide an appropriate baseline, in that alerting effects

associated with the presentation of a cue are not present in the no-cue condition. Although this inequality in alerting would, if anything, result in a conservative estimate of costs, we cannot rule out the possibility that some complicated interaction of alerting with target property and cue validity variables influenced the outcome in favor of the contingent involuntary orienting hypothesis.

One might argue that the center-cue condition is a more appropriate baseline in that it controls for the alerting effects of invalid and valid cues. If the center-cue condition is used as a baseline, the pattern of results would suggest that abrupt onsets are indeed special, in that invalid cues produced significantly longer response times than center cues in Experiment 1, regardless of target property (although note that the magnitude of the effect would still vary with target property). Using the center-cue condition as a baseline, however, is also extremely problematic. First, because such cues are presented foveally, they may produce greater alerting effects than peripheral cues, resulting in a liberal estimate of costs. Second, center cues necessarily occur at a spatial location, and therefore have the potential to draw attention to, or hold attention at, that location. Thus, because the target never actually occurs in the center box, this condition is similar to the 100% invalid condition, mitigating its usefulness as a baseline.

Given these concerns, and the importance of the hypothesis under investigation, Experiment 3 was conducted to provide supporting evidence for contingent involuntary orienting from a converging methodology.

Experiment 3

In the first two experiments, trials were blocked by cue validity to allow for a 100% invalid condition that would provide a strong test of the involuntariness of attention capture by the cue (see also Remington et al., 1992). Having provided evidence that involuntary attentional capture does occur in some conditions despite strong incentives to avoid it, we were able to relax this constraint somewhat to investigate the contingent involuntary orienting hypothesis under conditions in which the differential effects of alertness have been eliminated. In Experiment 3, onset and color cues were again paired with onset and color targets, but cues appeared on every trial and the location of the cue was uncorrelated with target location. With four possible target locations, this means the cue was valid on 25% of the trials and invalid on the remaining 75%. Thus, subjects should have very little incentive to voluntarily shift attention to the location of the cue. Any difference in response time as a function of cue validity would indicate the inability of a subject to withhold an attentional response to the cue (i.e., an involuntary shift of attention). The logic of this experiment was identical to that used by others investigating involuntary attention shifts (e.g., Jonides & Yantis, 1988; Müller & Rabbitt, 1989; Yantis & Jonides, 1984).

According to the contingent involuntary orienting hypothesis, there should be an effect of cue validity only when the cue shares the property used to locate the target. On the other hand, if abrupt onsets are special in their ability to involuntarily capture attention, significant cue validity effects should

Table 3
Response-Time Costs for 100% Invalid Cue Conditions (Invalid Cue – No-Cue Control) at Each Cue-Property-Target-Property Combination: Experiments 1 and 2

Cue property	Target property	
	Onset target (MS)	Color target (MS)
Onset cue	26*	4
Color cue	21*	42*

* $p < .05$.

be found for onset-cue conditions, regardless of whether the target is signaled by abrupt onset or a color discontinuity.

Method

Subjects. Twenty-four new subjects participated in this experiment. All were Villanova University undergraduates, ranging in age from 18 to 20, who participated to partially fulfill a requirement for a general psychology course. Subjects were tested for normal color vision and visual acuity of 20/30 or better using a Titmus II Vision Tester.

Apparatus and stimuli. The apparatus and stimuli were identical to those used in the first two experiments with the exception that center-cue and no-cue displays were not included in the design.

Design and procedure. The two cue types used in Experiments 1 and 2 (onset vs. color discontinuity) were crossed with the two target types (onset vs. color discontinuity) to yield four cue-property-target-property conditions that were presented in separate blocks of trials. All subjects participated in all four conditions. Condition order was counterbalanced across subjects using a Latin square. Subjects were randomly assigned to orders.

Each cue-property-target-property condition consisted of three contiguous blocks of 32 trials. Fourteen additional practice trials were presented at the beginning of each block and were not included in the data analysis. Within each block of trials, each of the two targets (X vs. =) appeared equally often in each of the four possible target locations. Cues also appeared equally often around each of the four locations but were completely uninformative with respect to predicting target location. Specifically, in each block, the target appeared at the cued location on exactly one fourth of the trials (valid cue) and at uncued locations on three fourths of the trials (invalid cue). For a given target location, invalid cues appeared equally often at each of the three remaining locations.

The sequence of events, response requirements, and data collection were identical to those used in Experiments 1 and 2. Prior to beginning the experiment, subjects were fully informed of the nature of the stimuli, conditions, and task through both written and oral descriptions. Each block began with the presentation of a screen indicating which of the four cue-property-target-property conditions, and which block of that condition, was about to begin. In all other procedural respects, the experiment was identical to the first two experiments.

Results

Mean response times and error rates for valid and invalid trials at each combination of cue and target property, pooled across target identity and block, are shown in Figure 2. A $2 \times 2 \times 2$ repeated measures ANOVA, with variables of target property (onset vs. color discontinuity), cue property (onset vs. color discontinuity) and cue validity (valid vs. invalid) was performed on the mean response times. The analysis revealed that response times for invalid cues were significantly longer than for valid cues, $F(1, 23) = 12.39, p < .01$, for the main effect of cue validity. Cue validity also interacted with target type, $F(1, 23) = 5.67, p < .05$, indicating that the overall effect of cue validity was greater for color targets than for onset targets. Most important, the analysis yielded a significant three-way interaction among target property, cue property, and cue validity, $F(1, 23) = 20.22, p < .001$. No other interactions were significant.

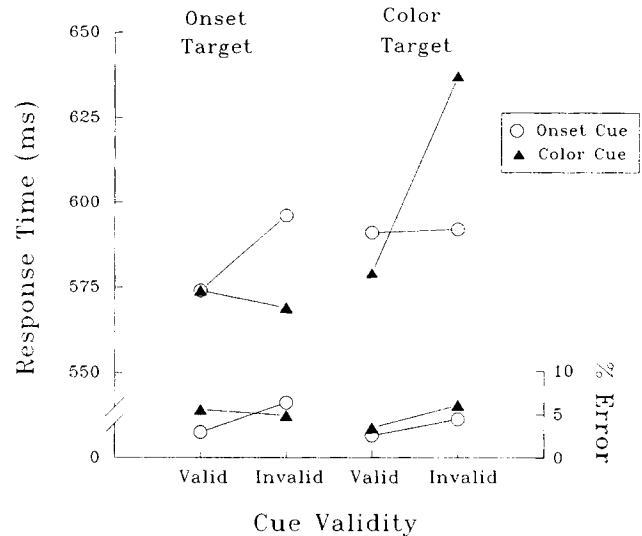


Figure 2. Average mean response times and error rates as a function of target property, cue property, and cue validity in Experiment 3.

The significant three-way interaction confirms what is evident in Figure 2. The effect of cue validity varied as a function of both cue property and target property. To explore this interaction, simple interaction comparisons were performed on the effect of cue property on cue validity at each level of target property. For onset targets, only the interaction between cue property and cue validity reached significance, $F(1, 23) = 6.11, p < .05$. Simple effects analyses confirmed that onset cues produced an effect of cue validity, $t(23) = 2.03, p < .05$ (one-tailed), whereas color cues did not, $t(23) = 0.14, p > .05$. For color targets, there was a main effect of cue validity, $F(1, 23) = 18.19, p < .001$, in addition to a significant interaction between cue property and cue validity, $F(1, 23) = 17.20, p < .001$. Simple effects analyses yielded the reverse pattern from that found with onset targets. Color cues produced a significant validity effect, $t(23) = 4.88, p < .001$ (one-tailed), and onset cues did not, $t(23) = 0.83, p > .05$.

Overall mean error rate averaged 4.6%. There was little evidence of a speed-accuracy trade-off; mean error rates were positively correlated with response times ($r = .45$). An ANOVA yielded results similar to those found with response time. There was a significant main effect of cue validity, $F(1, 23) = 11.90, p < .01$, as well as a significant three-way interaction among target property, cue property, and cue validity, $F(1, 23) = 7.21, p < .05$. Simple interaction comparisons of cue property and cue validity at each level of target property yielded a significant interaction between these two variables in the onset-target condition but not in the color-target condition, $F(1, 23) = 8.03, p < .01$, and $F(1, 23) = 0.41, p > .05$, respectively. The only other significant effect was of cue validity in the color-target condition, $F(1, 23) = 17.44, p < .001$. The results of simple effects analyses in the onset-target condition mimic the response time data. Onset cues produced significant validity effects and color cues did not, $t(23) = 4.62, p < .01$ (one-tailed), and $t(23) = 0.69, p > .05$, respectively.

Discussion

The results of Experiment 3 provide clear converging evidence for contingent involuntary orienting. Table 4 provides a summary of the effect of cue validity (i.e., invalid – valid cues) at each combination of cue and target property. When the locations of the cue and target were signaled by the same stimulus property, there was a significant effect of cue validity. When they were signaled by different properties, there was virtually no effect of cue validity.

At an empirical level, these results provide additional evidence that abrupt onsets (or perhaps dynamic discontinuities in general) do not always capture attention and that static discontinuities can capture attention when the task requires monitoring for such a discontinuity. The present results also indicate that the magnitude of cue validity effects is greater for color-cue-color-target conditions (56 ms) than for onset-cue-onset-target conditions (22 ms), $t(23) = 2.39, p < .05$. Despite the changes in methodology, these values are similar to the cost values for the analogous conditions in Experiments 1 and 2 (42 and 26 ms, respectively). Finally, the general convergence of these results with those of the first two experiments makes it unlikely that the pattern of costs across Experiments 1 and 2 was due to differential alerting effects associated with cued and no-cue conditions. Moreover, it suggests that the no-cue conditions were indeed appropriate baselines to assess involuntary attention shifts.

Experiment 4

The first three experiments suggest that involuntary shifts of attention to a given stimulus are contingent on whether that stimulus shares properties that are critical to performing the target task. However, the experiments were ambiguous with respect to the specific properties on which these contingencies are based. In the conditions in which static-discontinuity cues produced significant costs, the cue and target shared the same specific color as well as the same type of discontinuity. For example, in the color-cue-color-target conditions of Experiments 2 and 3, both cues and targets were red. Thus, the observed contingent orienting could be based not at the level of static (color) versus dynamic (abrupt onset) discontinuities, but at the level of specific feature values (e.g., the color red). Experiment 4 was conducted primarily to distinguish between these two possibilities.

In this experiment, the color-target condition from Experiment 2 was replicated with two modifications. As in Experi-

ment 2, the target was defined as the red character among white characters. The cue, however, consisted of green circles surrounding one of the four boxes with white circles around the others. With this design, cues and targets were both static, color discontinuities but differed with respect to the specific feature value defining the discontinuity. If contingent orienting operates at the level of color discontinuities in general, then the results should be similar to those found in Experiment 2; invalid cues should produce significant costs. If, however, the contingencies are based at the level of specific feature values (e.g., red), then a spatially incompatible green cue should not produce costs when monitoring for a red target.

A second modification of the experimental procedure was introduced in Experiment 4 in an effort to investigate the time course of shifts in visual attention to static discontinuities. A number of studies have traced the time course of voluntary or endogenous shifts of attention, as well as involuntary shifts in response to abrupt onsets, by manipulating the stimulus onset asynchrony (SOA) between cue onset and target onset in the spatial cuing paradigm (Müller & Findlay, 1987; Müller & Rabbitt, 1989; Remington et al., 1992; Shepard & Müller, 1989). Thus, to determine whether the time course of shifts to static discontinuities is different in any fundamental way from the time course of these other shifts, the cue-target SOA in the present study was systematically varied between 0 and 200 ms.

Method

Subjects. Subjects were 15 undergraduate students from Syracuse University who participated to partially fulfill a course requirement. Ages ranged from 18 to 20. All subjects reported normal color vision and normal or corrected-to-normal visual acuity.

Apparatus and stimuli. The apparatus was exactly the same as that used in Experiments 1 and 2. The stimuli were the same as those used in the color-target condition of Experiment 2 with the exception of the cue displays. The colored circles marking the cued location in the cue displays were high-contrast green (IBM color designation 10) instead of red. Thus, although the cue and target were both defined by a spatial discontinuity in color, they differed with respect to the particular color (i.e., green cue and red target).

Design and procedure. The experiment consisted of three cue conditions (no cue, valid cue, and invalid cue) crossed with five different cue-target SOAs measured from cue onset (0, 50, 100, 150, and 200 ms).⁷ At the 0-ms SOA, the cue and target occurred simultaneously for 50 ms. At the 50-ms SOA, the target display appeared immediately after the offset of the cue display. The three cue conditions were presented in separate blocks of 120 trials each. Each possible combination of SOA, target identity, cue position, and target position occurred equally often in each block. Subjects were allowed to rest halfway through each condition. Cue condition order was varied across subjects. In all other respects, the design and procedure were identical to that used in the color-target condition of Experiment 2.

Table 4
Effect of Cue Validity (Invalid Cue – Valid Cue) at Each Cue-Property-Target-Property Combination: Experiment 3

Cue property	Target property	
	Onset target (MS)	Color target (MS)
Onset cue	22*	1
Color cue	-5	56**

* $p < .05$. ** $p < .01$.

⁷ The center-cue condition was not incorporated into this experiment.

Results

Mean response times and error rates for each cue condition at each of the five SOAs are shown in Figure 3. A 3×5 repeated measures ANOVA was performed on the mean response times with cue condition (no cue, valid cue, or same-cue) and SOA (0, 50, 100, 150, or 200) as variables. The effect of cue condition was significant, $F(2, 28) = 6.28, p < .01$. A Newman-Keuls analysis ($\alpha = .05$) of this effect confirmed that invalid cues produced an overall cost relative to no-cue and valid-cue trials. This result replicates the findings of the color-target condition in Experiment 2.

The main effect of SOA was also significant, $F(4, 56) = 33.84, p < .0001$, confirming that across all cuing conditions response time decreased with increasing SOA. This effect is well documented and is generally assumed to reflect the alerting properties of the cue, independent of any spatial effects (Posner, 1980; Posner & Boies, 1971; Shulman, Remington, & McLean, 1979).

As is evident in Figure 3, the effect of SOA differed depending on cue condition, a result confirmed by a significant interaction between cue condition and SOA, $F(8, 112) = 5.23, p < .0001$. Simple main effects analyses yielded significant effects of cue condition at the 50-, 100-, 150-, and 200-ms SOAs (all $ps < .05$), but not at the 0-ms SOA, $F(2, 28) = 2.98, p > .05$. A Newman-Keuls test ($\alpha = .05$) was then conducted at each of these four SOAs to examine the costs and benefits associated with the valid and invalid cues relative to the no-cue condition. These costs and benefits are listed in Table 5. At all four SOAs, invalid cues produced a significant cost relative to the no-cue condition. In none of the conditions, however, did the benefit for the valid-cue condition relative to the no-cue condition reach significance. The Newman-Keuls analysis also established that response times for valid




Figure 3. Average mean response times and error rates as a function of cue condition and stimulus onset asynchrony in Experiment 4. (V = valid-cue condition; N = no-cue condition; I = invalid-cue condition).

Table 5

Mean Difference Values for Invalid and Valid Conditions Relative to No-Cue Control at Each Stimulus Onset Asynchrony in Experiment 4

Difference measure	Stimulus onset asynchrony				
	0	50	100	150	200
Invalid – no cue	19	43*	34*	45*	47*
Valid – no cue	37	-1	-6	-18	-10

Note. Positive numbers reflect a cost and negative numbers a benefit. All values are in milliseconds.

* $p < .05$.

cues were significantly less than invalid-cue response times at each SOA.

Analysis of mean error rates yielded a marginally significant main effect of SOA, with lower error rates for longer SOAs, $F(4, 56) = 2.25, p < .08$. The main effect of cue condition was not significant, $F(4, 56) = .74, p > .05$, nor did it interact with the SOA variable, $F(8, 112) = 1.70, p > .05$.

Discussion

Response times for invalid cues were significantly elevated with respect to no-cue trials at all but the 0-ms SOA. This pattern is similar to that found in Experiment 2 and is consistent with an involuntary shift of spatial attention to the location of the invalid cue. Moreover, costs were obtained even when the cue was not the exact same color as the target, suggesting that the contingencies for exogenous attention orientation operate at the level of static discontinuities in general rather than at the level of specific values along particular feature dimensions.

The pattern of interaction between cue condition and SOA in the present experiment is similar to that found for cues defined by dynamic discontinuities (Müller & Rabbitt, 1989; Remington et al., 1992). There was no significant effect of the cue with simultaneous presentation of cue and target. Significant costs were evident, however, as early as 50 ms after cue onset, and remained significant out to 200 ms. Benefits in response time for the valid cues relative to no-cue trials tended to increase with SOA in the present experiment, although at none of the SOAs was the magnitude of benefit significant. These patterns of costs and benefits as a function of SOA for static cues are generally consistent with the development of costs and benefits associated with abrupt-onset cues (Müller & Rabbitt, 1989; Remington, et al., 1992) and suggest that the time course of shifts elicited by these two different stimulus properties is similar.

In sum, the results of this experiment (a) provide further evidence that static discontinuities can elicit involuntary shifts of attention, (b) provide evidence that the contingencies for exogenous attention orientation are not necessarily set at the level of specific feature values, and (c) suggest that the time course of attention shifts to static discontinuities is similar to that of shifts to dynamic discontinuities.

General Discussion

The purpose of this series of experiments was to explore a new hypothesis concerning the conditions under which involuntary shifts of visual spatial attention occur. The contingent involuntary orienting hypothesis states that the involuntary orienting of attention to a stimulus event is contingent on whether that event shares a feature property that is critical to the performance of the target task. Consistent with this hypothesis, we have found conditions under which abrupt luminance changes over time (i.e., dynamic discontinuities) do not involuntarily summon spatial attention. Yantis and Jonides (1990) have previously shown that subjects can resist attention to an abrupt-onset stimulus when their attention is highly focused; the present results were obtained under conditions of spatial uncertainty, when spatial attention was not already focused. We have also shown that, contrary to the conclusions of previous research (e.g., Jonides & Yantis, 1988), conditions exist in which properties other than abrupt luminance changes (i.e., static discontinuities) do involuntarily summon attention. Thus, we have shown that the occurrence of involuntary attention shifts is systematically contingent on the relationship between the stimulus properties of the cue and the properties required to locate the target.

Before moving to the implications of these results for models of attention allocation, there are two potential qualifications to discuss regarding our conclusions. First, although it was reasonably assumed that the dimension defining static cues and targets in this series of experiments was a discontinuity in color, the luminance of the different colors used was not formally equated because of equipment limitations. Thus, although it seems unlikely that the observed contingency effects would be based on what are, at most, subjectively small differences in luminance, rather than the comparatively large differences in color, we cannot rule out this possibility with the present experiments.⁸ Whether the observed contingency effects were based on static discontinuities in color or luminance, however, is not central to our thesis. The important point is that we have established that covert involuntary attention shifts are contingent on task constraints.

Second, the distinction drawn between dynamic and static discontinuities, at this point, serves a primarily heuristic function. The terms as we use them are merely descriptive of the nature of the property that ultimately signals the location of the target or cue. Clearly, all of our displays contained both static and dynamic components. Color cues, for example, consisted of the abrupt onset of circles around each of the four boxes in addition to the discontinuity in color across circle sets. Thus, our strong claim that static discontinuities can capture attention must be qualified by the fact that all of our "static" displays contained "dynamic" components (i.e., abrupt onsets). In other words, although seemingly unlikely, we cannot rule out the logical possibility that the ability of static discontinuities to attract attention is dependent on the co-occurrence of an abrupt onset.

A Model of Exogenous Attention Control

The results of these experiments suggest a new view of the mechanisms underlying the involuntary orienting of spatial

attention. It had previously been suggested that exogenous, involuntary shifts of attention are the result of relatively inflexible, "hard-wired" mechanisms triggered by specific stimulus properties, namely abrupt luminance changes. The present results disconfirm this model. We propose that under conditions of spatial uncertainty, the exogenous allocation system can be "configured" or "set" to respond selectively to a property that signals the (unpredictable) location of stimuli that are relevant to optimal task performance. Thus, any particular system configuration, or "attentional control setting," is assumed to be a function of current behavioral goals. (In the absence of any current task goals, the system might default to settings based on long-term biases.) With a control setting established, events exhibiting the critical properties will involuntarily summon attention, whether or not the event is actually relevant to task performance. Stimuli not exhibiting these properties will not involuntarily summon attention.

One can draw an analogy between this type of attentional allocation system and a programmable interrupt in a digital computer. The range of external events that can cause interrupts is determined by the internal settings of the software, which, unlike "hard-wired" interrupt systems, can be varied by changing the code. Once the code is compiled and the program initiated, however, events that have these properties will be processed independently of the central processor (i.e., involuntarily). Similarly, in exogenous attention allocation, the control settings determining what properties will elicit shifts of attention can vary in accordance with task demands. Once the system is set and the processing sequence begun, however, the attentional response to properties compatible with current control settings is strongly involuntary.

This model can account for the present as well as previous spatial cuing data on exogenous orienting. According to the model, varying the nature of the property used to locate the target in the present experiments induced subjects to change their attentional control settings. Therefore, the occurrence of involuntary shifts of attention to a particular irrelevant cue was contingent on whether the system was currently set to respond to the properties exhibited by that cue. The apparent strongly involuntary shifts of attention to abrupt-luminance-change cues in previous spatial cuing studies (e.g., Müller & Rabbitt, 1989; Remington et al., 1992) can be accounted for by the fact that the occurrence and location of targets were signaled by a dynamic luminance change. As such, the exogenous allocation system would have been configured for luminance change, resulting in involuntary shifts to irrelevant, abruptly onset cues.

It is less clear how to interpret the results of the visual search studies (i.e., Jonides & Yantis, 1988; Miller, 1989; Yantis & Jonides, 1984) within the context of the model. Recall that in these studies, abrupt luminance changes, which were uncorrelated with target location, appeared to capture attention, whereas static discontinuities did not. With the spatial cuing paradigm used in the present experiments, spe-

⁸ We point out, however, that the conclusions drawn from previous experiments investigating the effects of static discontinuities in color on attention allocation have been similarly based on stimuli untested for equiluminance (Jonides & Yantis, 1988; Pashler, 1988).

cific control settings were induced by consistently signaling the target location with particular stimulus properties (i.e., a static or dynamic discontinuity). With the visual search studies just cited, however, it is difficult to infer what attentional control settings would have been appropriate because target location was not consistently signaled by any one property. Given our model, however, the results suggest that in all of these studies the system was set to respond to abrupt luminance changes. This could indicate that when there is little motivation to configure the system for any other property, abrupt luminance change (or dynamic discontinuities in general) may be instantiated as the “default” setting. Although we have shown that abrupt luminance changes have no special ability to elicit involuntary shifts of attention, we do not reject the possibility that the system may be biased toward dynamic discontinuities, given the important ecological information they may convey.

At What Level Is the System Configured?

The central tenet behind the model just outlined is that the exogenous allocation system can be configured to respond selectively to task-relevant stimulus properties. What are the functional stimulus properties for which the system can be configured? It is helpful to conceive of the set of possible properties as a hierarchy ranging from the very general to the very specific. At the most general level, the exogenous allocation system responds to discontinuities in preattentively processed visual information. As proposed earlier, these discontinuities can be divided into two basic types: dynamic discontinuities, which are defined by changes in properties over time, and static discontinuities, which are defined by changes in stimulus properties across space that do not vary over time. Each of these discontinuity types can, in turn, be defined in terms of any one of a set of stimulus dimensions. For example, a static discontinuity could result from a change in luminance, color, shape, orientation, binocular disparity, and so forth across a set of stimulus elements. Similarly, a dynamic discontinuity could result from a change across time in luminance, color, shape, orientation, spatial position (i.e., motion), and so forth. Each of these specific forms of discontinuity can further be defined in terms of particular features on the given dimension. For example, a static discontinuity in color might involve a red or green element and a dynamic discontinuity in luminance could be an onset or offset.

The broadest possible interpretation of our first three experiments is that the system can be configured at the level of static versus dynamic discontinuities. The results of Experiment 4 suggest that the system cannot be configured for specific feature values, at least within the color dimension. It is still unclear whether stimulus dimensions within the classes of static and dynamic discontinuities are functionally distinct with respect to configuring the exogenous allocation system. For example, would a static discontinuity defined by a difference in form elicit an involuntary shift of attention if a subject is monitoring for a target defined by a static color discontinuity? Similarly, would a cue defined by motion (a dynamic discontinuity in spatial position) capture attention when monitoring for an abrupt-onset target?

The results of Pashler (1988, Experiment 6), discussed earlier, suggest that the attentional control system may not, in fact, distinguish among feature dimensions within the class of static discontinuities. Recall that subjects were required to indicate which side of a textured display contained a form singleton (i.e., a single, static discontinuity in shape). Although irrelevant, random variation in color had no effect on performance, an irrelevant color singleton (i.e., a single, static discontinuity in color) significantly impaired performance on the task. This suggests that the exogenous system was configured at the level of static discontinuities in general rather than at the level of form discontinuity. Of course, the experiment was not designed to directly assess the occurrence of involuntary shifts of attention; a definitive conclusion would require methodology such as that used in the present experiments. Moreover, it is not clear that a similar pattern would generalize to other combinations of static or dynamic feature dimensions.

Relation to Existing Models of Attention Allocation

The model of exogenous attention allocation we have proposed is consistent with a number of recent computational models of attentional selection and allocation. These models provide specific mechanisms through which attentional controls could be set (e.g., Bundesen, 1990; Hoffman, 1979). Bundesen (1990), for example, has recently developed a theory in which the criteria for attentional selection (i.e., the various functionally distinct stimulus dimensions) are prioritized in accordance with current behavioral goals. The probability of selecting a given stimulus element is then a function of the weighted value of the element with respect to this dimensional prioritization. Similarly, Cave and Wolfe (1990), as well as others (e.g., Duncan & Humphreys, 1989; Koch & Ullman, 1985; Treisman & Sato, 1990; Treisman & Souther, 1985), have developed computational models in which attention is systematically allocated to the location of a display (or the information at that location) generating the highest activation in a “saliency” (Koch & Ullman, 1985), “attention” (Wolfe, Cave, & Franzel, 1989), or “discontinuity” (Treisman & Souther, 1985) map. The strength of the signals in these maps is a function of the “bottom-up” featural dissimilarity of the information at that location to the rest of the display and of “top-down” influence governed by current behavioral goals (Cave & Wolfe, 1990). Specifically, this top-down influence consists of increasing the relative magnitude of saliency signals from locations containing task-relevant stimulus properties.

One aspect of some of these computational models that is not consistent with a strong version of the model we propose is that the bottom-up saliency of an event, defined by its featural dissimilarity to the rest of the display, can directly influence the probability of that event eliciting a shift of attention, independent of whether the properties on which the saliency signal is based are task relevant or not. According to such a model, an event defined by a stimulus property for which the system is not configured could still elicit an involuntary shift of attention if its bottom-up saliency is high enough. Determining whether saliency does indeed play a role

in exogenous attention allocation or whether the top-down configuration of the system completely overrides any effects of salience will require further experiments in which both factors are systematically manipulated.

Endogenous-Exogenous Distinction

The results of this series of experiments suggest a refinement of the theoretical distinction between endogenous and exogenous control of attention allocation. Specifically, they suggest that the control of attention may be more consistently influenced by endogenous factors than previously thought. There may be no purely exogenous shifts of spatial attention; all shifts are mediated by "programmable" attentional control settings. The full range of influence of attentional control settings, however, remains to be seen. We cannot, at this time, rule out the possibility that there are circumstances under which stimulus salience can override internal control settings.

Although the present results complicate the distinction between exogenous and endogenous orienting, we maintain that there remains a critical difference in the role of external stimuli. Straightforward cases of endogenous control do occur in which higher mental processes direct spatial attention based on the interpretation of symbolic cues, statistical knowledge, or search strategies stored in memory, without any stimulus provocation. Moreover, the present experiments indicate that although exogenous shifts are modulated by endogenous factors, they are still driven by external stimuli; the attention shift is involuntary given the external stimuli and preestablished control settings. Finally, the distinction between exogenous and endogenous orienting is supported by growing evidence for corresponding differences between shifts generated by peripherally presented cues (typically exogenous) versus centrally presented symbolic cues (typically endogenous; see Henderson, 1991, for a review).

Although the hypothesis of contingent involuntary orienting was developed in the specific context of covert orienting of spatial attention, the concept may have important implications, at a more general level, for what it means to describe an act as "involuntary." Perhaps, in fact, all involuntary responses to stimuli have the potential to be modulated by programmable, internal control settings. These control settings, in turn, are a function of current behavioral goals, as well as past experience or enduring biases of the organism. Such a system represents a delicate, efficient balance between the rigidity necessary to ensure that potentially important environmental events do not go unprocessed and the flexibility to adapt to changing behavioral goals and circumstances.

References

- Banks, B., & Prinzmetal, W. (1976). Configurational effects in visual information processing. *Perception & Psychophysics*, *19*, 361-367.
- Beck, J. (1972). Similarity grouping and peripheral discriminability under uncertainty. *American Journal of Psychology*, *85*, 1-19.
- Beck, J. (1974). Relation between similarity grouping and peripheral discriminability. *Journal of Experimental Psychology*, *102*, 1145-1147.
- Beck, J. (1982). Textural segmentation. In J. Beck (Ed.), *Organization and representation in perception* (pp. 285-317). Hillsdale, NJ: Erlbaum.
- Bergen, J. R., & Julesz, B. (1983). Parallel versus serial processing in rapid pattern detection. *Nature*, *303*, 696-698.
- Broadbent, D. E. (1958). *Perception and communication*. Elmsford, NY: Pergamon Press.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, *97*, 523-547.
- Cave, K. R., & Wolfe, J. M. (1990). Modelling the role of parallel processing in visual search. *Cognitive Psychology*, *22*, 225-271.
- Colgate, R. L., Hoffman, J. E., & Eriksen, C. W. (1973). Selective encoding from multielement visual displays. *Perception & Psychophysics*, *14*, 217-224.
- Connor, J. M. (1972). Effects of increased processing load on parallel processing of visual displays. *Perception & Psychophysics*, *12*, 121-128.
- Donderi, D., & Zelner, D. (1969). Parallel processing in visual same-different decisions. *Perception & Psychophysics*, *5*, 197-200.
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*, 433-458.
- Egeth, H. (1977). Attention and preattention. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 7, pp. 277-320). San Diego, CA: Academic Press.
- Egeth, H., Jonides, J., & Wall, S. (1972). Parallel processing of multielement displays. *Cognitive Psychology*, *3*, 647-698.
- Folk, C. L., & Egeth, H. (1989). Does the identification of simple features require serial processing? *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 97-110.
- Hasher, L., & Zacks, R. T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General*, *108*, 356-388.
- Henderson, J. (1991). Stimulus discrimination following covert attentional orienting to an exogenous cue. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 91-106.
- Hoffman, J. E. (1979). A two-stage model of visual search. *Perception & Psychophysics*, *25*, 319-327.
- Jonides, J. (1980). Towards a model of the mind's eye's movement. *Canadian Journal of Psychology*, *34*, 103-112.
- Jonides, J. (1981). Voluntary vs. automatic control over the mind's eye's movement. In J. B. Long & A. D. Baddeley (Eds.), *Attention and performance IX* (pp. 187-203). Hillsdale, NJ: Erlbaum.
- Jonides, J., & Mack, R. (1984). On the cost and benefit of cost and benefit. *Psychological Bulletin*, *96*, 29-44.
- Jonides, J., & Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention. *Perception & Psychophysics*, *43*, 346-354.
- Julesz, B., & Bergen, J. R. (1983). Textons, the fundamental elements in preattentive vision and perception of textures. *The Bell System Technical Journal*, *62*, 1619-1645.
- Kahneman, D., & Henik, A. (1977). Effects of visual grouping on immediate recall and selective attention. In S. Dornic (Ed.), *Attention and performance VI* (pp. 307-332). Hillsdale, NJ: Erlbaum.
- Kahneman, D., & Henik, A. (1981). Perceptual organization and attention. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 181-211). Hillsdale, NJ: Erlbaum.
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Toward the underlying neural circuitry. *Human Neurobiology*, *4*, 219-227.
- Miller, J. (1989). The control of attention by abrupt visual onsets and offsets. *Perception & Psychophysics*, *45*, 567-571.
- Müller, H. J., & Findlay, J. M. (1987). Sensitivity and criterion effects in the spatial cuing of visual attention. *Perception & Psychophysics*, *42*, 383-399.
- Müller, H. J., & Rabbitt, P. M. A. (1989). Reflexive and voluntary orienting of visual attention: Time course of activation and resistance to interruption. *Journal of Experimental Psychology: Human*

- Perception and Performance*, 15, 315–330.
- Neisser, U. (1967). *Cognitive psychology*. New York: Appleton-Century-Crofts.
- Pashler, H. (1988). Cross-dimensional interaction and texture segregation. *Perception & Psychophysics*, 43, 307–318.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–25.
- Posner, M. I., & Boies, S. J. (1971). Components of attention. *Psychological Review*, 78, 391–408.
- Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and performance X* (pp. 531–556). Hillsdale, NJ: Erlbaum.
- Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition* (pp. 55–85). Hillsdale, NJ: Erlbaum.
- Remington, R. W., Johnston, J. C., & Yantis, S. (1992). Involuntary attentional capture by abrupt onsets. *Perception & Psychophysics*, 51, 279–290.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1–66.
- Shepard, M., & Müller, H. J. (1989). Movement versus focusing of visual attention. *Perception & Psychophysics*, 46, 146–154.
- Shulman, G. L., Remington, R. W., & McLean, J. P. (1979). Moving attention through visual space. *Journal of Experimental Psychology: Human Perception and Performance*, 5, 522–526.
- Theeuwes, J. (1991). Exogenous and endogenous control of attention: The effect of visual onsets and offsets. *Perception & Psychophysics*, 49, 83–90.
- Treisman, A., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95, 15–48.
- Treisman, A., & Sato, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 459–478.
- Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, 114, 285–310.
- Ullman, S. (1984). Visual routines. *Cognition*, 18, 97–159.
- Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 419–433.
- Yantis, S., & Jonides, J. (1984). Abrupt visual onsets and selective attention: Evidence from visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 601–621.
- Yantis, S., & Jonides, J. (1990). Abrupt visual onsets and selective attention: Voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 121–134.

Received September 6, 1990

Revision received December 10, 1991

Accepted December 20, 1991 ■