Journal of Experimental Psychology:
Human Perception and Performance
2003, Vol. 29, No. 5, 835–845

0096-1523/03/$12.00 DOI: 10.1037/0096-1523.29.5.835

Symbolic Control of Visual Attention: The Role of Working Memory and Attentional Control Settings

Jay Pratt
University of Toronto

Bernhard Hommel
University of Leiden

This study examined how 1 symbol is selected to control the allocation of attention when several symbols appear in the visual field. In Experiments 1–3, the critical target feature was color, and it was found that uninformative central arrows that matched the color of the target were selected and produced unintentional shifts of attention (i.e., involuntary, initiated slowly, producing long-lasting facilitatory effects). Experiment 4 tested whether such selection is the result of an attentional filter or of a competition bias due to a match of incoming information against integrated object representations stored in working memory. Here, the critical feature was shape and color was irrelevant, but matching color arrows were still selected. Thus, features of objects in working memory will bias the selection of symbols in the visual field, and such selected symbols are capable of producing unintentional shifts of attention.

On a nearly continuous basis, people are faced with complex visual scenes containing information that is critical to their goals and information that is irrelevant to those goals. Because people cannot simultaneously process all of the information in the visual field, it is crucial that they selectively attend to the goal-relevant information while ignoring goal-irrelevant information. This selection can be accomplished in many ways. In some cases, it proceeds in a bottom-up manner, such as when the abrupt onset of a new object or event in the visual field automatically captures attention. In some cases, the selection is accomplished in a top-down manner, with attention being volitionally allocated to specific locations in the visual field that have a high probability of yielding important information. In addition, experience often dictates where people should shift their attention in certain situations. For example, when a driver approaches a traffic intersection, experience should lead him or her to shift attention up- or rightward in order to determine the status of the traffic lights. In other situations, visual symbols are specifically placed in the visual field to aid people in determining where they should direct their attention. Indeed, through the communicative value of visual symbols, which can range from simple shapes to words, it is possible to indicate to others where attention should be allocated in certain situations. In this sense, many visual symbols can be considered "nothing more than a social convention by means of which persons who know the convention direct one another's attention to particular aspects of their shared world" (Tomasello & Call, 1997, p. 408).

The communicative value of symbols has long been known to visual attention researchers, and there is a considerable body of evidence to indicate that attention can be volitionally oriented to specific portions of the visual field in response to visual symbols. For example, arrows and numbers that indicate the most likely location of an upcoming target have often been used to produce shifts of attention (e.g., Jonides, 1981; Posner, 1980). In such experiments, shifts of attention are typically revealed by faster responses to targets that appear at the location indicated by the symbol (cued or valid location) than to targets that appear at other (uncued or invalid) locations. In addition, several studies have shown that some types of symbols can produce involuntary (or automatic) shifts of attention. Friesen and Kingstone (1998) used schematic faces in which the eyeballs suddenly changed from looking straight ahead to looking to one side or the other. They found that subsequent peripheral targets were detected faster on the side toward which the schematic eyeballs had looked, indicating that the observers had shifted their attention in response to the change of gaze by the schematic face. These findings were replicated with realistic faces (Driver et al., 1999; Langton & Bruce, 1999) and with pictorial changes in pointing gestures instead of gaze (Langton & Bruce, 2000). Moreover, very young infants (Hood, Willen, & Driver, 1998) and chimpanzees (Povinelli & Eddy, 1996) also tend to automatically attend to locations at which human faces are gazing. These studies strongly suggest that certain types of visual information can produce reflexive changes in the orienting of attention in observers. It is important to note, however, that hand pointing and eye movements have a strong ecological basis and in this sense are unlike symbols, which are ecologically neutral and acquire their communicative value through experience only.

Jay Pratt, Department of Psychology, University of Toronto, Toronto, Ontario, Canada; Bernhard Hommel, Department of Psychology, University of Leiden, Leiden, the Netherlands.

This research was supported by Natural Sciences and Engineering Research Council of Canada Grant 419206 to Jay Pratt. We thank Tara McAuley, Tijmen Moerland, Shalu Saini, and Menno van der Woude for their work in collecting the data. We are also grateful to Chip Folk, Brad Gibson, Yehoshua Tsal, and an anonymous reviewer for constructive, insightful comments and suggestions of alternative interpretations.

Correspondence concerning this article should be addressed to Jay Pratt, Department of Psychology, University of Toronto, 100 Saint George Street, Toronto, Ontario M5S 3G3, Canada. E-mail: pratt@psych.utoronto.ca

Recently, Hommel, Pratt, Colzato, and Godijn (2001) showed that involuntary shifts of attention are not limited to viewing hand or eye movements but also occur with simple, overlearned communicative symbols. Hommel et al. presented centrally located arrows and directional words (*up*, *down*, *left*, *right*) which were followed by a peripheral target. It is important to note that these symbols were entirely irrelevant to the task, and the observers were explicitly told that arrows and words contained no useful information about the location of the upcoming target. Nevertheless, in the first and third experiments, targets were detected faster when they occurred at the location indicated by the arrow or word (compatible condition) than at another location (incompatible condition). The second experiment extended this finding by showing that the compatibility effect also affected (i.e., counteracted) inhibition of return (IOR), a long-lasting attentional inhibitory effect. Finally, the fourth experiment showed that the compatibility effect was present even when the observers were told which location was the most likely target location. Thus, even in the presence of a volitional shift of attention to a specific location, the noninformative arrows and directional words still influenced the responses. Moreover, the compatibility effects occurred for targets at both the volitionally attended and the volitionally unattended locations. From these findings, Hommel et al. concluded that overlearned visual symbols can produce involuntary shifts of attention.

However, it seems unlikely that all visual symbols capable of producing automatic shifts of attention do so in all circumstances. Consider what would happen if one was to encounter a typical city intersection where two arrows (e.g., indicating a left turn lane and a straight-through lane) and two directional phrases (e.g., "merge left," "exit 200 m right") appear on the various signs in the visual field. Given the assumption that reflexive shifts of attention cannot be made in two or more directions simultaneously, where might attention be allocated? There is good evidence to suggest that the symbol that is relevant to the current goal will be selected, and this task-relevant symbol will then produce a reflexive shift of attention. The evidence for this hypothesis comes from two separate sources: research on the role of control settings in attentional capture and research on the interaction of working memory and the orienting of attention.

Folk and Remington, with their colleagues, have built an extensive program of research showing the importance of control settings on the involuntary allocation of attention (e.g., Folk, Remington, & Johnston, 1992, 1993; Folk, Remington, & Wright, 1994; Remington & Folk, 2001). Previous to their work, it was generally considered that the abrupt onset of a peripheral cue typically produced an automatic shift of attention. Folk et al. (1992, 1994) tested this notion using a display consisting of four peripheral placeholders surrounding a fixation placeholder in which one of two types of cues preceded one of two types of targets. Onset cues consisted of four white dots around one of the four peripheral placeholder boxes, whereas color cues consisted of four red dots around one placeholder box and four white dots around each of the other three peripheral boxes. In a similar manner, onset targets consisted of a single white stimulus in one of the peripheral boxes, and color targets consisted of a red stimulus in one of the boxes and white stimuli in each of the other three boxes. Folk et al.'s critical finding was that onset cues captured attention when the task was to identify the onset target but not when the task was to identify the color target, and vice versa. They proposed that the attentional control setting, based on the specific goals of the task, allows only cues that are consistent with the control setting to capture attention. Taken in the context of the present study, if an attentional control setting is present, then only the symbol that is consistent with the control setting should capture attention—an example of (goal-) "contingent automaticity" (Bargh, 1989; De Maeght, Hommel, & Schneider, 2003).

Recent evidence has shown that the allocation of attention is influenced by the contents of working memory in addition to attentional control settings. This evidence was obtained by Downing (2000), whose study was based on the notion that objects in the visual field compete for attention, and the strongest competitors will be objects that have already had their representations activated (cf. Desimone & Duncan, 1995; Duncan & Humphreys, 1989). At the beginning of each trial, Downing presented participants with a to-be-remembered item at fixation, followed by two simultaneously presented peripheral items (the memorized item and a distractor). Then, a probe appeared at one of the peripheral locations, requiring a speeded response based on the orientation of the probe. Following this probe response, working memory was tested by a same–different response. Downing hypothesized that a peripheral item that matched the item held in working memory would attract attention to that location and speed responses to probes that appeared at that location. Indeed, probes were reacted to faster if presented at the same peripheral location as the memorized item, indicating that attention had been attracted to that location. In addition, this advantage was not found when the memory test was removed from the task, suggesting that shifts of attention did not occur simply because the peripheral items physically matched the initially presented items. Rather, the results were consistent with the notion that objects that have active representations (i.e., are in the contents of working memory) tend to win the competition for allocation of attention.

In the context of the present study, these results suggest that the potency of symbols to evoke reflexive shifts of attention may vary with their relevance to the current task or their match with the contents of working memory (we return to this distinction in Experiment 4). Accordingly, we investigated whether the relative impact of competing symbolic stimuli (arrows, in this case) on attentional control depends on top-down selection processes and current task goals. The four experiments of the present study used similar designs that consisted of three major events. First, a stimulus appeared at the fixation point (called a *target cue*) and provided information about the type of stimulus that would serve as the target for that trial. The target cue could indicate the color or the shape of the forthcoming target. Then four differently colored arrows, originating from the fixation point and each pointing to a different peripheral location, were presented. The final stimulus appeared in one of the peripheral locations and either matched the target cue on the critical dimension (i.e., was a target and to be responded to) or did not match (i.e., was a distractor and not to be responded to). On the basis of our previous findings (Hommel et al., 2001), we assumed that all four arrows would compete for attentional control—that is, for inducing an automatic shift of attention. However, given our present considerations, we expected that the arrow that shared features with the target's representation (activated by the target cue) would tend to win this competition and determine which peripheral location attention would be oriented to. If this was the case, responses to the target should have been faster or more accurate if it appeared at this location and, hence, if the location of the target corresponded to that indicated by the arrow with which it shared features.

One final point, worth addressing before the details of the four experiments are presented, has to do with the traditional taxonomy of attention processes being either *automatic* or *controlled.* In the attentional-orienting literature, this can be seen as a distinction between *exogenous* and *endogenous* orienting (e.g., Müller & Rabbitt, 1989). Uninformative peripheral cues generate exogenous shifts of attention; such shifts are thought to be reflexive and quickly initiated, producing facilitatory effects at short stimulus onset asynchronies (SOAs) and IOR at long SOAs. Informative central cues generate endogenous shifts that are thought to be under some amount of volitional control, are slower to be initiated, and produce long-lasting facilitatory effects without any IOR. Perhaps a more descriptive definition would be that attention is "pulled" exogenously toward a peripheral cue and "pushed" endogenously in the direction indicated by a central cue. The uninformative central cues used by Hommel et al. (2001) do not fit easily into either category but, rather, have characteristics of both exogenous shifts (involuntarily initiated) and endogenous shifts (attention pushed in the direction of the arrow, slow but long-lasting facilitatory effects). As suggested by Yehoshua Tsal (personal communication, December 17, 2002), we refer to the shifts of attention produced by uninformative central cues as *unintentional* to avoid confusion with the existing terminology.

Experiment 1

In the first experiment, color was used as the task-relevant dimension. Accordingly, the target cue and the target were chosen from a set of four colors (blue, green, red, and white). Each of the four arrows that intervened between the target cue and target were also presented in one of the four colors. For example, a trial might consist of a blue target cue (indicating that a response should be made only if a blue target appears), followed by four arrows (one blue, one green, one red, and one white), followed by a blue target presented in the periphery. Because the blue arrow shares a task-relevant feature with the representation active in working memory, it should be selected and should control attentional orienting. If this is so, attention should be unintentionally oriented to the peripheral location at which the blue arrow points, resulting in faster responses to targets at that location than to targets at the locations pointed at by nonselected (i.e., other than blue) arrows. Thus, finding faster responses at the color-compatible locations (locations pointed at by arrows of target-matching color) implies the occurrence of two critical processes: the top-down selection of an arrow based on the task goals and the selected arrow's production of an unintentional shift of attention.

Method

Participants. Eighteen paid students from the University of Leiden, Leiden, the Netherlands, participated in this experiment. None were aware of the purpose of the experiment.

Apparatus. All of the experiments in this study were conducted on an IBM-compatible PC with a color monitor. Participants were tested individually in a darkened room, seated about 44–50 cm from the computer monitor.

Procedure. The trial sequence is shown in Figure 1. At the beginning of each trial, a yellow fixation dot (0.2° in diameter) and four yellow peripheral boxes (located above, below, right, and left of the fixation dot; each box subtending 1°) were presented on a black background. The boxes were 5.5° from the fixation dot (center to center). After 750 ms, the initial fixation dot was replaced with a target cue (a larger fixation dot, 0.4° in diameter, presented in one of four colors: blue, green, red, or white) for 500 ms. The target cue was then replaced with the small yellow fixation dot, which remained for the duration of the trial. After the fixation dot had reappeared for 500 ms, four arrows (each 1° in length and pointing toward one of the peripheral boxes) were presented, with each arrow appearing in a different color (blue, green, red, and white). Following a delay of either 300 ms or 1,000 ms from the onset of the arrows, a circular target (0.5° in diameter) appeared in one of four colors (blue, green, red, or white) at one of the four peripheral locations. The participants were instructed to (a) pay attention to the target cue, because that would indicate which color target would be responded to; (b) ignore the colored arrows, because they did not indicate the upcoming target location; (c) press the computer keyboard's spacebar as quickly as possible if the target was the same color as the large colored fixation dot; and (d) withhold response if the color of the target did not match the color of the target cue. The target was removed when the

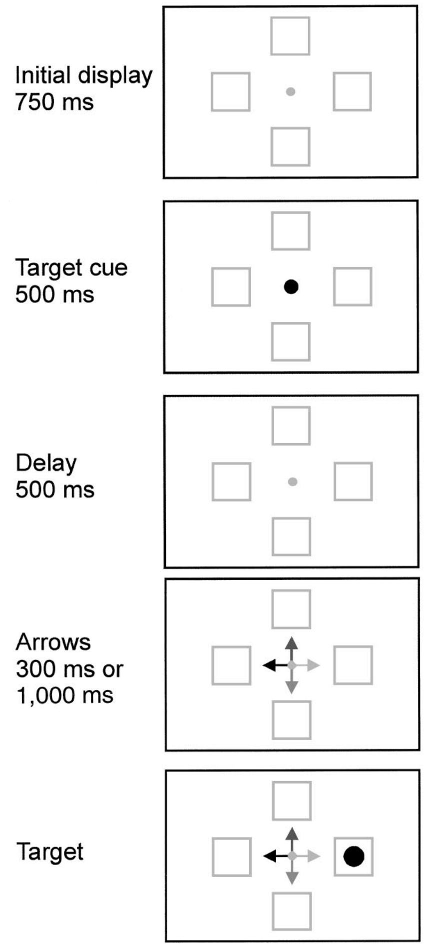

Figure 1. The trial sequence and display durations for Experiment 1. The fixation dot (gray circle) and the peripheral boxes were yellow. In the actual experiment, (a) the target cue (smaller black circle) was equally likely to be blue, green, red, or white; (b) the color of each arrow was equally likely to be blue, green, red, or white; and (c) the location of the target was equally likely at any of the four locations.

participant responded or 1,000 ms had elapsed. The intertrial interval was 1,000 ms.

Design. The color of the target cue (blue, green, red, or white), the direction in which each of the colored arrows (blue, green, red, and white) pointed, and the location of the target (up, down, left, or right) were randomized on each trial (i.e., the variables were fully crossed, then randomized). Across the experiment, the target cue and target were the same color on 75% of the trials. Each participant completed 480 trials, with a short break every 160 trials.

Results and Discussion

Before the reaction time (RT) data were analyzed, error trials were identified and removed. If the color of the target matched the target cue, errors occurred when RTs were less than 100 ms (anticipation responses) or greater than 1,000 ms (miss: no response to the target was made within the response window). If the color of the target did not match the target cue, errors occurred when RTs were less than 1,000 ms (false alarm: a response to a nontarget was made). The error rates for the three types of errors are shown in the top portion of Table 1.

With the error trials removed, mean RTs for respond trials (with same-colored target cue and target) were calculated per participant for compatible trials (in which the arrow of the same color as the target pointed toward the location of the target) and incompatible trials (in which the arrow of the same color as the target did not point toward the location of the target) for both the 300-ms and 1,000-ms SOAs (see Figure 2 for overall means). The mean RTs were analyzed with a 2 (arrow: compatible or incompatible) × 2 (SOA: 300 ms or 1,000 ms) analysis of variance (ANOVA). Main effects were found for SOA, $F(1, 17) = 13.7$, $MSE = 712.3$, $p < .003$ (longer RTs at the 300-ms SOA), and for arrow compatibility, $F(1, 17) = 89.3$, $MSE = 55.8$, $p < .00001$ (shorter RTs on compatible trials). There was no SOA × Arrow Compatibility interaction, $F(1, 17) < 1$.

The three types of errors (anticipations, misses, and false alarms) were analyzed separately. Because no anticipation errors were made, no statistical analysis was required. Miss rates were analyzed with a 2 (arrow) × 2 (SOA) ANOVA; no effects were found ($ps > .22$). False alarms were analyzed with a t test comparing error rates at the two SOAs (because no target was present in these conditions, there were no compatibility issues), and no difference in false alarms was found between the short and the long SOA ($t < 1$).

Table 1
Percentages of Errors by Stimulus Onset Asynchrony (SOA) and Cue–Target Compatibility in Experiments 1–4

Experiment and error	Compatible	Incompatible	Compatible	Incompatible
	300-ms SOA		1,000-ms SOA	
Experiment 1				
Anticipation	0	0	0	0
Miss	1.9	1.8	1.3	1.1
False alarm	3.6		4.6	
	300-ms SOA		1,000-ms SOA	
Experiment 2				
Anticipation	0	0	0	0
Miss	1.0	1.5	0.7	0.8
False alarm	2.4		2.4	
	50-ms SOA		300-ms SOA	
Experiment 3: Color target				
Anticipation	0	0	0	0
Miss	0	1.5	0.8	1.2
False alarm	1.0		0.5	
Experiment 3: *X* target				
Anticipation	0	0	0	0
Miss	0.8	1.5	0	0.8
False alarm	1.3		1.1	
	300-ms SOA		1,000-ms SOA	
Experiment 4				
Anticipation	0	0	0	0
Miss	2.5	3.2	3.0	2.6
False alarm	6.9		7.8	

Note. False-alarm error rates are not differentiated by compatibility. Anticipation = target present, reaction time (RT) < 100 ms; Miss = target present, RT > 1,000 ms; False alarm = target absent, RT < 1,000 ms.

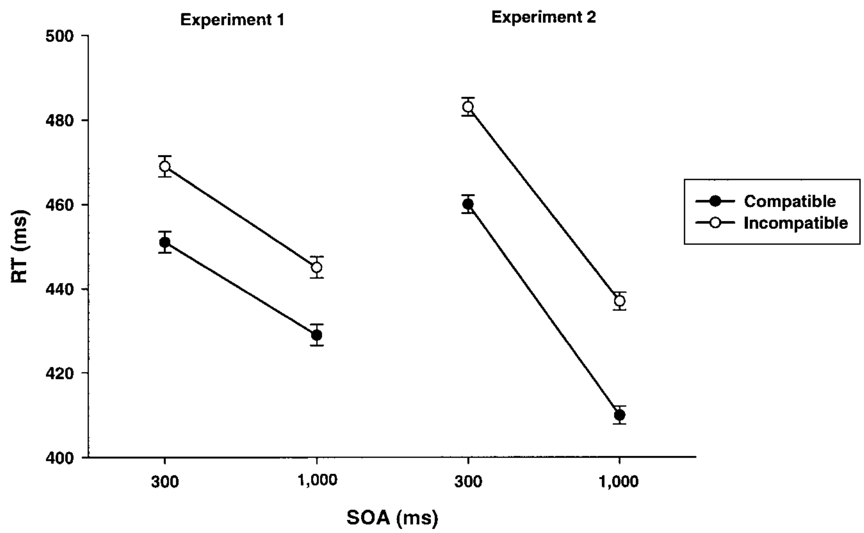

Figure 2. Mean reaction times (RTs) for trials in which the target cue and target were the same color for Experiments 1 and 2. Compatible trials occurred when a target appeared in a location pointed to by an arrow that matched its color. Incompatible trials occurred when the target appeared in a location pointed to by an arrow that did not match its color. Error bars, which represent within-subject standard error, were calculated as suggested by Loftus and Masson (1996) for within-subject designs. SOA = stimulus onset asynchrony.

As noted earlier, for faster responses to be found at the compatible location, the arrow matching the target cue color had to be selected, and the selected arrow had to produce an unintentional shift of attention to the indicated peripheral location. Both of these processes appear to have taken place because responses were faster to targets at compatible locations. Thus, the results from Experiment 1 suggest that the symbolic control of visual attention is indeed mediated by task goals. However, this conclusion is not watertight. Regarding the first process, our scenario assumed that a particular arrow was selected because it shared a task-relevant feature with the target and, hence, received top-down support. There is an alternative view.

An alternative explanation is that the compatibility effect may well reflect a task-independent bottom-up feature priming, rather than a top-down influence produced by the current goal (as suggested by our scenario). That is, selection of compatible targets may have occurred because the arrows were primed by the preceding, feature-overlapping target cue, which would imply a bottom-up process. To examine this possibility, we carried out the next experiment.

Experiment 2

In Experiment 1, the first event in the trial sequence was the appearance of the colored target cue, which was quickly followed by the four differently colored arrows. One of those arrows was the same color as the target cue and, hence, might have been perceptually primed by it. If this was the case, the selection of one arrow over the others (and the consequent attentional shift) would have reflected a bottom-up rather than a top-down process. In order to determine whether bottom-up processes might account for the selection of an arrow in this paradigm, the present experiment used the same basic sequence of events as did the first experiment, but now the target cue consisted of a word (*BLUE*, *GREEN*, *RED*, or *WHITE*) indicating which target was to be responded to. If the selection seen in the first experiment was due to bottom-up processes, no compatibility effects should have been found in the present experiment because the target cue and the arrows did not have any common physical features. However, if the selection was top-down, based on the task demands, then similar compatibility effects to those of the first experiment should have been observed.

Method

Participants. Fifteen students from the University of Toronto participated in this experiment in exchange for course credit. None were aware of the purpose of the experiment.

Apparatus and procedure. The apparatus was similar to that used in the first experiment. The procedure was the same as that used in the first experiment, except that a word (*BLUE*, *GREEN*, *RED*, or *WHITE*) served as the target cue instead of a colored dot. The word used as the target cue was presented in yellow and appeared for 1,000 ms.

Design. The color indicated by the target cue (blue, green, red, or white), the direction in which each of the colored arrows (blue, green, red, and white) pointed, and the location of the target (up, down, left, or right) were randomized on each trial. Across the experiment, the target cue and target were the same color on 75% of the trials. Each participant completed 480 trials, with a short break every 160 trials.

Results and Discussion

The data were treated as in Experiment 1. Error rates are shown in Table 1, and mean RTs are shown in the right panel of Figure 2. In RTs, main effects were found for SOA, $F(1, 14) = 53.9$, $MSE = 633.0$, $p < .0001$ (longer RTs at the 300-ms SOA), and for arrow compatibility, $F(1, 14) = 51.9$, $MSE = 171.1$, $p < .00001$ (shorter RTs on compatible trials). There was no SOA × Arrow Compatibility interaction, $F(1, 14) < 1$. Anticipation errors were absent, and neither miss rates ($ps > .15$) nor false alarms ($t < 1$) yielded reliable effects.

The results are very similar to those of Experiment 1; targets were responded to faster at locations pointed to by the arrows that were consistent with the task demands. Thus, as before, this indicates that one arrow was selected from the others, and this arrow produced an unintentional shift of attention to a peripheral location. As the present design prevented bottom-up perceptual priming of arrows by target cues, arrow selection must have been due to top-down processes—processes presumably governed by task demands.

Another important point is that we again found comparable compatibility effects at a short (300-ms) and a long (1,000-ms) SOA. The finding of this long-lasting facilitatory effect is consistent with the notion that certain uninformative central cues (e.g., arrows and directional words) produce unintentional shifts of attention. In other words, overlearned symbols can generate shifts of attention that share most of the characteristics of endogenous shifts but are initiated involuntarily.

Experiment 3

Although the results of Experiment 2 indicate that the compatibility effect is not due to perceptual priming between the target cue and the target, there remains the possibility of a kind of flanker or integration effect. Consider the situation of a red target, say, paired with a red arrow pointing toward the target location—hence, a compatible trial. According to our interpretation, maintaining information about the target color (i.e., red) in working memory primes the processing of the red arrow, which again pushes spatial attention toward the location that the arrow indicates. If this is so, the eventual target appears at a location that is attended already, which speeds up processing of and response to the target. However, faster processing of the target might also result from a kind of location-specific integration process. Assume, for the sake of the argument, that matching the target against the stored "search template" occurs on the basis of perceptual evidence sampled not only from the target location—an obvious assumption—but also from the target's vicinity. If so, an arrow (or any other stimulus) that appears in the same color as and in a location nearby the target might provide additional evidence for the presence of the target (i.e., increase the evidence counter for "red") and thereby speed up the decision process. Along these lines, it would be possible to account for the arrow-compatibility effects reported in Experiments 1 and 2 without the assumption that arrows actually push spatial attention in the direction that they indicate.

We admit that the plausibility of this alternative account is low. For one thing, it would not explain why Hommel et al. (2001) found benefits of compatibility between arrow direction and target location—and even between the meaning of directional words and target location—under conditions in which all arrows were centrally located and the overlap between arrows or directional words and the target was zero with regard to decision-related features. For another, one would expect that such a location-sensitive feature-sampling process would have a limited integration window that considers perceptual evidence from temporally close stimuli only. This does not seem to fit with our observation that the compatibility effect was obtained with an SOA of no less than 1 s and that the size of this effect was not even numerically smaller than it was with an SOA of 300 ms. But be this as it may, we thought that it would be helpful to provide more direct evidence that target-matching arrows do have an impact on the allocation of spatial attention. To do so, in Experiment 3 we complicated the task of our participants slightly. Not only were participants asked to carry out a go response to precued color targets, as in Experiments 1 and 2, but they were also to respond to a visual probe, a yellow *X*.

With regard to the precued color targets, we of course expected to replicate the findings obtained in Experiments 1 and 2. Results for the probe were particularly diagnostic, however. Along the lines of our involuntary attention-shift hypothesis, we expected that the results would mimic those obtained for color targets: faster responses when the probe location matches the direction of the arrow that shares the color of the precued target. In contrast, a location-specific integration account would predict compatibility effects only if the critical arrow shares a task-relevant feature with the target stimulus. This is the case with precued color targets, for which predictions from the two accounts do not differ, but it should not be the case for the probe, which should not benefit from any perceptual evidence that an arrow might provide. In other words, probe-related compatibility effects should be consistent with an attention-shifting account but not with a location-specific integration account.

A further interesting contrast refers to the SOA between arrow and target, which for diagnostic reasons we reduced to 50 and 300 ms. As carrying out an arrow-induced attention shift should take time, an attention-shifting account would lead one to expect that compatibility effects would be greater in magnitude at the long than at the short SOA. Reversely, an integration account suggests that the likelihood of integrating flanker- and target-related perceptual evidence decreases with increasing temporal distance between the two events, resulting in a prediction that compatibility effects (for color targets) should be smaller in magnitude at the long than at the short SOA.

Method

Participants. Eighteen students from the University of Toronto participated in this experiment in exchange for course credit. None were aware of the purpose of the experiment.

Apparatus and procedure. The apparatus was similar to that used in the first experiment. The procedure was the same as that used in the first experiment, save for three exceptions. The first exception was that on 20% of the trials, a yellow *X* (the probe) appeared in one of the four placeholder boxes instead of the color target. The participants were told to respond when the target matched the color target or when the target was a yellow *X*. The second exception was that the two SOAs used were 50 and 300 ms. The third exception was that a closed-circuit television camera was used to monitor eye position. When a shift in the point of gaze was detected, the experimenter issued a warning to the participant to remain fixated on the center of the display throughout each trial.

Design. The color of the target cue (blue, green, red, or white), the direction in which each of the colored arrows (blue, green, red, and white) pointed, and the location of the target (up, down, left, or right) were randomized on each trial. Across the experiment, the target cue and target were the same color on 55% of the trials, and the yellow *X* target appeared on 20% of the trials. Each participant completed 480 trials, with a short break every 160 trials.

Results and Discussion

Error rates are shown in Table 1, and mean RTs are shown in the left and middle panels of Figure 3. Very few fixation failures were noted (<0.5% of the trials). The mean RTs were analyzed with a 2 (arrow: compatible or incompatible) × 2 (target: color or *X*) × 2 (SOA: 50 ms or 300 ms) ANOVA. In RTs, main effects were found for SOA, $F(1, 17) = 418.5$, $MSE = 221.3$, $p < .00001$ (longer RTs at the 50-ms SOA); target, $F(1, 17) = 493.7$, $MSE = 1009.6$, $p < .00001$ (shorter RTs for color targets); and arrow compatibility, $F(1, 17) = 18.3$, $MSE = 427.6$, $p < .0008$ (shorter RTs on compatible trials). The target main effect indicates that the participants were searching for the color target and took longer to respond when the less likely *X* target appeared. It is important to note that there was an SOA × Arrow Compatibility interaction, $F(1, 17) = 6.2$, $MSE = 226.8$, $p < .03$, with larger compatibility effects at the 300-ms SOA. No other interactions reached significance (*ps* > .20). Anticipation errors were absent, and false alarms ($t < 1$) did not yield any reliable effect. However, there was a trend for compatibility in the misses ($p > .07$), but this was in the opposite direction of a speed–accuracy trade-off.

Two results from the present experiment indicate that shifting attention, not an integration effect, produced the compatibility effects noted in the previous two experiments. First, the pattern of responses to the *X* probe was almost identical to that for the color targets, despite the fact that the probe did not share any feature with the colored arrows. Second, the compatibility effect for both probe and color targets was larger at the 300-ms SOA than at the 50-ms SOA, a pattern of results that points to an internally generated shift of attention.

The small compatibility effects found at the short SOA may have been due to the involuntary nature of the selection process. Although a particular arrow is selected because it matches the task goal (i.e., finding a red object), this selection is performed involuntarily (as in the attention set notion put forward by Folk and Remington; see Folk et al., 1992, 1994). The involuntary nature of this selection may allow the shift in attention to begin earlier than is typically seen when an arrow at fixation indicates the probability of where a target will appear. It is clear, however, that the substantial RT benefit provided by shifting attention to a location in advance of a target was found only at the longer SOA. It is worth noting that the early occurring compatibility effect is unlikely to have been due to priming because it occurred for both the color targets and the probe.

The results of the previous experiments clearly implicate some type of top-down process as the mechanism for arrow selection, but they are silent as to how this process works in detail. What seems clear is that the process reflects the demands of the task. Simply put, if a person's goal is to find a blue object, blue objects will tend to be selected and to take control of spatial attention. Apparently, then, intending to perform a task somehow leads to the preparation of a "cognitive reflex" (Hommel, 2000) that enables the automatic processing of events matching the goal representation. Let us consider two similar but not entirely identical ways in which these cognitive reflexes might be enabled.

One way that the unintentional processing of attention-controlling objects might be implemented can be derived from the attentional control setting notion developed by Folk, Remington, and colleagues (e.g., Folk et al., 1992). As envisioned by these and other authors (e.g., Bundesen, 1990; Duncan & Humphreys, 1989; Wolfe, Cave, & Franzel, 1989), preparation for a specific task (e.g., searching for a red target) may include the implementation of attentional filters (or "attentional weights," in Bundesen's terminology). These filters check incoming information against task-specific selection criteria, such as the presence of a particular

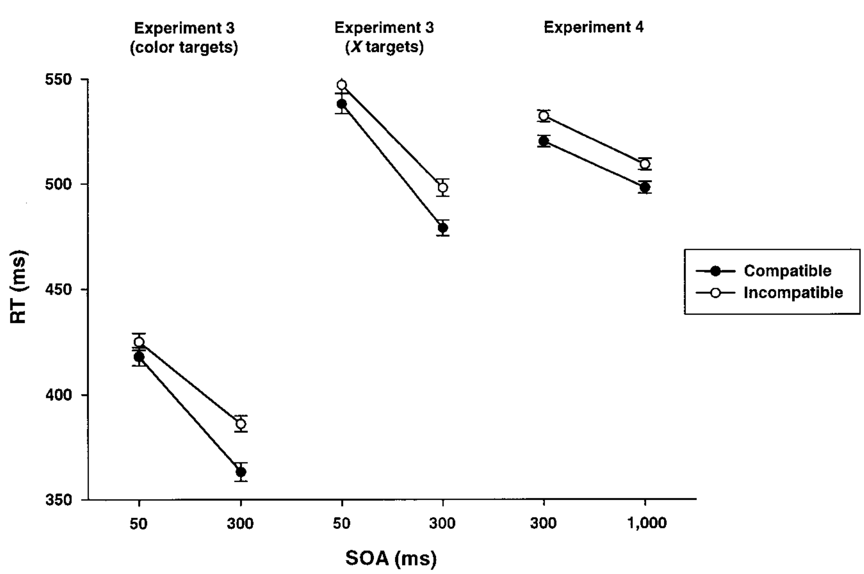

Figure 3. Mean reaction times (RTs) for trials in which the target cue and target were the same color for Experiments 3 and 4. Compatible trials occurred when a target appeared in a location pointed to by an arrow that matched its color. Incompatible trials occurred when the target appeared in a location pointed to by an arrow that did not match its color. Error bars, which represent within-subject standard error, were calculated as suggested by Loftus and Masson (1996) for within-subject designs. SOA = stimulus onset asynchrony.

feature. Only stimuli possessing this feature can pass and enter working memory, which is assumed to be a necessary precondition for these stimuli to affect information processing and overt behavior. According to this scenario, one would expect that only cues that share a critical feature with a target would be allowed through the filter to capture attention (see Figure 4A), so that cue processing is (conditionally) automatic and goal-directed at the same time. In fact, Remington and Folk (2001) have demonstrated that when a multidimensional object does not carry a defining property specified by top-down control settings, only the task-relevant dimension of an attended object is selected for deeper processing, while the task-irrelevant dimension of the same object is ignored. Obviously, the notion of an attentional filter makes for an appealing account of the findings from the previous experiments: Preparation for a target may have been accompanied by the implementation of a corresponding blue, green, red, or white filter that allowed the arrow of the same color to pass but not the other colored arrows. Once an arrow enters working memory, it can exert direct influence on spatial attention, thereby facilitating the processing of arrow-compatible targets. We call this the *attentional-filter hypothesis.*

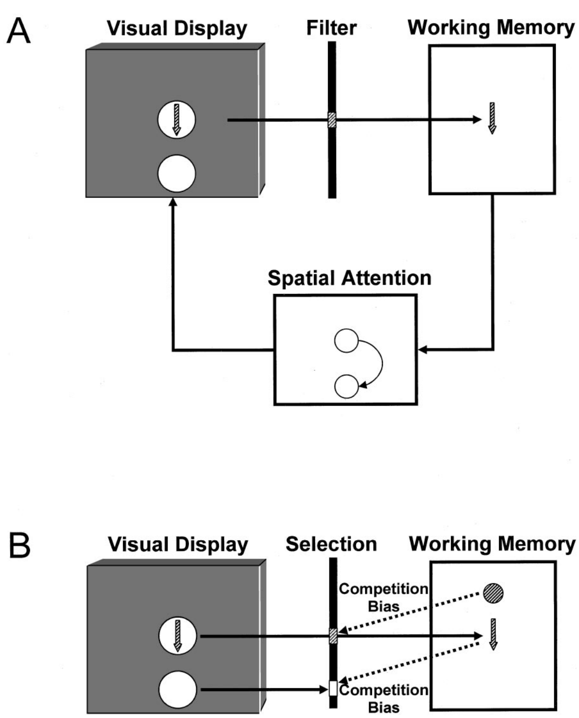

Figure 4. Alternative accounts of the arrow-compatibility effect and its dependency on task goals. A: An attentional-filter account, which assumes that preparing to react to a stimulus of a particular color (indicated by diagonal stripes) tunes a passive input filter in such a way that only stimuli possessing that color can pass and enter working memory. If, and only if, an arrow enters working memory will it impact the top-down control of spatial attention by pushing attention in the direction the arrow indicates. B: A memory-match account, which assumes a reciprocal relationship between the content of working memory and the criteria used to select incoming stimuli: Objects entering working memory bias the selection in favor of objects with which they share features. Storing the goal object (indicated by striped circle) in working memory favors the selection of the arrow that matches its color, which again favors the selection of objects in locations that match the arrow's direction.

However, the work of Downing (2000) is also consistent with a slightly different, more interactive scenario (that Downing himself apparently favors). As pointed out earlier, this work was motivated by the idea that competition in bottom-up information processing is biased by the contents of working memory. Along the lines of Folk et al. (1992), one might assume that this working-memory content consists of the task-relevant stimulus feature or feature dimension, so that incoming stimulus information will receive attentional support to the degree that it matches the represented goal attribute(s). Indeed, this is the interpretation commonly suggested by theorists who favor variants of the biased-competition approach (Bundesen, 1990; Desimone & Duncan, 1995; Duncan & Humphreys, 1989). However, in everyday life, goals are not usually defined in terms of features or feature dimensions but in terms of objects: People look for an apple, not for a conjunction of round shape and red color (see Duncan & Humphreys, 1989, for similar considerations). Accordingly, goal representations in working memory may refer to integrated objects, not just to the task-relevant feature(s) these objects possess, so that stimuli are selected because they match the representation of a memorized goal object, not (only) because they pass through a preinstalled relevance filter. Applied to our task, this interpretation suggests that preparing for a blue target, say, might be associated with maintaining a representation of this goal object in working memory. Maintaining this object might then "work back" on the input filter and favor any other stimulus matching it in whatever respect—that is, with respect to any matching feature (see Figure 4B). If so, when the arrows appear, it is only the blue arrow that matches the goal object stored in working memory and, hence, it will be the only arrow selected. As soon as the arrow enters working memory, it will support the selection of objects matching its direction, which again favors arrow-compatible targets. We call this the *memory-match hypothesis.*

These two accounts are not necessarily mutually exclusive but may refer to alternative ways to select information for action.[1] In many circumstances they will also make comparable predictions, especially if the structure of the experimental task enforces or strongly suggests a definition of goal objects in terms of features, such as in most visual-search tasks. The same is true for the results of Downing (2000): His participants may have stored the to-be-remembered objects in working memory, and these integrated object representations may have biased the system to process the same or similar objects. But it may also be that participants selected and stored only those features that they would later need to recognize the item in the memory test.

However, the two accounts do have contrasting implications for the irrelevant target dimensions. Consider the case in which participants are presented with a particular visual object and asked to react to one, and only one, feature that this object possesses (e.g., color). According to the attentional-filter account, a color filter will be established, and objects consisting of that color will be

[1] One might even be tempted to assume that these are just alternative descriptions of the same mechanism. For instance, it may be that attentional filters are implemented, controlled, and maintained by storing information about the critical dimension in working memory. However, recent observations that loading visual working memory with objects does not impair the search for visual features (Woodman, Vogel, & Luck, 2001) cast some doubts on this possibility.

selected. It is important to note that it should be only the task-relevant feature dimension that determines the selection of stimuli. But not so according to the memory-match account. This account implies that the presented target object is stored in working memory and that stimuli are selected if they match this maintained representation. However, this representation comprises both relevant and irrelevant features, so stimuli might be selected even if they match only an irrelevant feature of the goal object. We tested this unique prediction in our next experiment.

Experiment 4

Experiment 4 was designed to test the different implications of the two candidate selection mechanisms and, thus, to determine which one was responsible for the compatibility effects found in the previous experiments. To that end, the target cue in this experiment appeared in one of the four colors used in the experiment (as it did in Experiment 1), but it also appeared in one of two shapes (square or circle). The important feature of this experiment was that the shape of the cue indicated which stimuli should be responded to, but the color of the target cue was completely irrelevant to the task. As before, the target cue was followed by the four differently colored arrows, but now the target stimulus always appeared in yellow and its shape was either consistent with the shape of the target cue (i.e., respond) or inconsistent (i.e., do not respond).

If the top-down process is implemented as a passive attentional filter, the task should motivate participants to set the filter for a shape, making it insensitive to the task-irrelevant color. Accordingly, no arrow should be more likely to be selected—and, hence, no compatibility effects should be found—because the shape filter would treat all arrows equally. However, if the top-down processes reflect the direct impact of goal objects stored in working memory, all perceivable attributes of the target cue (e.g., shape, color, spatial position) should be more or less active regardless of their relevance to the task. Accordingly, arrows matching the color of the target object should receive top-down support and, thus, win the competition and determine where attention will be allocated. In other words, finding that arrows have a stronger impact on spatial attention (i.e., speed up reactions if the target appears where the arrow points) if their color matches that of the expected target would provide support for a memory-match account, whereas the absence of an arrow-compatibility effect would suggest an attentional-filter interpretation.

Method

Participants. Twenty students from the University of Toronto participated in this experiment in exchange for course credit. None were aware of the purpose of the experiment.

Apparatus and procedure. The apparatus was similar to that used in Experiment 1. The basic procedure was similar to that used in Experiment 1, but the target cue and target stimuli were different. In this experiment, the target cue appeared as either a colored (blue, green, red, or white) square or circle, and the target appeared as a yellow square or circle. In addition, the instructions were different from those of the previous experiments in that participants were instructed to pay attention to the shape of the target cue and to respond only when the shape of the target matched the shape of the target cue.

Design. The shape of the target cue (square or circle), the color of the target cue (blue, green, red, or white), the direction in which each of the colored arrows (blue, green, red, or white) pointed, and the location of the target (up, down, left, or right) were randomized on each trial. Across the experiment, the target was 66.6% likely to be of the same shape as that indicated by the target cue. Each participant completed 480 trials, with short breaks every 160 trials.

Results and Discussion

The data were treated as in Experiment 1. Error rates are shown in the bottom portion of Table 1 and mean RTs in the right panel of Figure 3. RTs revealed main effects for SOA, $F(1, 19) = 26.0$, $MSE = 431.4$, $p < .0002$ (longer RTs at the 300-ms SOA), and for arrow compatibility, $F(1, 19) = 5.1$, $MSE = 396.8$, $p < .035$ (shorter RTs on compatible trials), but no SOA × Arrow Compatibility interaction, $F(1, 19) < 1$. Again, anticipation errors were absent, and neither miss rates (ps $> .21$) nor false alarms ($t < 1$) yielded reliable effects.

The results show that, again, responses were faster to targets that appeared at locations pointed to by arrows that matched the color of the target cue. It is important to note that this was true despite the fact that the response was based on the shape of the target cue. Apparently, then, the task-irrelevant feature of the target cue biased the selection process in favor of one of the four arrows, indicating that selection was based on the integrated contents of working memory, as the memory-match hypothesis claims, and not (or not exclusively), on dimension-specific attentional control settings, as the attentional-filter hypothesis suggests.

General Discussion

There is strong evidence indicating that when people encounter an overlearned symbol in the visual field, such as an arrow or a directional word, that symbol produces an unintentional shift of attention (Hommel et al., 2001). In the present study, we addressed the question of what happens when people face more than one symbol at a time, a common situation in everyday life. Experiment 1 indicated that symbols are more likely to affect attentional control if they possess the task-relevant feature of an expected target. Thus, the unintentional effects that symbols have are not independent of the cognitively represented task goal, suggesting a sort of "conditional automaticity" in the sense meant by Bargh (1989). Experiment 2 strengthened this conclusion by ruling out an alternative account of the results of Experiment 1 in terms of perceptual priming—in fact, the compatibility effects were larger in Experiment 2 than they were in any of the other three experiments. Experiment 3 ruled out an account in terms of location-specific integration in providing direct evidence that selecting a particular arrow pushes spatial attention in the direction the arrow indicates. Finally, Experiment 4 examined in more detail whether the top-down support of goal-matching stimuli is best characterized as the impact of preestablished attentional filters or, rather, as the result of a competition bias due to a match of incoming information against integrated object representations stored in working memory. All in all, our observations are partly compatible with a filter approach, but a memory-match approach provides a more complete account of the data.

The filter approach that we considered represents an extension of Folk et al.'s (1992) concept of attentional control settings. Originally, this concept had been developed to account for distracting effects of task-irrelevant stimuli, such as spatial cues appearing with an abrupt onset (Posner, 1980) or visual oddballs in multistimulus displays (Theeuwes, 1994). These distractors attract attention to their location and thereby impair the processing of target information presented elsewhere. According to the control-setting approach, this is because participants in the particular experiments have specified their selection criteria in ways that are insufficient to filter out all distractors; for example, defining the to-be-attended target in terms of its visual oddity (e.g., instructions to select the deviant shape) may make it difficult to exclude oddball distractors (e.g., stimuli with a deviant color). Applying the control-setting approach to our findings, one can argue that our original observation—that task-irrelevant arrows and words exert an apparently direct influence on where attention is directed (Hommel et al., 2001)—reflected the intention of our participants to react to abrupt visual onsets. That is, participants may have defined "visual onset" as the task-relevant feature or "perceptual category" (Bundesen, 1990) so that any stimulus possessing this feature was selected.[2] The results of the present Experiments 1–3 are also consistent with a somewhat generalized version of the control-setting approach, according to which processing a stimulus event is "contingent on whether that event shares a feature property that is critical to the performance of the task at hand" (Folk et al., 1992, p. 1032). In our case, precuing a particular color—which made that color "critical to the performance"—led to the selection of the arrow sharing this color so that this arrow's direction could influence attentional control.

Regarding the outcome of Experiment 4, however, an account in terms of a passive attentional filter or attentional control setting is not entirely sufficient. Here, we varied targets and arrows on different feature dimensions with the arrow-related dimension being not "critical to the performance," so that a control-setting view would have predicted no effect of arrow–target matches. Because such an effect was observed, it is unlikely that stimulus selection was successfully restricted to the task-relevant stimulus dimension. Of course, this does not mean that the notion of attentional control settings is incorrect, especially in view of the fact that it does a good job in explaining the impact of distractors as such. It only means that this notion is unable to account for all aspects of our findings, so we need to consider additional factors. Indeed, at the moment, it is not at all clear how quickly attentional control settings can be implemented, how quickly they can be altered, or how long they last. It is not even certain that such control settings can be induced to occur on a trial-to-trial basis, because almost all experiments examining these settings induce a single control setting over a block of trials (e.g., Folk et al., 1992, 1994). Accordingly, one may speculate that control settings in the sense meant by Folk and colleagues establish only the general cognitive context for task-specific processing, whereas other, more flexible mechanisms take care of the details and more time-critical adaptations.

We next discuss the biased-competition approaches, à la Desimone and Duncan (1995) and Downing (2000), which attribute a great deal of attentional control to working memory and its content. As we pointed out earlier, these approaches assume that objects in the visual field compete for selection. Attention tends to be allocated to the winners of the competition, and the selected objects are given priority access to perception and action systems. On the one hand, biased-competition approaches are comparable to the control-setting account in their assumption that task goals have a strong impact not only on stimulus selection but also on whether and which distractors can interfere with performance. On the other hand, though, two of the implications of the biased-competition approaches distinguish them from the control-setting account. First, the biasing content of working memory is likely to consider, but is not restricted to, the task-relevant features of a goal-related object. Indeed, if what is stored in working memory is integrated object representations, not just individual features (e.g., Duncan & Humphreys, 1989), memorizing a goal object can be expected to support the processing of both stimuli overlapping with relevant features of the goal object and stimuli overlapping with its irrelevant features. This is exactly what we found in Experiment 4, in which the target cue's irrelevant color apparently was stored in working memory and, thus, produced top-down priming of same-colored arrows.

A second implication of the biased-competition approach that, in our view, distinguishes it from the control-setting account refers to the time scale on which these two mechanisms are likely to operate. Even if neither approach is yet well defined with regard to the temporal characteristics of the processes thought to be involved, we take attentional control settings to reflect the general task goals that people develop or are instructed to follow. If this is so, one may think that those settings are implemented before a task begins and then maintained more or less until it stops, which again implies that they refer to task-invariant characteristics only. In contrast, the concept of working memory refers to a system that is very flexible, of limited temporal capacity (if codes are not actively maintained), and updated constantly (e.g., Baddeley, 1986). Such a system is well suited to operate on a trial-to-trial basis and, hence, might be suspected to be involved when a participant prepares him- or herself to react to a precued target object. Accordingly, we tend to think that it was this preparation—and the fact that it referred to the whole target object, not just to its relevant features—that was responsible for the arrow-compatibility effect in Experiment 4 and, by inference, in Experiments 1–3 as well.

We feel that, taken altogether our findings provide strong support for goal-based models of attentional selection in general and for biased-competition models in particular. Minimally, we need to assume that people translate their goals into attentional control settings that specify which stimulus (and, perhaps, response) features are relevant (e.g., Folk et al., 1992). If, and inasmuch as, stimuli satisfy these specified criteria, they enter the competition for attentional resources and, as our arrow-compatibility effect suggests, for attentional control. In some cases, this type of selection may suffice, especially if there are only a few stimuli or the feature overlap between distractors and targets is low (see Duncan & Humphreys, 1989). However, if more than one stimulus enters the competition, more specific top-down constraints are needed. This, in our view, is where working memory comes into play. In most cases, working memory will contain more information about the sought-for target object than a single relevant feature, so that supporting competitors, according to their match with the content

[2] More speculatively, one may even argue that the differential effect sizes Hommel et al. (2001) observed were due to control settings: Symbols with spatial meaning (arrows and words) had a greater impact in tasks where the target was accompanied by distractors (Experiments 1A and 1B), hence, under conditions where target selection was more difficult and presumably was based on location information or stimulus–location conjunctions (Treisman, 1988; Wolfe, 1998). In other words, "spatial" distractors of any kind may generally be more potent if the task set is likely to consider spatial information.

of working memory, can fine-tune the selection process and settle into a solution. Once a competitor receives sufficient top-down support, such as a target-compatible arrow in our experiments, it will enter working memory and thereby acquire a status comparable to other working-memory contents. Hence, it now can also bias stimulus processing in a top-down fashion, just like the more “intentional” members of working memory. As our findings and those of Hommel et al. (2001) suggest, this bias involves spatial aspects of stimulus selection so that selecting a target stimulus in a location that matches the new, not so “intentional” working-memory member is facilitated. In other words, we think that symbols acquire the potency to control stimulus selection by entering working memory and that they do so by matching the stored target representation to at least some degree. This, in turn, suggests that examining the interaction between attentional selection and working memory processes will be a fruitful course for future research aimed at understanding the extensive communicative ability that is so vital to the success of our species.

References

Baddeley, A. D. (1986). *Working memory.* Oxford, England: Oxford University Press.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3–51). London: Guilford Press.

Bundesen, C. (1990). A theory of visual attention. *Psychological Review, 97,* 523–547.

De Maeght, S., Hommel, B., & Schneider, W. X. (2003). *The attentional demands of detecting a simple feature.* Manuscript submitted for publication.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18,* 193–222.

Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science, 11,* 467–473.

Driver, J., Davis, G., Ricciardelli, P., Kidd, P., Maxwell, E., & Baron-Cohen, S. (1999). Gaze perception triggers reflexive visuospatial orienting. *Visual Cognition, 6,* 509–540.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96,* 433–458.

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 1030–1044.

Folk, C. L., Remington, R. W., & Johnston, J. C. (1993). Contingent attentional capture: A reply to Yantis (1993). *Journal of Experimental Psychology: Human Perception and Performance, 19,* 682–685.

Folk, C. L., Remington, R. W., & Wright, J. H. (1994). The structure of attentional control: Contingent attentional capture by apparent motion, abrupt onset, and color. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 317–329.

Friesen, C. K., & Kingstone, A. (1998). The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic Bulletin & Review, 5,* 490–495.

Hommel, B. (2000). Intentional control of automatic stimulus–response translation. In Y. Rossetti & A. Revonsuo (Eds.), *Beyond dissociation: Interaction between dissociated implicit and explicit processing* (pp. 221–242). Amsterdam: John Benjamins.

Hommel, B., Pratt, J., Colzato, L., & Godijn, R. (2001). Symbolic control of visual attention. *Psychological Science, 12,* 360–365.

Hood, B. M., Willen, J. D., & Driver, J. (1998). Adult’s eyes trigger shifts of visual attention in human infants. *Psychological Science, 9,* 131–134.

Jonides, J. (1981). Voluntary versus 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.

Langton, S. R. H., & Bruce, V. (1999). Reflexive visual orienting in response to the social attention of others. *Visual Cognition, 6,* 541–567.

Langton, S. R. H., & Bruce, V. (2000). You must see the point: Automatic processing of cues to the direction of social attention. *Journal of Experimental Psychology: Human Perception and Performance, 26,* 747–757.

Loftus, G. R., & Masson, M. E. J. (1996). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review, 1,* 476–493.

Müller, H. J., & Rabbitt, P. M. (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.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32,* 3–25.

Povinelli, D. J., & Eddy, T. J. (1996). Chimpanzees: Joint visual attention. *Psychological Science, 7,* 129–135.

Remington, R. W., & Folk, C. L. (2001). A dissociation between attention and selection. *Psychological Science, 12,* 511–515.

Theeuwes, J. (1994). Stimulus-driven capture and attentional set: Selective search for color and visual abrupt onsets. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 799–806.

Tomasello, M., & Call, J. (1997). *Primate cognition.* New York: Oxford University Press.

Treisman, A. (1988). Features and objects: The Fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 40*(A), 201–237.

Wolfe, J. (1998). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13–73). Hove, England: Psychology Press.

Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model of visual search. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 419–443.

Woodman, G. F., Vogel, E. K., & Luck, S. J. (2001). Visual search remains efficient when visual working memory is full. *Psychological Science, 12,* 219–224.

Received March 5, 2002
Revision received March 2, 2003
Accepted March 6, 2003 ■