Annu. Rev. Psychol. 2001. 52:629–51

ATTENTION AND PERFORMANCE

Harold Pashler,[1] James C. Johnston,[2] and Eric Ruthruff[2]
[1]*Department of Psychology, University of California, San Diego, La Jolla, California 92093; e-mail: hpashler@ucsd.edu*
[2]*NASA Ames Research Center, Moffett Field, California 94035;*
e-mail: jcjohnston@mail.arc.nasa.gov, eruthruff@mail.arc.nasa.gov

Key Words dual-task interference, psychological refractory effect, task switching, attentional control, reflexive orienting

■ **Abstract** Recent progress in the study of attention and performance is discussed, focusing on the nature of attentional control and the effects of practice. Generally speaking, the effects of mental set are proving more pervasive than was previously suspected, whereas automaticity is proving less robust. Stimulus attributes (e.g. onsets, transients) thought to have a "wired-in" ability to capture attention automatically have been shown to capture attention only as a consequence of voluntarily adopted task sets. Recent research suggests that practice does not have as dramatic effects as is commonly believed. While it may turn out that some mental operations are automatized in the strongest sense, this may be uncommon. Recent work on task switching is also described; optimal engagement in a task set is proving to be intimately tied to learning operations triggered by the actual performance of a new task, not merely the anticipation of such performance.

CONTENTS

0066-4308/01/0201-0629$14.00

INTRODUCTION

The phrase "attention and performance" has come to refer to a venerable research tradition that began during World War II in the United Kingdom. Research within this tradition has sought to illuminate basic questions about the architecture of the human mind by examining the human performance in relatively simple tasks. Although some of the topics that once occupied attention and performance researchers are no longer studied very intensively (e.g. choice reaction time and matching), the core issues of attentional mechanisms, limits, and control have probably undergone more intensive study within the past five years than in any comparable period of time in the past. There has also been a great increase in efforts to relate attentional mechanisms characterized at a functional level to brain activity and neural circuits (Parasuraman 1998).

It is not possible, given space limitations, to do justice to all of the important results and ideas that have emerged from this research. This review has a much more modest aim: to review recent work as it bears on two important themes around which some notable progress has occurred recently. These themes would seem to have broad implications for psychologists beyond those primarily focusing on attention and performance. The two themes that will be discussed in this chapter are, first, new evidence painting a very different view of attentional control from what has previously been assumed, and second, new analyses of the effects of practice.

Both themes involve issues often discussed with reference to the concept of automaticity, which was widely discussed in the late 1970s and early 1980s. However, the new findings emerging in the past 5 or 10 years suggest quite a different picture from early formulations of automaticity (such as Schneider & Shiffrin 1977, Posner & Snyder 1975).

TOP-DOWN CONTROL AND BOTTOM-UP PROCESSING

Human behavior emerges from the interaction of the goals that people have and the stimuli that impinge on them. Behaviors are commonly viewed as lying along a continuum reflecting the relative influence of these two factors in their causation. At one extreme are reflexes, which occur in direct response to certain classes of stimuli unless actively inhibited (e.g. very hot temperatures on the fingers cause retraction of the hand). At the other extreme are voluntary behaviors such as speaking a certain word at a capricious moment of one's own choosing, which are not directly and reliably triggered by any particular incoming stimulus, as far as can be discerned. Most human behavior would seem to lie in between these two extremes, reflecting the joint impact of high-level goals (so-called top-down influences) and recent stimuli (so-called bottom-up influences). Teasing apart the principles that govern the interaction of top-down and bottom-up forces is critical to understanding any type of human behavior.

In the area of selective attention, the issue of top-down versus bottom-up control has played an important role in many influential theoretical formulations. Several theorists have hypothesized separate modes of attention control that correspond closely to the general categories of reflexive and voluntary behavior discussed above. Two influential examples are Posner's (1980) distinction between exogenous and endogenous attention control, and Jonides's (1981) distinction between automatic and nonautomatic attention control. Note that Posner's concept of exogenous control (explicitly modeled on the neurological concept of a reflex) and Jonides's concept of automatic control both suggest that attention control lies outside of the organism, namely in the stimuli impinging on the organism. If the appropriate stimulus arises, the response of attending to that stimulus will occur. Because of the key role attributed to the stimulus, exogenous control has often been characterized as stimulus-driven or bottom-up control.

In the other hypothesized mode of attention control (endogenous/nonautomatic/voluntary), control is thought to reflect the same sort of decisions that we make to initiate any other (overt) voluntary action. Just as I can decide to reach out with my hand to grasp a coffee cup, I can decide to move my spatial attention to the location of the coffee cup. Note that whereas a cognitive decision is hypothesized to be necessary for an endogenous attention shift, no particular stimulus event is necessary. Thus, whereas exogenous attention control is characterized as stimulus driven, endogenous control is typically characterized as cognitively driven. These suggestions echo those of early writers on attention. Titchener (1908), for example, remarked that any sudden change or movement, including a change in pitch, could distract a person who was trying to concentrate on something else (p. 192), and James (1890) made similar suggestions.

In order to empirically test predictions about voluntary attention, it is usually assumed that subjects will make voluntary decisions in accord with systematic rewards and punishments imposed by the experimenter. Thus, in practice it is expected that voluntary attention allocation will follow experimenter-imposed incentives. If, on the other hand, attention is systematically drawn[1] to objects when there is no incentive for such an allocation, that allocation of attention is involuntary (and hence by exclusion must be accomplished by the automatic/exogenous/reflexive control mode).

Although, in principle both hypothesized modes of attention control would appear equally deserving of study, in practice most research has focused on reflexive control of attention.[2] We review two recent lines of research. The first has attempted to characterize what critical stimulus properties allow an object to

---

[1] The use of the term drawn should not suggest that visual attention is a single spotlight which, when allocated to one stimulus, must depart from another (cf. Bichot et al 1999, Kramer & Hahn 1995).

[2] Why this should be so is not clear. Perhaps the category of reflexive control is more readily investigated because it emphasizes the role of the stimuli, whose nature and timing are readily manipulated by researchers.

capture attention reflexively. The second has examined whether reflexive attention is indeed involuntary in the way that is traditionally claimed.

What Stimulus Properties Reflexively Attract Attention?

Various theorists have offered related but somewhat different suggestions about what critical stimulus properties might reflexively draw attention. A number of researchers have hypothesized that changes in stimuli over time ("transients") are prone to capture visual attention, especially when these occur rapidly. Yantis & Jonides (1984, 1990) theorized that stimuli with abrupt onsets have a special propensity to attract attention. Their visual search experiments showed that abrupt onset targets were found unusually rapidly, as if attention were preferentially drawn to those items. Yantis & Hillstrom (1994) later refined this view, suggesting that what should be ecologically most useful is for attention to be drawn to new objects in the field because these may well present either an important threat to be avoided (like a predator) or an important opportunity to be sought out (like prey). Yantis & Hillstrom noted that in the earlier experiments of Yantis & Jonides the stimuli shown to especially attract attention not only had abrupt luminance onsets but also marked the appearance of new objects. To unconfound these properties, they created special stimuli that constituted new objects but did not appear abruptly (the stimuli were made visible by manipulation of texture, motion, or binocular disparity). Visual search was unusually fast for these new object items. In another experiment, Yantis & Hillstrom showed that stimulus objects that did not generate the percept of a new object, but did contain a strong luminance burst, yielded no significant speedup in search. Thus, they concluded that the key stimulus feature that captures attention is not an abrupt luminance change, as previously hypothesized, but rather the appearance of a new object.

Whereas Yantis and colleagues have emphasized the importance of temporal discontinuities in stimulus properties or objects present in a scene, others have emphasized the importance of discontinuities in the spatial distribution of stimulus properties. Some researchers have argued that attention is strongly attracted to "singletons," stimuli whose properties differ strongly from all of their immediate neighbors (Wolfe 1992, Wolfe et al 1989, Theeuwes 1991; but see Pashler 1988, Bacon & Egeth 1994). Other researchers have hypothesized that all stimuli have at least some tendency to attract attention. Stimuli differ in the strength of this tendency (usually dubbed salience), and those stimuli having the highest relative salience are most likely to attract attention (cf. Lee et al 1999, Theeuwes 1992, Braun 1999).

In spite of the heterogeneity of these hypotheses, they share an emphasis on the properties of stimuli (whether individually or in relation to other stimuli) that attract attention, without reference to any cognitive goals of the observer. These goals are presumed to play a role in the control of attention, but only when attention is controlled voluntarily rather than involuntarily (reflexively).

Is Reflexive Attention Control Involuntary?

One of the key properties attributed to the exogenous or reflexive mode of attention control is that it is involuntary. That is, attention capture depends only on the occurrence of the proper stimulus, not on the goals of the observer. Visual search times decrease with abrupt-onset targets and increase with abrupt-onset distractors even when, across trials, the location of abrupt onset stimuli is uncorrelated with the location of the target (e.g. Yantis & Jonides 1990, Mueller & Rabbitt 1989). Thus, it was argued that abrupt-onset stimuli can capture attention independently of the will of the subject. Remington et al (1992) pushed the argument for involuntary capture of attention further by using abrupt-onset distractors that were guaranteed to be in nontarget locations on all trials. With this arrangement, allocating attention to abrupt onsets could only hurt performance (rather than merely failing to help, on average, as with the uncorrelated-locations design[3]); hence, subjects should have had the strongest possible motivation to prevent attention from being captured by these distractors. Nevertheless, abrupt-onset distractors still impaired performance. This result provides strong support for involuntary attention capture by abrupt-onset stimuli.

By the early 1990s there appeared to be a broad consensus that the reflexive mode of attention control is free of influence from top-down cognitive processes. According to this view, attentional control is at the mercy of the appearance of certain classes of stimuli (be they transients, abrupt-onset stimuli, new objects, or more generally, relatively salient stimuli). Such stimuli would draw attention in a completely bottom-up manner, overriding or at least delaying the cognitively driven, top-down control of attention.

Yantis & Jonides (1990), however, noted one special case in which voluntary processes could prevent capture. If subjects know with certainty the location of a target and are given sufficient time to focus attention on that target, then distractor stimuli are at least momentarily unable to capture attention (see also Theeuwes 1991). This finding suggests that if an observer's attention is already locked onto one location, the power of other stimuli to draw attention is nullified. (It is unclear whether this nullification occurs because perceptual processing of the distractors is prevented, or due to a "lock-up" of the attention-shifting mechanism, which prevents a new shift.) In any event, this exception appears restricted to situations in which attention is tightly focussed on a object.

New View of Involuntary Attention Capture

To summarize, the allocation of spatial attention has commonly been claimed to operate through two modes. In the voluntary mode, high-level cognitive processes determine where attention is needed and shift attention accordingly in a

---

[3] Phenomena such as probability matching suggest that people may fail to optimize whenever consequences have some random variation from trial to trial.

top-down, cognitively driven manner. In the reflexive mode, stimuli that are high in absolute or relative attention-attracting properties cause an involuntary shift of attention. The reflexive mode, unlike the voluntary mode, was hypothesized to operate independently of the goals of the observer.

This consensus was challenged by Folk et al (1992), who argued that, paradoxically, even the involuntary capture of attention is subject to top-down cognitive control. Folk et al noted that previous studies of involuntary attention capture always used distractor stimuli sharing properties that defined the class of targets subjects were searching for. In particular, studies such as Remington et al (1992) that documented attention capture by abrupt-onset distractors had always used abrupt onset targets (inadvertently, one might say). Under the display conditions tested, it is thus conceivable that subjects were relying on the abrupt-onset property to help find targets. If so, the ability of abrupt-onset distractor stimuli to capture attention might have derived not from their stimulus properties per se, but from the match between their stimulus properties and the subject's voluntary set. What had not been studied were conditions in which distractor properties clearly failed to match the properties subjects used to find targets.

Folk et al (1992) set out to study attention capture with careful control of the relation between distractor properties and the properties used to find targets. They tested two unique distractor properties (color or abrupt onset) against each of two unique target-finding properties (again, color or abrupt onset). Confirming previous results, they found that when subjects are looking for abrupt onset targets, abrupt-onset distractors capture attention and color distractors do not. However, when subjects are looking for color targets (e.g. a red object in a background of green objects), abrupt-onset distractors do not capture attention, and color distractors (e.g. red distractors in a green background) do. Folk et al (1992) concluded that it is not the stimulus properties of distractors per se that determine whether attention capture will occur, but rather the relationship of distractor properties to the target-finding properties. They hypothesized that involuntary attention capture occurs if and only if distractors have a property that subjects are using to find targets.

To explain these results Folk et al (1992) proposed a new theory of attention capture. The theory describes the workings of attention as being analogous in some respects to that of a thermostat. Once set, a thermostat automatically activates the furnace when the temperature falls below a certain temperature without requiring further intervention by the person who set the thermostat. The person does, in one sense, control the thermostat, but in computer parlance the control is exercised off-line. That is, the person presets the critical temperature but has no role in the (on-line) events that trigger the furnace to turn on when the temperature falls below the preset value. Similarly, in the case of involuntary attention capture, Folk et al proposed that cognitive goals determine attentional control settings in advance (off-line); given the attention control settings in place at any particular time, the appearance of stimuli matching that setting will (on-line) capture attention with no further involvement of cognitive processes. Folk et al dubbed their hypothesis contingent orienting because the reflexive allocation of attention is contingent on the match between the stimulus properties of distractors and the attentional control settings.

Recent Findings on Contingent Orienting

Folk et al (1994) confirmed the generality of the claim that distractors draw attention involuntarily if and only if their distinctive property matches the property subjects are looking for. They tested various pairwise combinations of the three properties of apparent motion (present or absent), color, and abrupt onsets (present or absent). In each case, the principle was upheld. From the same perspective, Folk & Remington (1998) recently examined the claim by Yantis & Hillstrom (1994), described above, that attentional capture is triggered by the appearance of a new object (not necessarily marked by an abrupt onset). Folk & Remington found that the power of new-object distractors to capture attention was modified by the attentional control setting of the subject in the same manner as previously shown for abrupt onsets. Most importantly, they found that when observers were set for color targets, new-object distractors did not involuntarily capture attention.

The studies described thus far involve small effects of a single precue and relatively uncrowded displays. H Pashler (2000, submitted) attempted to magnify any effects of irrelevant transient distractors by presenting a very large number of them and presenting them concurrently with the target display. Pashler's task required observers to search for a particular red digit (e.g. "3") in a crowded display of 30 stationary red digits plus an additional 30 green distractors. To avoid introducing an incentive to search for onsets, observers previewed displays for a short while until a voice told them which digit to look for. The question asked was how performance would be affected if all of the green distractors were made to flash, twinkle, or "shimmy" rather than remaining stationary. One might expect the presence of all these potentially distracting stimulus changes to drastically increase search times, but the data showed that performance was little affected or actually improved. These results show a striking ability of top-down processes to tune selective attention so effectively that even large numbers of changing distractors will not capture attention, provided that their properties do not fit what the subject is looking for.

Recent research has shown that the "cognitive penetration" of involuntary attention capture by top-down processes is subtle and wide-reaching. Folk et al (1992) were able to experimentally demonstrate how the properties required to distinguish targets from distractors could alter control settings and thereby alter attention capture. They conjectured that control settings might also be tuned for other properties of targets, even those properties shared by distractors. For example, if the target display consists of an abrupt onset of targets (red) and abrupt-onset distractors (green), the most efficient strategy might be to first locate abrupt-onset stimuli and then search among them for red objects. Gibson & Kelso (1998) have since elegantly confirmed this conjecture. They showed that manipulations of the property used to signal the onset of the target field (shared by all stimuli in that field) did in fact produce corresponding changes in which properties showed attentional capture. For instance, when color as well as onsets signaled the target display, attentional capture was found for color stimuli as well as onset stimuli. This line of research has served to clarify several conflicting results regarding the relative ability of onsets and offsets to capture attention (e.g. Miller 1989, Martin-Emerson & Kramer 1997). Atchley et al (2000) found that they could produce attention capture by onsets, by offsets, or by either, simply by manipulating whether subjects needed to use onsets, offsets, or either to find candidate objects that might be targets.

We referred above to the common suggestion that objects that differed strongly from a relatively homogenous field of neighbors—so-called singletons—sometimes draw attention even when this is disadvantageous to performance of the task (Pashler 1988, Theeuwes 1992). These results, too, have also been shown to require re-interpretation. It is critical that in most experiments with distractor singletons, the subject is looking for unique targets that also constitute singletons in some different stimulus dimension (isolated elements different from surrounding stimuli). Hence, even when, for example, a color singleton captures attention when subjects are looking for a form singleton (Pashler 1988), it may be that subjects have a very general set for singletons. Bacon & Egeth (1994) confirmed this conjecture, showing that the ability of singletons to capture attention is dependent upon whether observers are or are not looking for target singletons. Gibson & Jiang (1998) have also reported that color singletons do not attract attention unless their relevance has become clear to the observer. An experiment by Jonides & Yantis (1988), using a different methodology (visual search), also supports this conclusion: Search was not slowed by the presence of certain properties in distractors that made them singletons, so long as the distribution of those properties was not correlated with the target location (ensuring subjects had no incentive to look for the distractor property at issue).

Yantis & Egeth (1999) have succeeded in closing an apparent "loophole" in the case for top-down control of involuntary attention capture. Is it not possible, they asked, that certain stimuli (e.g. feature singletons) do capture attention in a completely bottom-up, stimulus-driven fashion, but that this mechanism can be inhibited by top-down control? This hypothesis would still leave a core of attention-capture behavior that was not actually caused by top-down control. Yantis & Egeth (1999) set up an ingenious test of this hypothesis. If the key stimuli (e.g. feature singletons) that might capture attention were presented only rarely, it seems unlikely that capture would come to be inhibited. What Yantis & Egeth (1999) found, however, is that if feature singletons were presented only rarely, little or no capture occurred. The inhibition hypothesis is therefore wrong. In the absence of any top-down influence, feature singletons simply do not capture attention. Overall, as the consequences of the contingent capture hypothesis have been explored, further successes have been achieved and the generality of the model has been extended considerably.

In recent years there has also been progress clarifying what is actually happening when attention is captured by distractor stimuli. Much of our evidence comes from a cueing paradigm in which the distractors are presented for a short period (e.g. 50 ms), followed by a brief interval (e.g. 100 ms), followed by the target field (which might remain on until the response). The usual evidence for capture by a distractor is a delay—typically 40–50 ms or so—in the response to the target. This distractor-induced delay provides strong evidence that as a result of processing the distractor *something* happens to delay processing of the target. Because the distractor-induced delay only occurs if the distractor is in a different location than the target, it is natural to hypothesize that attention has been temporarily drawn to the distractor location. However, it is plausible that attention might only have been drawn *toward* the target without dwelling there; it is even possible that attention to the target was in some way impaired without the focus of attention being drawn in any particular direction. Remington et al (2000) showed that attention really *is* allocated to the cued location. They hypothesized that if attention capture really does involve a shift of attention to the distractor object, then that object should be processed to a deeper level. As a measure of distractor processing depth, they tested whether response-compatibility of the distractor would influence response time (RT) to the target. Whenever they found the traditional evidence for capture (a distractor-induced delay) they also found evidence for deep processing of distractors (a response compatibility effect); when traditional evidence for capture was lacking, there was no evidence for deep processing of distractors. Hence it appears likely that what we have been calling attention capture does indeed involve a brief allocation of attention to the distractor.

This research has—rather long after the introduction of the concept of attention capture—finally placed the metaphor on strong ground. Spatial attention is indeed allocated to distractor stimuli to which the subjects are not voluntarily trying to pay attention. What now seems fairly clear, however, is that though the ability of the distractor stimuli themselves to draw attention is unwanted in a local sense, it is properly understood as an inevitable by-product of what the subject *is* voluntarily trying to attend.

CONTROL OF PROCESSING: Task Switching and Cueing

Task Switching Effects

In the previous section we saw that the phenomenon of involuntary attention capture, often presumed to be entirely bottom-up and stimulus driven, turns out to be subtly influenced by top-down goals. We turn now to another area within the field of attention and performance that has undergone a remarkable resurgence of research in the past few years: task switching. Here, by contrast to attention capture, we will be discussing a phenomenon that appears at first glance to be entirely under top-down cognitive control, but that turns out to be heavily influenced by bottom-up, stimulus-driven processing.

In most studies of basic information-processing mechanisms, people are instructed to perform a particular task, and the focus of the researcher is on the events occurring between the presentation of an imperative stimulus and the occurrence of the response. The actions that a person undertakes in order to prepare to perform a task are not illuminated in such studies. Research on task switching, by contrast, sheds light on these preparatory activities.

The study of executive control of task-set was pioneered by the educational psychologist Arthur Jersild (1927). To measure the cost of switching task sets, Jersild had subjects perform various speeded tasks involving lists of items. Some lists ("pure lists") required the subject to perform the same task on each item in turn. Other lists ("alternating lists") required the subject to alternate tasks, e.g. performing task A on the first item, task B on the second, task A on the third, and so forth. Jersild found that when the two tasks involved a common stimulus set (e.g. addition versus subtraction of digits), subjects responded much more slowly on the alternating lists. The extra time required, often referred to as the switch cost, was very large by the standards of research on simple tasks (sometimes exceeding 500 msec/item). In the alternating-list condition, subjects not only had to keep close track of which task was due to be performed next (the stimulus itself was no help in this regard), but also needed to avoid any tendency to perform the (irrelevant) task-set used on the previous trial. In contrast, when the two tasks had distinct stimulus sets (e.g. adding digit pairs versus naming an antonym of a word) Jersild found no switch cost (or even a small switch benefit). In this condition the stimulus typically provides a strong task cue, and there is little tendency to perform the wrong task. This basic pattern of results has since been replicated using modern methods of timing individual responses (e.g. Allport et al 1994, Fagot 1994, Spector & Biederman 1976).

One question left unanswered by Jersild's work is whether the slower RTs sometimes observed in the alternating-task blocks are caused by task-switching per se, or rather by the "overhead" involved in keeping two task sets active at the same time. To address this question, Rogers & Monsell (1995) recently added a new twist to the task-switching paradigm. Instead of presenting the switch and nonswitch conditions in separate blocks of trials, they mixed these two conditions together within the same block of trials. In their "alternating-runs" paradigm, subjects performed tasks A and B in runs of two or more trials (e.g. AABBAABB). By comparing switch and nonswitch trials within the same block of trials, Rogers & Monsell were able to measure switch costs uncontaminated by unwanted differences between conditions in executive overhead, arousal, etc.

Using the alternating-runs paradigm, Rogers & Monsell (1995) verified Jersild's finding that switching between cognitive tasks (even very simple ones) can result in a substantial time cost (300 msec or more). However, the results differed from those of Jersild (1927) in that reliable switch costs were found even with distinct stimulus sets ($\sim$50 msec). This empirical discrepancy might be partly due to the fact that Rogers & Monsell presented individual stimuli, whereas Jersild presented lists of stimuli; Jersild's lists allowed for preview of upcoming items, which might have been more useful in the mixed list condition (see Spector & Biederman 1976).

What causes the switch cost? The most obvious explanation is that task switch trials require an active, top-down process that changes the current task set (also known as task-set reconfiguration). There are at least two lines of evidence in favor of this hypothesis. First, Rogers & Monsell (1995) found that switch costs can be reduced by allowing the subjects extra time to prepare for the upcoming task switch (see also Biederman 1973, Fagot 1994, Sudevan & Taylor 1987). Interestingly, preparation time reduced the switch cost only when this variable was manipulated between blocks rather than within blocks; it appears that subjects failed to reconfigure their task set when not guaranteed sufficient time to complete the process. This finding suggests that switch-cost reduction reflects an active (top-down) process, rather than a passive decay of interfering representations from the previous task (see Allport et al 1994). Further evidence for an active top-down reconfiguration process comes from Meiran (1996). By presenting subjects with a task cue either early or late within a fixed inter-trial interval, he was able to vary the available preparation time while holding constant the time elapsed since the previous trial (and hence the opportunity for passive decay). The results showed that the switch cost was indeed much smaller when subjects had more time to prepare for the upcoming task. This result confirms that the switch cost is reduced by an active top-down preparatory process (e.g. a task-set reconfiguration) during the inter-trial interval, rather than by a passive decay of interfering representations from the previous trial. If the top-down task-set reconfiguration were the only source of the switch cost, then providing a very long inter-trial interval should allow subjects to finish the reconfiguration in advance, completely eliminating the switch cost. This prediction, however, has been disconfirmed in several studies (Allport et al 1994, Fagot 1994, Rogers & Monsell 1995) that found residual switch costs at long inter-trial intervals. Responses to a task switch are especially slow when that response is incompatible with the response suggested by the irrelevant task (Meiran 1996, Rogers & Monsell 1995, Sudevan & Taylor 1987). Hence, the task set from the previous trial remains somewhat active even when subjects know it will not be needed on the upcoming trial and have plenty of time to suppress it.

In order to answer the question of why task switching hurts performance, it is helpful to reverse the question and ask why task repetition helps performance. From this perspective, residual switch costs occur because no amount of top-down task preparation can provide benefits as strong as the bottom-up effect of having just performed the task on the previous trial.

Explaining Residual Switch Cost

To explain the residual switch cost, Rogers & Monsell (1995) conjectured that top-down attempts at reconfiguring the task set are inherently incomplete. Specifically, they hypothesized that completion of the task-set reconfiguration on switch trials can only be triggered exogenously, by the appearance of the actual stimulus. Thus, no amount of preparation can eliminate the residual switch cost. According to this hypothesis, performing a task once leaves the system optimally prepared to perform that task again on the next trial. Therefore, when subjects are presented with alternating runs of four trials (e.g. AAAABBBB) the switch cost will be borne entirely by the first trial within a run of four. Rogers & Monsell have confirmed this prediction (however, see Salthouse et al 1998).

Rogers & Monsell's explanation of the residual cost offers one plausible analysis, but other alternatives deserve consideration. One attractive possibility is that certain critical mental processes (e.g. response selection) are performed more slowly on task switch trials than on task repetition trials. Allport et al (1994), for instance, proposed that performance on task switch trials is slowed owing to interference from the task set implemented on the previous trial. According to their "task-set inertia" model, switch costs are reduced over time (on the order of minutes) owing to passive decay. This hypothesis is difficult to reconcile with the evidence discussed above that the reduction in the switch cost during the inter-trial interval is due to an active (top-down) process (see Meiran 1996, Rogers & Monsell 1995). In addition, this hypothesis seems to predict that performance within a run of trials of one task should show gradual improvement over a period of several minutes. However, as noted above, Rogers & Monsell found that the switch cost was borne entirely by the first trial within the run.

Ruthruff et al (2000c) have also attempted to explain the residual switch cost in terms of a processing slowdown. These authors used a variant of the Rogers & Monsell (1995) alternating-runs paradigm, with long inter-trial intervals in which the presented task was not always the expected task. Thus, subjects usually received task sequences such as AABBAA, but occasionally received sequences such as AABBAB. This allowed Ruthruff et al to examine the interaction between the effects of task expectancy and the effects of task recency (i.e. the residual switch cost). These authors found additive effects between task expectancy and task recency in several experiments involving single-affordance conditions. Furthermore, two experimental factors thought to influence the duration of response selection (namely, stimulus repetition and the complexity of the stimulus-repetition mapping) had effects that were generally additive with task expectancy but overadditive[4] with task recency. This pattern of data is most easily reconciled with a model in which task recency influences the stage of response selection, whereas task expectancy influences some earlier stage (see Sternberg 1969). The inference that task recency influences response selection is incompatible with Rogers & Monsell's suggestion that the task-recency effect is caused by a task-set reconfiguration stage inserted on task switch trials. [One caveat, however, is that Ruthruff et al (2000c) studied only conditions in which the two tasks had distinct

---

[4] Two factors, A and B, are said to interact when the effects of factor A depend on the level of factor B (and vice versa). When the effects of factor A tend to be larger in the "difficult" levels of factor B than in the "easy" levels of factor B, the interaction is said to be overadditve.

stimulus sets, so it is not clear whether their conclusions apply equally well to conditions in which the two tasks share the same stimulus set.] Ruthruff et al did argue for the existence of a task-set reconfiguration stage, but they proposed that this stage is needed only when the task is unexpected, not when it is merely a task switch. Ruthruff et al went on to speculate that the role of top-down, goal-driven processes is to ensure that the proper task set is in place, but that these top-down processes have little or no direct, on-line influence over elementary cognitive processes.

Further work will be needed to determine which of the competing models of the residual switch cost are correct. In any case, the data at least permit us to draw several conclusions regarding executive control. First, switch costs clearly are much greater when the tasks being studied involve overlapping stimulus sets, so that the stimulus itself does not unambiguously indicate which task needs to be performed upon it. Second, deliberate (top-down) mechanisms can be successfully employed to ready critical mental machinery to carry out a particular task in the near future, substantially reducing reaction times to that task. Third, it appears that these top-down mechanisms alone cannot always achieve the full state of task readiness that can be achieved by actually having performed the task on the previous trial. Thus, there is a lingering bottom-up component to the switch cost, the exact cause of which is not yet known.

Taken together, this section and the previous section on attention capture show that recent research is making some progress toward unraveling the complex and subtle interplay of top-down, goal-directed and bottom-up, stimulus-driven processes in human performance. In retrospect, perhaps it should have been expected that the interplay would be rather complicated. The distinction between bottom-up and top-down control closely parallels the distinction in development between learned behaviors (controlled by the demands of the environment) and innate behaviors (controlled by an internal genetic program). Several decades ago it became clear that the development process typically involves an intricate and subtle interaction over time of environmental influences and genetic programming; few processes can be neatly classified as either learned or innate. An analogous picture is emerging from the areas of recent research reviewed here. It is possible that few, if any, classes of human behavior will turn out to be neatly classified as purely top-down or purely bottom-up.

EFFECTS OF PRACTICE

A number of influential theories proposed in the late 1970s suggested that tasks become “automatized” after they have enjoyed sufficient consistent practice (e.g. Schneider & Shiffrin 1977). Automatization is normally understood to mean that a task or operation acquires several properties. One feature is a lack of voluntary control—that the operation proceeds more or less reflexively given the appropriate stimulus input. Another is the lack of interference with other ongoing mental operations. Highly practiced tasks, it was said, can be performed simultaneously with other tasks without interference (except for structural conflicts such as common reliance on the same effectors, requiring foveation, etc.).

Effects of Practice on Discrete Dual-Task Performance

The psychological refractory period (PRP) effect refers to a ubiquitous and often large slowing effect observed when people try to perform two speeded tasks close together in time (each task involving a choice of responses based on a distinct stimulus). While responses to the first-presented stimulus are often little affected by temporal proximity to the second task, responses to the second stimulus are usually slowed as the interval between stimuli is reduced. The increase in second-task RTs as the interval between stimuli is reduced is the most commonly used measure of the magnitude of the PRP effect. Several decades of work have provided strong support for the view that the PRP effect arises from postponement of central processing stages in the second task—a processing bottleneck. On this account, central stages in task 2 cannot commence until corresponding stages of the first task have been completed, whereas perceptual and motoric stages in the two tasks can overlap without constraint. This "central bottleneck" theory (Welford 1967) is supported by a variety of evidence, including the existence of PRP slowing in tasks that do not share input or output modalities, and the results of chronometric studies in which different stages of the two tasks are manipulated (see Pashler 1994, for a review).

The PRP effect is a particularly clearcut example of true dual-task interference, and according to the central bottleneck theory, the interference is severe, in the sense that mental operations in one task are said completely to displace mental operations in the other task. Thus, the PRP effect provides a natural experimental design with which to ask whether practice produces automatization.

Does practice dramatically reduce the size of the PRP effect, and if so, does it completely circumvent the central bottleneck? Many early studies of practice in PRP designs concluded that although practice substantially reduces RTs, it has only a modest effect on the amount of PRP interference (e.g. Bertelson & Tisseyre 1969, Borger 1963, Dutta & Walker 1995, Karlin & Kestenbaum 1968). From these results it would appear that the cognitive limitations responsible for the PRP effect are robust across practice. There is reason to question the generality of these findings, however, owing to the specific task combinations employed in these experiments. Most used pairs of tasks that both required a manual response (perhaps due to the technical difficulties involved in automatically detecting and scoring vocal responses). Recent research suggests that it is difficult or even impossible to control the two hands independently (De Jong 1993; McLeod 1977, 1978). Hence, it is possible that subjects in these early practice studies did not have a fair chance to learn to perform both tasks at the same time.

To see if a dramatic reduction in the PRP effect is possible with distinct responses, Van Selst et al (1999) recently studied the effects of 36 practice sessions in a PRP design combining an auditory-vocal task (saying aloud whether a tone is high or low in pitch) with a visual-manual task (pressing a button to indicate the identity of a letter on the computer screen). Van Selst et al observed a rather dramatic reduction in the magnitude of dual-task interference with practice. The initial PRP effect (increase in second-task RT as the interval between the tone and the letter was reduced) was 353 msec—typical of PRP studies with unpracticed tasks—but after extensive practice the effect had shrunk to a mere 50 msec. (When the experiment was repeated with manual responses on both tasks, a much larger residual PRP effect remained after practice, confirming that the choice of response modalities can be important.)

Although practice dramatically reduced the PRP effect when input and output conflicts were minimized in this way, the effect had not entirely disappeared. Based on the effects of manipulating variables targeted to slow particular stages in the two tasks, Van Selst et al (1999) were able to show that a central bottleneck was still present after extensive practice. Thus, although practice reduced stage durations (as one would have expected), it did not allow stages that previously operated in series to begin operating in parallel. If the central bottleneck model holds both before and after practice, then the reduction in the PRP effect with practice reflects the speedup in the perceptual and central stages of task 1. Speedups in later stages of task 1, or in any stages of task 2, should not modulate the effect. To test this prediction, Ruthruff et al (2000a) had the highly trained subjects of Van Selst et al complete two additional transfer experiments. The first transfer experiment paired a new task 1 judgment with the old (highly practiced) task 2 judgment. The second transfer experiment paired the old (highly practiced) task 1 judgment with a new task 2. As predicted by the bottleneck model with stage-shortening, the PRP effect was relatively large in the first experiment (with a new task 1) but relatively small in the second experiment (with a new task 2).

The results of these transfer experiments argue against other interpretations of the effects of practice on PRP interference. For instance, it has often been supposed that when subjects practice doing two tasks at the same time, what they really learn is how to efficiently integrate this particular pair of tasks (e.g. Neisser 1976). If this account holds true in the present context, subjects should have produced relatively large PRP effects in both transfer experiments because they involved new pairs of tasks. Alternatively, if subjects had completely automatized the two tasks, so that they no longer required any limited resources, then subjects should produce small or nonexistent PRP effects in any transfer experiment involving either one of these highly practiced tasks. Contrary to this prediction, Van Selst et al (1999) found that the PRP effect shrunk if and only if the previously practiced task 1 was retained.

It should be noted, however, that while most of the results of Van Selst et al (1999) and Ruthruff et al (2000a) suggested that practice merely shortened the time required for central processing stages, there were persistent hints that practice may also have produced some small degree of true automatization on some trials, at least for some subjects.

Practice in Continuous Dual-Task Performance

So far we have discussed practice in the context of the discrete dual-task (PRP) design because this procedure allows one to test hypotheses about the effects of practice with exceptional clarity. However, the best known studies of practice effects have used continuous tasks and aggregate performance measures. A number of these studies have yielded extraordinary improvements in dual-task processing with practice. One commonly used approach has been to study the effects of practice in dual-task designs in which accuracy is the primary dependent measure. For example, Hirst et al (1980) found striking improvements in people's ability to read while taking dictation (though dual-task interference never completely disappeared). These results are sometimes taken to mean that even difficult mental operations (seemingly more sophisticated than any involved in the typical PRP task) became able to operate simultaneously and without interference. This conclusion is unwarranted, however, because such designs may be quite insensitive to processing delays. It is quite possible that critical mental operations on the two tasks never proceed at the same time, but subjects are nonetheless able to smoothly switch back and forth between the two tasks (see Broadbent 1982; McCann & Johnston 1992; Pashler 1990, 1998; Pashler & Johnston 1998). As an analogy, from the point of view of a user, time-sharing computers appear capable of performing two tasks at once without loss, even though the central processing unit never performs more than one operation at a time. In order to reliably detect capacity limitations, it is necessary to measure RT on individual tasks (as in the PRP design). Many other studies commonly cited as demonstrating that practice causes automatization (e.g. Schneider & Shiffrin 1977) involved tasks such as holding memory loads and visual search, which are not necessarily prone to queuing, even at the onset of practice.

Other investigators have studied the effects of practice on dual-task designs in which RT was the dependent measure, but all dual-task trials used the same stimulus-onset asynchrony (SOA) (as opposed to the typical PRP design, in which SOA is systematically varied within blocks). Schumacher et al (1997), for instance, trained subjects for 5 sessions in a dual-task design to compare single-task performance to dual-task performance at a 0 msec SOA (i.e. simultaneous stimulus presentation). One task was to say "one," "two," or "three," to low-, medium-, and high-pitched tones, respectively. The other task was to press one of three response keys that corresponded to the position of a disk on the computer screen. By the fifth session, subjects were able to perform these tasks together with no significant impairment in either RT or accuracy on the two tasks.

Schumacher et al's (1997) data, like that of Van Selst et al (1999) and Ruthruff et al (2000a), indicate that dual-task interference effects can be greatly reduced with extensive practice. Their data might also appear to indicate that practice in their design allowed subjects to bypass the central bottleneck. This might or might not be the case. One feature of their design (specifically, the use of simultaneous stimulus presentations) might have encouraged subjects to conjoin the two tasks into one. For instance, subjects may have learned to map each stimulus pair {S1, S2} onto the appropriate response pair {R1, R2}. A second limitation is that their visual-manual task involved a highly compatible mapping of stimuli onto responses, and thus might not require any central operations. It has been known for a long time that certain highly compatible tasks produce little or no PRP interference, perhaps because they do not require response selection or any other central processes (Greenwald 1972, Greenwald & Shulman 1973, Halliday et al 1959, Johnston & Delgado 1993, Pashler et al 1993). It will be interesting to see whether the results extend to more difficult tasks with neutral or incompatible response mappings.

Several recent studies using similar designs but less compatible response mappings have found sizable dual-task interference effects (e.g. Levy & Pashler 2000, Ruthruff et al 2000b). It is also not clear whether Schumacher et al's (1997) subjects had bypassed the bottleneck (i.e. central operations can operate in parallel with no interference) or whether there was what Van Selst et al (1999) call a latent bottleneck. A latent bottleneck occurs when the central operations on the two tasks cannot overlap, in principle, but due to the timing of the stimulus presentations they rarely "need to" occur simultaneously (i.e. central operations on one task generally finish before central operations on the other task are set to begin).

Hazeltine et al (2000) have found ways to address several of the concerns with Schumacher et al's (1997) study. They first replicated the basic finding of no-interference after practice. To evaluate the possibility that subjects converted the two tasks into a single conjoint task, Hazeltine et al trained subjects on only six of the nine possible stimulus pairs created by combining the three task 1 stimuli with the three task 2 stimuli. Then, in a subsequent testing phase, they presented the three stimulus combinations that had never appeared during the training phase, along with the six combinations that were trained. Subjects responded just as quickly to the new stimulus pairs as to the old stimulus pairs. This finding argues against the possibility that the learning gained in the training phase was specific to the particular stimulus combinations presented. However, this finding does not rule out the possibility that subjects learned a more abstract integration of the two tasks that was generalizable to all stimulus pairs, even those not yet experienced. To test this more general hypothesis, Hazeltine et al conducted a follow-up study in which they included −50 and +50 msec SOA conditions, in addition to the usual 0 msec SOA condition. The asynchronous presentations should force subjects to first wait (approximately 50 msec) for both stimuli to arrive before beginning the conjoint response selection. However, RTs were not slowed substantially in any of the SOA conditions. To address the concern that the visual-manual task was ideo-motor compatible, Hazeltine et al transferred subjects to a condition with an incompatible mapping of screen positions to fingers. This manipulation increased RTs to the visual-manual task, as expected, but had very little effect on the amount of dual-task interference.

The results of Hazeltine et al (2000) and Schumacher et al (1997), therefore, provide tantalizing evidence that practice allows subjects to bypass the central bottleneck under certain conditions. It remains to be seen whether this finding will generalize to other pairs of tasks and to conditions with large variations in SOA (as is the case in many real-world tasks).

CONCLUSION

Recent studies show that practice can dramatically reduce dual-task interference (Hazeltine et al 2000, Ruthruff et al 2000a, Schumacher et al 1997, Van Selst et al 1999). Thus, the very large interference effects observed with novel tasks (300–400 msec) might overestimate the amount of interference observed between pairs of highly practiced tasks in the real world. In many cases, however, it appears that practice merely reduces stage-durations, without allowing subjects to bypass the processing bottleneck (Ruthruff et al 2000a, Van Selst et al 1999). At the same time, some new evidence suggests that under some conditions practice can eliminate the processing bottleneck entirely (Hazeltine et al 2000, Schumacher et al 1997). Further work is needed to better define the boundary conditions for these two outcomes.

Locus of Practice Effects in Skill Learning

In the previous section, we discussed recent research on how practice affects the PRP. The key question was whether practice merely reduced the duration of component stages of the central bottleneck or allowed subjects to carry out more central operations in parallel on two tasks. Recent evidence suggests that the great bulk of the improvement with practice comes from reducing the duration of central stages, whereas relatively little of the improvement reflects being able to perform central operations on both tasks in parallel.

The question of how practice reduces RTs has been intensively studied in a separate but closely related research area, that of skill learning. Over the past two decades much interest has focussed on the so-called power law of learning. Rosenbloom & Newell (1987) proposed this law after observing that in a wide variety of RT tasks the shortening of RT with practice closely followed a power function. Numerous theories have been proposed to account for this empirical relation; one of the best known is the instance theory of Logan (1988). Logan theorized that reductions in RT with practice do not come about from improvements in the speed of executing a basic algorithm used by subjects at low practice levels. Rather, the reduction in RT with practice occurs because subjects employ an alternative strategy of retrieving memory traces left by previous performances (“instances”) and directly choosing the retrieved solution without need for calculation. Logan theorized that the algorithmic process and instance retrieval process operate in parallel in accordance with a “race” model—the first process to finish determines the RT. Because both of the processes (algorithm-based retrieval and instance retrieval) hypothesized to operate in parallel would seem to be examples of central processes in the sense discussed in the previous section, this theory appears inconsistent with the conclusions of PRP research (cf. Pashler & Johnston 1998). Thus, assessing the validity of this race model is relevant to the goal of providing a unified theory of information processing limitations.

An empirical challenge to Logan's race model of improvement with practice has emerged from studies of numerosity judgments about dot patterns. Palmeri (1997) reported that judgments in this paradigm improved rapidly with practice, closely following the power law, and data collected using this design initially appeared to be consistent with Logan's model. Rickard (1999), however, reanalyzed this data set and argued that although the data showed the usual pattern of a decline in RT over time with a positive second derivative (i.e. practice had diminishing returns over sessions of practice), careful examination showed systematic deviations from a perfect power law fit. Rickard argued that the data more closely fit predictions from his component power law theory. According to this theory, performance during the course of practice reflects a mixture of instance-based and rule-based performance. The essence of Rickard's model is that on any given trial, subjects perform only the rule-based or the exemplar-based strategy—but not both—and that, over practice, the mixture ratio becomes increasingly shifted toward the exemplar-based strategy. If Rickard's mixture model is correct, contrary to Palmeri's race model, the data can be explained without assuming that subjects can execute both strategies in parallel. Rickard's analysis would be entirely congenial to the findings of the PRP literature suggesting that only one central process can operate at a time.

Palmeri (1999), in a reply to Rickard, argued that more complex versions of the race model can fit the data, but the complexities needed to achieve such fits seem rather uninviting. They include an assumption that the algorithmic process can speed up with practice; it was an important motivation for the Logan model to avoid this assumption. The matter is presently not resolved, and it will be very intriguing if results from future skill-learning studies support or overturn the idea that multiple memory retrieval processes generally cannot, or at least do not, operate in parallel.

CONCLUDING COMMENTS

We have surveyed several particularly active research areas within the scope of the attention and performance tradition, focusing on studies of attention capture and studies of the limits of practice effect. One cannot examine this literature without being struck that the traditional focus on distinguishing parallel and serial processing remains very much a focus of research, despite skepticism expressed from time to time (e.g. Newell 1973). It is certainly the case that serial and parallel models may mimic each other in certain contexts, particularly with respect to patterns of mean response times (Townsend 1974). However, a variety of techniques now exist for successfully distinguishing these types of processes, and the question of whether different sorts of mental events can or cannot operate simultaneously is so fundamental that it simply cannot be evaded. Indeed, it may turn out that behavioral data can provide us with better answers to questions about the time course of processing than it can on questions of mental representation or underlying computational architecture.

At a more substantive level, one of the main themes to emerge from the research surveyed above is the idea that the effects of mental set are more pervasive than had previously been thought. As described in the first sections of this chapter, a variety of proposals for "wired-in" attention capture by particular stimulus attributes have been effectively challenged; attention, it turns out, is subject to a far greater degree of top-down control than was suspected 10 years ago. A second broad conclusion is that the effects of practice appear much more circumscribed than was previously appreciated. Although it may turn out that some mental operations are truly subject to automatization in the strongest sense, it appears quite unlikely that this holds nearly as broadly as was once suspected.

LITERATURE CITED

Allport A, Styles EA, Hsieh S. 1994. Shifting intentional set: exploring the dynamic control of tasks. In *Attention and Performance XV: Conscious and Nonconscious Information Processing*, ed. C Umilta, M Moscovitch, pp. 421–52. Cambridge: MIT Press

Atchley P, Kramer AF, Hollstrom AP. 2000. Contingent capture for onsets and offsets: attentional set for perceptual transients. *J. Exp. Psychol.: Hum. Percept. Perform.* 26:594–606

Bacon WF, Egeth HE. 1994. Overriding stimulus-driven attentional capture. *Percept. Psychophys.* 55:485–96

Bertelson P, Tisseyre F. 1969. Refractory period of c-reactions. *J. Exp. Psychol.* 79:122–28

Bichot N, Cave K, Pashler H. 1999. Visual selection mediated by location: feature-based selection of noncontiguous locations. *Percept. Psychophys.* 61:403–23

Biederman I. 1973. Mental set and mental arithmetic. *Mem. Cogn.* 1:383–86

Borger R. 1963. The refractory period and serial choice reactions. *Q. J. Exp. Psychol.* 15:1–12

Braun J. 1999. On the detection of salient contours. *Spat. Vis.* 12:211–25

Broadbent DE. 1982. Task combination and the selective intake of information. *Acta Psychol.* 50:253–90

De Jong R. 1993. Multiple bottlenecks in overlapping task performance. *J. Exp. Psychol.: Hum. Percept. Perform.* 19:965–89

Dutta A, Walker BN. 1995. *Persistence of the PRP effect: evaluating the response-selection bottleneck*. Presented at 36th Annu. Meet. Psychonom. Soc., Los Angeles

Fagot C. 1994. *Chronometric investigations of task switching*. PhD thesis. Univ. Calif., San Diego

Folk CL, Remington RW. 1998. Selectivity in distraction by irrelevant featural singletons: evidence for two forms of attentional capture. *J. Exp. Psychol.: Hum. Percept. Perform.* 24:847–58

Folk CL, Remington RW, Johnston JC. 1992. Involuntary covert orienting is contingent on attentional control settings. *J. Exp. Psychol.: Hum. Percept. Perform.* 18:1030–44

Folk CL, Remington RW, Wright JH. 1994. The structure of attentional control: contingent attentional capture by apparent motion, abrupt onset, and color. *J. Exp. Psychol.: Hum. Percept. Perform.* 20:317–29

Gibson BS, Jiang Y. 1998. Surprise! An unexpected color singleton does not capture attention in visual search. *Psychol. Sci.* 9:176–82

Gibson BS, Kelsey EM. 1998. Stimulus-driven attentional capture is contingent on attentional set for displaywide visual features. *J. Exp. Psychol.: Hum. Percept. Perform.* 24:699–706

Greenwald A. 1972. On doing two things at once: time sharing as a function of ideomotor compatibility. *J. Exp. Psychol.* 94:52–57

Greenwald A, Shulman HG. 1973. On doing two things at once: II. Elimination of the psychological refractory period effect. *J. Exp. Psychol.* 101:70–76

Halliday AM, Kerr M, Elithorn A. 1959. Grouping of stimuli and apparent exceptions to the psychological refractory period. *Q. J. Exp. Psychol.* 11:72–89

Hazeltine E, Teague D, Ivry R. 2000. *Dual-task performance during simultaneous execution: evidence for concurrent response selection processes*. Presented at Annu. Meet. Cogn. Neurosci. Soc., San Francisco

Hirst W, Spelke ES, Reaves CC, Caharack G, Neisser U. 1980. Dividing attention without alternation or automaticity. *J. Exp. Psychol. Gen.* 109:98–117

James W. 1890/1950. *The Principles of Psychology,* Vol. 1. New York: Dover

Jersild AT. 1927. Mental set and shift. *Arch. Psychol.* 89:whole issue

Johnston JC, Delgado DF. 1993. *Bypassing the single-channel bottleneck in dual-task performance*. Presented at 34th Annu. Meet. Psychonom. Soc., Washington, DC

Jonides J. 1981. Voluntary versus automatic movements of the mind's eye. In *Attention and Performance IX*, ed. J Long, A Baddeley pp. 197–203. Hillsdale, NJ: Lawrence Erlbaum

Jonides J, Yantis S. 1988. Uniqueness of abrupt visual onset in capturing attention. *Percept. Psychophys.* 43:346–54

Karlin L, Kestenbaum R. 1968. Effect of number of alternative on the psychological refractory period. *Q. J. Exp. Psychol.* 20: 167–78

Kramer AF, Hahn S. 1995. Splitting the beam: distribution of attention over noncontiguous regions of the visual field. *Psychol. Sci.* 6:381–86

Lee DK, Itti L, Koch C, Braun J. 1999. Attention activates winner-take-all competition among visual filters. *Nat. Neurosci.* 2:375–81

Levy J, Pashler HE. 2000. Is dual task slowing instruction dependent? Accepted pending revision. *J. Exp. Psychol.: Hum. Percept. Perform.* In press

Logan GD. 1988. Toward an instance theory of automatization. *Psychol. Rev.* 95:492–527

Martin-Emerson R, Kramer AF. 1997. Offset transients modulate attentional capture by sudden onsets. *Percept. Psychophys.* 59: 739–51

McCann RS, Johnston JC. 1992. Locus of the single-channel bottleneck in dual-task interference. *J. Exp. Psychol.: Hum. Percept. Perform.* 18:471–84

McLeod P. 1977. A dual task response modality effect: support for multiprocessor models of attention. *Q. J. Exp. Psychol.* 29:651–67

McLeod P. 1978. Does probe RT measure central processing demand? *Q. J. Exp. Psychol.* 30:83–89

Meiran N. 1996. Reconfiguration of processing mode prior to task performance. *J. Exp. Psychol.: Learn. Mem. Cogn.* 22:1423–42

Miller JO. 1989. The control of attention by abrupt visual onsets and offsets. *Percept. Psychophys.* 45:567–71

Mueller HJ, Rabbitt PM. 1989. Reflexive and voluntary orienting of visual attention: Time course of activation and resistance to interruption. *J. Exp. Psychol.: Hum. Percept. Perform.* 15:315–30

Neisser U. 1976. Cognition and reality: principles and implications of cognitive psychology. San Francisco: Freeman

Newell A. 1973. You can't play twenty questions with nature and win: projective comments on the papers of this symposium. In *Visual Information Processing*, ed. WG Chase, pp. 283–308. New York: Academic

Palmeri TJ. 1997. Exemplar similarity and the development of automaticity. *J. Exp. Psychol.: Learn. Mem. Cogn.* 23:324–54

Palmeri TJ. 1999. Learning categories at different hierarchical levels: a comparison of category learning models. *Psychonom. Bull. Rev.* 6:495–503

Parasuraman R. 1998. *The Attentive Brain*. Cambridge, MA: MIT Press

Pashler H. 1988. Cross-dimensional interaction and texture segregation. *Percept. Psychophys.* 43:307–18

Pashler H. 1990. Do response modality effects support multiprocessor models of divided attention. *J. Exp. Psychol.: Hum. Percept. Perform.* 16:826–42

Pashler H. 1994. Dual-task interference in simple tasks: data and theory. *Psychol. Bull.* 16: 220–44

Pashler H. 1998. *The Psychology of Attention*. Cambridge, MA: MIT Press

Pashler H, Carrier M, Hoffman J. 1993. Saccadic eye movements and dual-task interference. *Q. J. Exp. Psychol. A* 46:51–82

Pashler H, Johnston J. 1998. Attentional limitations in dual-task performance. In *Attention*, ed. HE Pashler, pp. 155–89. Hove, UK: Psychology Press/Erlbaum

Posner MI. 1980. Orienting of attention. *Q. J. Exp. Psychol.* 32:3–25

Posner MI, Snyder CRR. 1975. Facilitation and inhibition in the processing of signals. See Rabbitt & Dornic 1975, pp. 669–81

Rabbitt PMA, Dornic S, eds. 1975. *Attention and Performance V*. London: Academic

Remington RW, Folk CL, McLean JP. 2000. Contingent attentional capture or delayed allocation of attention? *Percept. Psychophys.* In press

Remington RW, Johnston JC, Yantis S. 1992. Involuntary attentional capture by abrupt onsets. *Percept. Psychophys.* 51:279–90

Rickard T. 1999. A CMPL alternative account of practice effects in numerosity judgment tasks. *J. Exp. Psychol.: Learn. Mem. Cogn.* 25:532–42

Rogers RD, Monsell S. 1995. Costs of a predictable switch between simple cognitive tasks. *J. Exp. Psychol.: Gen.* 124:207–31

Rosenbloom P, Newell A. 1987. Learning by chunking: a production system model of practice. In *Production System Models of Learning and Development*, ed. D Klahr, P Langley, pp. 221–86. Cambridge, MA: MIT Press

Ruthruff E, Johnston JC, Van Selst MV. 2000a. Why practice reduces dual-task interference. *J. Exp. Psychol.: Hum. Percept. Perform.* In press

Ruthruff E, Pashler HE, Klaassen A. 2000b. Processing bottlenecks in dual-task performance: structural limitation or voluntary postponement? *Psychonom. Bul. Rev.* In press

Ruthruff E, Remington RW, Johnston JC. 2000c. Switching between simple cognitive tasks: the interaction between top-down and bottom-up factors. *J. Exp. Psychol.: Hum. Percept. Perform.* In press

Salthouse TA, Fristoe N, McGuthry KE, Hambrick DZ. 1998. Relation of task switching to speed, age, and fluid intelligence. *Psychol. Aging* 13:445–61

Schneider W, Shiffrin RM. 1977. Controlled and automatic human information processing: I. Detection, search and attention. *Psychol. Rev.* 84:1–66

Schumacher EH, Seymour T, Glass JM, Lauber EJ, Kieras DE, Meyer DE. 1997. *Virtually perfect time sharing in dual-task performance*. Presented at 38th Annu. Meet. Psychonom. Soc., Philadelphia

Spector A, Biederman I. 1976. Mental set and mental shift revisited. *Am. J. Psychol.* 89:669–79

Sternberg S. 1969. The discovery of processing stages: extensions of Donders' method. In *Attention and Performance II*, ed. WG Koster, pp. 276–315. Amsterdam: North Holland

Sudevan P, Taylor DA. 1987. The cueing and priming of cognitive operations. *J. Exp. Psychol.: Hum. Percept. Perform.* 13:89–103

Theeuwes J. 1991. Exogenous and endogenous control of attention: the effect of visual onsets and offsets. *Percept. Psychophys.* 49:83–90

Theeuwes J. 1992. Perceptual selectivity for color and form. *Percept. Psychophys.* 51:599–606

Titchener EB. 1908. *Lectures on the Elementary Psychology of Feeling and Attention*. New York: Macmillan

Townsend JT. 1974. Issues and models concerning the processing of a finite number of inputs. In *Human Information Processing: Tutorials in Performance and Cognition*, ed. BH Kantowitz, pp. 133–68. Hillsdale, NJ: Erlbaum

Van Selst MV, Ruthruff E, Johnston JC. 1999. Can practice eliminate the Psychological Refractory Period effect? *J. Exp. Psychol.: Hum. Percept. Perform.* 25:1268–83

Welford AT. 1967. Single-channel operation in the brain. *Acta Psychol.* 27:5–22

Wolfe JM. 1992. The parallel guidance of visual attention. *Curr. Dir. Psychol. Sci.* 1:124–28

Wolfe JM, Cave KR, Franzel SL. 1989. Guided search: an alternative to the feature integration model of visual search. *J. Exp. Psychol.: Hum. Percept. Perform.* 15:419–33

Yantis S, Egeth HE. 1999. On the distinction between visual salience and stimulus-driven attentional capture. *J. Exp. Psychol.: Hum. Percept. Perform.* 25:661–76

Yantis S, Hillstrom A. 1994. Stimulus-driven attentional capture: evidence from equiluminant visual objects. *J. Exp. Psychol.: Hum. Percept. Perform.* 20:95–107

Yantis S, Jonides J. 1984. Abrupt visual onsets and selective attention: evidence from visual search. *J. Exp. Psychol.: Hum. Percept. Perform.* 10:601–21

Yantis S, Jonides J. 1990. Abrupt visual onsets and selective attention: evidence from visual search. *J. Exp. Psychol.: Hum. Percept. Perform.* 16:121–34