Journal of Experimental Psychology:
Learning, Memory, and Cognition
1987, Vol. 13, No. 2, 327–334

0278-7393/87/$00.75

Selective Attention and the Suppression of Cognitive Noise

W. Trammell Neill and Richard L. Westberry
University of South Florida

The *distractor-suppression effect* is the relative slowing of Stroop (1935) color naming when the current appropriate response is identical to the inappropriate response activated by the distractor word appearing in the immediately preceding trial. Two experiments investigated aspects of the time course of distractor suppression. Experiment 1 found the suppression effect when subjects were instructed to maintain strict accuracy but not when subjects were encouraged to sacrifice some accuracy for greater speed. Experiment 2 traced the recovery from suppression by varying the interval between successive trials (20, 520, 1020, or 2020 ms). The suppression effect was found to persist for at least a second; by 2 s the effect was completely dissipated. The results support the view of selective inhibition as an active, time-dependent control process that develops over time following the activation of distracting information and that is released after response to the task-appropriate information has been made. The results are interpreted in the context of a model in which widespread automatic activation in memory is followed by a process of "narrowing down" the range of activated representations to those specifically appropriate to current task demands (Keele & Neill, 1978).

In situations requiring selective attention, we respond to some subset of present or potential stimulus information while avoiding response to other information that is potentially or actually distracting to performance. Most theories have viewed attention as the direct allocation of some processing resource or capacity to the relevant, attended information (e.g., Broadbent, 1958; Deutsch & Deutsch, 1963; Kahneman, 1973; Keele, 1973; Norman, 1968; Norman & Bobrow, 1975; Posner, 1978; Shiffrin & Schneider, 1977). In other words, the mechanism of selective attention is not presumed to operate directly on information about unattended events. Thus, the activation of information about unattended events is expected to simply dissipate passively at whatever levels such activation has occurred. As a consequence, recently activated but irrelevant information should not be any less available for later processing than other such information not as recently activated.

Neill (1977, 1977/1978a, 1979; Keele & Neill, 1978) has argued that, in addition to the facilitation of processing relevant information, selective attention entails the active inhibition of specific distracting information. If the degree of such inhibition exceeds the degree of initial activation, then recently ignored information may become temporarily less available for later processing than other information not as recently activated. This possibility allows an empirical distinction between selective facilitation of processing relevant information and selective inhibition of processing irrelevant information.

Selective inhibition differs operationally and theoretically from other forms of inhibition that have been proposed to affect information processing. For example, Posner and Snyder (1975a, 1975b) refer to "*inhibition*" of processing unattended signals caused by the commitment of processing capacity to an attended signal (see also Neely, 1976, 1977). Inhibition here is de facto and nonselective, and it does not reflect an operation on specific unattended information (Posner, 1982). Other proposed forms of nonselective inhibition, such as *attenuation* of all unattended information (Treisman, 1964), or *dampening* of spreading activation (Anderson, 1976) also do not predict a specific bias against more recently activated information. Similarly, models based on an automatic *lateral inhibition* between cognitive or perceptual units (e.g., Brown, 1979; Estes, 1972; McClelland & Rumelhart, 1981; Walley & Weiden, 1973) do not make this prediction; because degree of inhibition would be directly determined by proximity (similarity) between activated units, more recently activated units should remain relatively more activated if such proximity is equated.

The Stroop (1935) color-word task is prototypical of selective attention paradigms: The subject must respond to a subset of available stimulus information (color in which a word is written) and avoid response to other, distracting information (meaning of the word).[1] Dalrymple-Alford and Budayr (1966) found that total time to name colors in a list of Stroop words was especially slow if each color corresponded to the distractor

We are indebted to Colleen Bridges, Lisa Cargerman, Kris Cavendish, Joy Dorchak, Manny Garcia, David Kaplan, Vicki Lowe, Roy McAllister, Jerome Nadel, Elfriede Niedhardt, Doreen Nicoff, Jim Quarterero, Carol Roth, Shiffra Royffé, Diane Schiff, Alan Shaffren, Phil Spoto, Doug Stokes, Margo Tillinghast, Myles Tonnacliff, Nancy Traylor, and Barry Tritt for assistance in these and related experiments. Comments by Douglas Lowe and two anonymous reviewers greatly facilitated reorganization of this manuscript.

Correspondence concerning this article should be addressed to W. Trammell Neill, who is now at the Department of Psychology, Adelphi University, Garden City, New York 11530.

[1] It has been suggested that selective attention is more difficult when relevant and irrelevant dimensions are integrated in the same stimulus object (Broadbent, 1982; Treisman, 1969). It should be noted, however, that Stroop color-word interference also occurs when color and word are spatially separated (Dyer, 1973). More important, suppression of distracting information has been demonstrated with spatially separated relevant and irrelevant stimuli (Allport, Tipper, & Chmiel, 1985; Tipper & Cranston, 1985).

word immediately preceding it in the list. Because this result was a potential artifact of the massed-list procedure, Neill (1977, Experiment 1) replicated this effect by measuring vocal naming latencies to randomized, individually presented Stroop words. For example, if the word RED written in color *green* were presented on trial *n*, reaction time on trial *n + 1* was found to be slower if the required response were "red" than if it were unrelated, e.g., "blue." If selective attention on trial *n* involved only the direct facilitation of "green", the response "red" should either remain more available than "blue" or (if activation failed to persist) not differ from it. Similarly, any broad, nonselective or lateral inhibition could not account for the diminished availability of "red" relative to "blue." The inhibition appears to be specific to the activated, distracting information.

This *distractor-suppression effect* has been replicated in several studies using the Stroop color-naming task (Lowe, 1979, 1985; Neill, 1982; Westberry, 1983), as well as in analogous designs with other materials. Allport, Tipper, and Chmiel (1985, Experiments 1 & 2) and Tipper and Cranston (1985) required subjects to name letters written in a specified color while ignoring simultaneous letters written in another color: When a target letter to be named had been presented as a distractor to be ignored on the preceding trial, naming time was slowed. Allport et al. (1985, Experiments 3–9) and Tipper (1985) obtained similar effects when subjects attended to line drawings of objects while ignoring simultaneous drawings in another color. A particularly interesting finding by Tipper (1985, Experiment 3) is that the suppression effect appears to generalize to the semantic category of the ignored object. Thus, having to ignore a pictured cat, for example, subsequently slows naming a pictured dog.

There are, however, conditions in which the distractor-suppression effect is not obtained. In particular, Neill (1977, Experiment 2) found a relative facilitation, rather than inhibition, when the current color matched the immediately preceding distractor word. Unlike Neill (1977, Experiment 1), this experiment used manual key-press responding. Several investigators (Allport et al., 1985; Lowe, 1985; Tipper & Cranston, 1985) have cited the difference in response modalities as support for alternative interpretations of the suppression effect. However, both Neill (1982) and Westberry (1983) have replicated the suppression effect with manual key-press responses, so response modality does not appear to be the critical factor. A more likely explanation for the conflicting results of Neill (1977) lies in different demands for speed versus accuracy in the two experiments: In the first experiment, subjects were instructed to respond as accurately as possible, so as to minimize practical difficulties in scoring vocal errors; in the second, subjects were encouraged to sacrifice some accuracy for greater speed, in order to obtain analyzable error rates.

That relative demand for speed versus accuracy should affect selective inhibition is implicit in the model of attention proposed by Keele and Neill (1978). Consistent with most current theories, it is assumed that information can be activated in memory automatically, either directly by external stimuli (cf. Eriksen & Schultz, 1979; Keele, 1973; Posner, 1978; Shiffrin & Schneider, 1977) or indirectly by association to other activated information (cf. Anderson, 1976; Anderson & Bower, 1973; Collins & Loftus, 1975; Fischler, 1977; Neely, 1976, 1977; Schvaneveldt & Meyer, 1973). However, if such activation occurs automatically, then it must do so regardless of appropriateness to current task demands. Consequently, irrelevant activations may lead to incompatible decisions or responses and interfere with task performance, as in the Stroop effect. In other words, automatic activation may contribute "noise" to decision processes at some point or points in the processing of information.

Keele and Neill (1978) argue that the automatic activation of information in memory must be followed by some process of "narrowing down" the range of activated memory structures to those specifically appropriate to current task demands. It is proposed that this narrowing-down may be accomplished through the inhibition of the activated, but task-inappropriate, memory structures. Assuming this process to be time dependent, it follows that the degree of inhibition of distracting information will depend on how much time is allowed for such inhibition to occur. If subjects are encouraged to respond very rapidly, at a cost to accuracy, then irrelevant memory structures may remain relatively activated, which will yield facilitation should such structures subsequently become task relevant.

Evidence that selective inhibition depends on task demands for speed versus accuracy was found by Neill (1979) in a variation of the priming design developed by Posner and Snyder (1975a, 1975b). Subjects were required to judge paired letters or digits as "same" or "different." Prior to each pair, a letter or digit warning signal indicated that character as likely to appear in the upcoming pair. On some trials, subjects received unexpected stimuli drawn from either the same or opposite category (letters or digits) as the expected stimulus. Insofar as members of a category are assumed to be more highly associated to each other than to members of other categories, current theories of associative activation predict that stimuli in the same category as the primed stimulus should be responded to more easily than stimuli in the opposite category. When subjects were encouraged to sacrifice accuracy for speed, such was indeed the case. However, when subjects were given strict accuracy instructions, unexpected stimuli from the same category as the expected stimulus were matched less accurately than unexpected stimuli from the opposite category. These results were interpreted as indicating that the activation of intracategory associates may hamper either identification or matching of particular category members. For example, a pair like AA might be misrecognized as AB, or vice versa. Under strict accuracy instructions, such interference may be circumvented by the selective suppression of intracategory associates to the expected character.

In the Stroop task, the source of irrelevant activation is primarily the externally presented distractor word, rather than internal association to relevant activations in memory. A similar narrowing-down process is assumed to be necessary, regardless of whether the source of interference is external or internal. Thus, it is of theoretical importance to demonstrate that the suppression of such cognitive noise in either case is affected by similar variables. Consequently, Experiment 1 directly manipulated instructional emphasis on speed versus accuracy in the Stroop task, similar to the procedure of Neill (1979, Experiment 2).[2] In addition, this manipulation should shed light on

[2] Neill (1982) attempted to manipulate speed versus accuracy emphasis in the Stroop task between separate groups of subjects. Unfortunately, the instructional manipulation did not produce overall performance differences between the two groups. However, when the data were

the conflicting results of Neill (1977), which have contributed to controversy concerning the nature of the distractor suppression effect.

Experiment 1

Method

Subjects. Subjects were 16 University of South Florida undergraduates, each of whom participated in an individual session of approximately 1.5 hr, and received extra credit in a psychology course.

Apparatus. Stimuli were presented on a Sony KV-1514 color television, controlled by an Apple II Plus microcomputer via a SUP'R'MOD II TV interface unit. The microcomputer also recorded key-press responses from four microswitch response keys and measured response latencies from a Mountain Hardware Apple Clock.

Stimuli. Stroop conflict stimuli consisted of the words *red, blue, green,* and *yellow* printed in red, blue, green, or yellow color, with the constraint that a word not be written in the color named by that word. Nonconflict stimuli were generated by substituting *o* for each letter in the conflict words (i.e., *ooo, oooo, ooooo, oooooo*), preserving the same color constraint. Consequently, there were 12 possible conflict stimuli and 12 possible neutral stimuli. Stimuli measured from 7.5 to 13.8 cm in length and were viewed from a distance of approximately 120 cm.

Procedure. Each subject participated in 6 blocks of 100 trials each under instructions emphasizing speed but strict accuracy and in 6 blocks of 100 trials each under instructions emphasizing speed at a cost to accuracy. The first block under each instructional condition was regarded as practice and was not included in the data analysis. Half of the subjects performed under strict accuracy instructions first, and half performed under lax accuracy instructions first.

Strict accuracy instructions emphasized speed, but with the constraint of making as few errors as possible. At the end of a block, if a subject made more than five errors, the message PLEASE TRY TO RESPOND MORE CAUTIOUSLY was displayed. If the subject made five or fewer errors in that block, the message PLEASE CONTINUE TO RESPOND QUICKLY was displayed.

Lax accuracy instructions emphasized sacrificing some caution to achieve greater speed. At the end of a block, if a subject made fewer than 8 errors, the message PLEASE RESPOND FASTER, LESS CAUTIOUSLY was displayed. If the subject made more than 12 errors, the message PLEASE TRY TO RESPOND MORE CAUTIOUSLY was displayed. If the subject made between 8 and 12 errors, the message PLEASE CONTINUE TO RESPOND QUICKLY was displayed. Thus, strict accuracy instructions encouraged an error rate of under 5%, whereas lax accuracy instructions encouraged an error rate of approximately 10%.

On each trial the subject was shown a single conflict or nonconflict stimulus, selected randomly but with .75 probability of a conflict stimulus on any trial. Subjects were required to make a manual key-press response to each stimulus. All subjects had the following response assignment: "red," left middle finger; "green," left index finger; "blue," right index finger; "yellow," right middle finger. Each stimulus remained in view until the subject's response. If the response was correct, the next stimulus followed the response by approximately 20 ms; if incorrect, the word ERROR was displayed for approximately 2 s prior to the next stimulus.

Table 1
Mean Reaction Times (in Milliseconds) and Percentage of Error as a Function of Stimulus Type and Relation to Preceding Conflict Stimulus, Under Strict or Lax Accuracy Demand: Experiment 1

	Strict		Lax	
Stimulus	RT	%E	RT	%E
Conflict				
Related	882	5.1	775	11.0
Unrelated	851	4.3	794	10.3
Nonconflict				
Related	839	4.0	744	6.0
Unrelated	819	3.2	760	7.5

Data analysis. Each trial was classified according to the relation of the current stimulus color and word to the color and word presented on the preceding trial. Of present concern are four relations: (a) Conflict trials following conflict trials, in which the current color matches the preceding distractor word, with no other direct relation (e.g., *blue* written in green, following *green* written in red); (b) Conflict trials following conflict trials, in which stimuli are unrelated (e.g., *blue* written in yellow, following *green* written in red); (c) Nonconflict trials following conflict trials, in which the current color matches the preceding distractor word (e.g., *oooo* written in green, following *green* written in red); (d) Nonconflict trials following conflict trials, in which stimuli are unrelated (e.g., *oooo* written in yellow, following *green* written in red).[3] Preliminary analysis indicated that the lengths of the nonconflict stimuli, which were matched to conflict stimuli (e.g., *ooo—red*), did not affect reaction times and so were ignored for the final classification. The stimulus probabilities resulted in each subject receiving approximately 47 conflict-related, 47 conflict-unrelated, 23 nonconflict-related, and 47 nonconflict-unrelated trials under each instructional condition. Error trials and trials immediately following error trials were not included in the reaction time analysis.

A logarithmic transformation was applied to the reaction time data in order to reduce the influence of exceptionally long latencies (see Winer, 1971). Logarithms of the raw latencies were averaged in each of the eight conditions for each subject. The antilogs of these means were then subjected to a $2 \times 2 \times 2$ analysis of variance (ANOVA) with within-subjects variables of stimulus type (conflict vs. nonconflict), relatedness (related vs. unrelated) and instructional emphasis (strict vs. lax accuracy).

Results and Discussion

The results of Experiment 1 are displayed in Table 1. Conflict stimuli produced longer reaction times (826 ms) than nonconflict stimuli (791 ms), $F(1, 15) = 18.03$, $MS_e = 2{,}175$, $p < .001$. Strict accuracy instructions produced longer reaction times (848 ms) than lax accuracy instructions (768 ms), $F(1, 15) =$

collapsed over groups, a significant correlation was found between the suppression effect and an independent measure of overall reaction time, Pearson $r = .57$, $p < .01$. That is, subjects who responded more slowly and cautiously tended to show a greater suppression effect. In the present experiment, it was expected that within-subjects manipulation would more effectively contrast the instructional sets for the subjects; the instructional manipulation by Neill (1979, Experiment 2) was similarly within subjects.

[3] There are 14 conceptually different relations possible between successive conflict and/or nonconflict stimuli in this design. The software developed for this study isolated the reaction times only for the four conditions critical to the present hypotheses. Consequently, data for other possible relations are not available from these experiments. However, some of these relations are relevant to other issues, and have been discussed elsewhere (Lowe, 1979, 1985; Neill, 1977/1978a, 1978b).

26.31, $MS_e = 7{,}687$, $p < .001$. In addition, relatedness interacted significantly with instructions, $F(1, 15) = 12.99$, $MS_e = 1{,}010$, $p < .005$. A Fisher's least significant difference (LSD) of 24 ms (two-tailed, $p < .05$) indicates that this interaction is due to a significant suppression effect (related–unrelated difference) of 26 ms under strict accuracy instructions, whereas a nonsignificant trend in the opposite direction, −17 ms, was found under lax accuracy instructions. Neither the main effect of relatedness nor any other interactions were significant.

Conflict stimuli produced more errors (7.7%) than nonconflict stimuli (5.2%), $F(1, 15) = 6.30$, $MS_e = 31.8$, $p < .025$. Strict accuracy instructions resulted in fewer errors (4.2%) than lax accuracy instructions (8.7%), $F(1, 15) = 8.88$, $MS_e = 74.6$, $p < .01$. In addition, a marginal interaction ($p < .10$) of these two variables was obtained, $F(1, 15) = 4.26$. This marginal interaction appears to reflect a larger conflict–nonconflict difference under lax accuracy instructions, and may be attributable to ceiling effects on accuracy under strict accuracy instructions.

The color-word interference effect reported by Stroop (1935) and many others was replicated here in the conflict–nonconflict differences in both reaction time and errors. In addition, the instructional sets were clearly effective in producing speed–accuracy trade-off, as reflected in both measures. Of critical importance to the present hypothesis is that instructions interacted dramatically with relatedness: As predicted, the distractor-suppression effect found by Neill (1977, Experiment 1) was replicated under strict accuracy instructions. Under lax accuracy instructions, a nonsignificant trend was found in the opposite direction, consistent with the facilitation found by Neill (1977, Experiment 2). The present experiment confirms that the distractor-suppression effect is obtainable when manual key-press responses are required (Neill, 1982; Westberry, 1983). Hence, the failure of Neill (1977, Experiment 2) to find the effect is not simply attributable to use of manual responses, but, rather, is probably due to a relatively lax demand for accuracy in that experiment.

Note that the distractor-suppression effect and its interaction with instructions generalizes here to nonconflict stimuli preceded by conflict stimuli. If the effect does indeed reflect a suppression of the preceding distractor response, that response should be hampered regardless of the current stimulus type. Neill (1982) found distractor suppression when probed by the same nonconflict stimuli (colored zeroes) as used here. Lowe (1979, Experiment 4) found the suppression effect to generalize to random letter strings intermixed with conflict words, confirming that the probe stimulus does not itself have to generate conflict. On the other hand, when the nonconflict stimuli were simple color patches, Lowe found facilitation rather than suppression when the required response matched the preceding distractor word. Allport et al. (1985, Experiment 9) and Tipper and Cranston (1985, Experiment 3) also found facilitation when no distractor was presented on the probe trial. That suppression is sometimes not found when probed by nonconflict stimuli (even when present for conflict stimuli) is of theoretical importance and will be addressed more fully in the General Discussion.

Experiment 2

The results of Experiment 1 were predicted by the assumption that the suppression of distracting information requires time to develop after the presentation of a conflict stimulus. Further support for the time-dependent growth of suppression was found by Lowe (1985, Experiments 2 and 3): Subjects were instructed to attend (but not respond) to the color of a briefly presented conflict word, and then respond to the color of a second conflict word whose relation to the first was varied. When stimulus onset asynchrony (SOA) between the first and second stimuli was manipulated (50, 100, 200, or 400 ms), Lowe found suppression only for the three longer SOAs. That this result is not due to backward masking at the shortest SOA is clear from other effects (the advantage of identical stimulus repetition over color-only repetition) fully established at that SOA.

Although distractor suppression appears to develop over time, the effect obviously cannot persist indefinitely. If it did, all responses would be equally inhibited after a few trials, making it impossible to detect an overt suppression effect as manifested in the related–unrelated difference. Rather, it seems likely that the distractor response should be released from inhibition after response to the appropriate information has been made. Thus, whereas Lowe (1985) found distractor suppression to increase as a function of SOA, it is predicted to decrease as a function of delay after response to the preceding stimulus. Accordingly, Experiment 2 tested this hypothesis by varying the response–stimulus interval (RSI) between successive trials.

Method

Subjects. Subjects were 12 University of South Florida undergraduates, each of whom participated in an individual session of approximately 1 hr and received extra credit in a psychology course.

Stimuli and apparatus. Stimuli and apparatus were identical to Experiment 1.

Procedure. Each subject participated in 11 blocks of 100 trials each. The first block was regarded as practice and was not analyzed. After a subject's response on a given trial, the presentation of the next stimulus was delayed randomly by approximately 20, 520, 1,020, or 2,020 ms. Subjects were given strict accuracy instructions, as defined in Experiment 1. Other aspects of procedure were similar to Experiment 1.

Data analysis. Reaction times were logarithmically transformed as in Experiment 1. Stimulus type and relatedness conditions were as defined in Experiment 1. Antilogs of the means for each subject in each condition were subjected to a 2 × 2 × 4 ANOVA with within-subjects variables of stimulus type (conflict vs. nonconflict), relatedness (related vs. unrelated), and RSI (20, 520, 1,020, or 2,020 ms).

Results and Discussion

The mean reaction times for the various conditions are displayed in Table 2. An ANOVA yielded significant main effects of stimulus type, $F(1, 11) = 203.56$, $MS_e = 436$, $p < .001$; relatedness, $F(1, 11) = 8.14$, $MS_e = 905$, $p < .025$; and RSI, $F(3, 33) = 88.05$, $MS_e = 770$, $p < .001$. In addition, RSI interacted significantly with stimulus type, $F(3, 33) = 2.99$, $MS_e = 784$, $p < .05$; and with relatedness, $F(3, 33) = 3.11$, $MS_e = 721$, $p < .05$.

Interpretation of the reaction time results is facilitated by simultaneously considering the accuracy data, which is displayed in Table 3. An ANOVA of the error percentages yielded significant main effects of stimulus type, $F(1, 11) = 7.00$, $MS_e = 14.9$, $p < .025$; and RSI, $F(3, 33) = 5.20$, $MS_e = 25.3$, $p < .01$. No interactions reached significance.

The significant effect of stimulus type on reaction time again reflects the typical Stroop interference effect, with conflict stim-

Table 2
Mean Reaction Times (in Milliseconds) as a Function of Stimulus Type, Relation to Preceding Conflict Stimulus, and Response–Stimulus Interval (RSI; in Milliseconds): Experiment 2

	RSI			
Stimulus	20	520	1,020	2,020
Conflict				
Related	870	862	903	916
Unrelated	860	839	893	924
Nonconflict				
Related	837	801	871	884
Unrelated	811	771	860	886

uli producing slower reaction times (883 ms) than nonconflict stimuli (840 ms). This pattern is supported by the significant main effect of stimulus type on the error percentages, with more errors occurring to conflict stimuli (4.4%) than to nonconflict stimuli (2.9%).

The significant effect of RSI appears to reflect somewhat faster reaction times at shorter intervals. Average reaction times, in order of increasing RSI, were 845, 818, 882, and 903 ms. It should be noted, however, that the effect of RSI on accuracy in part compensates for the effects on speed: Average error percentages, in order of increasing RSI, were 5.9%, 4.0%, 2.2%, and 2.6%. The apparent speed–accuracy trade-off is consistent with other studies of foreperiod effects on reaction time and errors (e.g., Posner, Klein, Summers, & Buggie, 1973).

The interaction of stimulus type with RSI in reaction time reflects differences in Stroop interference at different RSIs, although not in a particularly systematic fashion. Collapsing over relatedness, the interference effects (conflict–nonconflict differences) were 41, 65, 22, and 35 ms, in order of increasing RSI. Increased caution at the longer RSIs (as reflected in the main effects on speed and accuracy) may have also contributed to reduced Stroop interference. However, it should be noted that direct manipulation of speed–accuracy trade-off in Experiment 1 did not produce an interaction with stimulus type. The dependence of Stroop interference on time between trials may warrant further investigation.

Table 3
Mean Percentages of Error Responses as a Function of Stimulus Type, Relation to Preceding Conflict Stimulus and Response–Stimulus Interval (RSI; in Milliseconds): Experiment 2

	RSI			
Stimulus	20	520	1,020	2,020
Conflict				
Related	8.3	4.9	3.0	3.1
Unrelated	5.7	4.9	2.5	2.9
Nonconflict				
Related	4.9	3.2	1.6	2.1
Unrelated	4.6	3.0	1.7	2.4

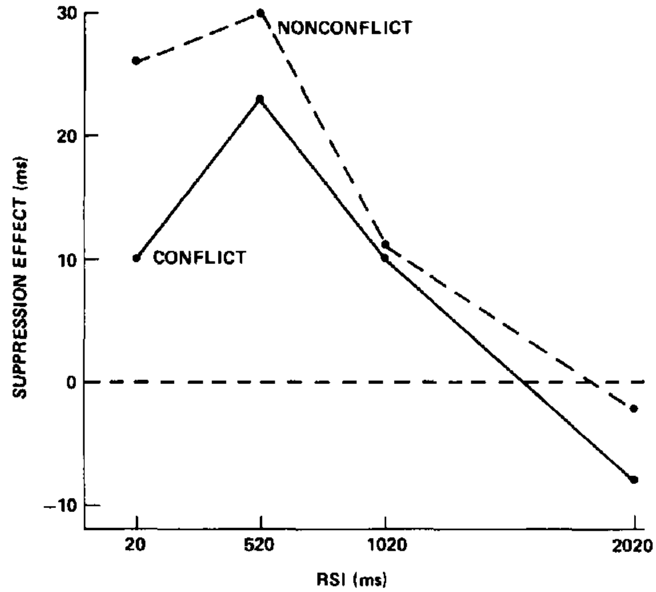

Figure 1. Suppression effect (related–unrelated difference) for conflict and nonconflict stimuli as a function of interval between preceding response and presentation of current stimulus.

The main effect of relatedness reflects longer reaction times overall for the related condition (868 ms) than for the unrelated condition (855 ms), which again replicates the distractor-suppression effect. Of greatest present concern is the interaction between relatedness and RSI. Figure 1 displays the suppression effects (related–unrelated differences) for both conflict and nonconflict stimuli, as a function of RSI. Collapsing over stimulus type, the suppression effects were 18, 26, 11, and −5 ms, in order of increasing RSI. Although the triple interaction of relatedness, stimulus type, and RSI did not achieve significance ($F < 1$), note in Figure 1 that the suppression effect appears to increase for conflict stimuli from RSI = 20 to RSI = 520. Because conflict stimuli in the related condition at RSI = 20 showed an exceptionally high error rate (8.3%), this may be an artifact of speed–accuracy trade-off on those trials. On the other hand, some small increase in the suppression effect over early RSIs would not be unexpected if the hypothesized inhibitory mechanism were prone to some "inertia" (i.e., continuing to suppress distracting information for a brief period following emission of the appropriate response). At any rate, the systematic decline from RSI = 520 to RSI = 2020 is consistent with the hypothesis that the suppression effect should dissipate with longer RSIs.

General Discussion

The two experiments reported here explore different aspects of the time course of the distractor-suppression effect. Keele and Neill (1978) proposed that the activation of information in memory must be followed by a process of narrowing down the range of activations to those specifically appropriate to current task demands. It has been suggested that this focusing process is accomplished through the direct, selective inhibition of distracting information. If such information subsequently becomes task relevant, it follows that processing may be hampered by the prior inhibition. However, the suppression effect requires time to develop, as demonstrated by Lowe (1985). Consequently, if a response is selected before the distracting information is completely inhibited, then that information may remain highly available, thereby facilitating processing when such information subsequently becomes task relevant.

Experiment 1 directly confirmed the prediction that the distractor suppression effect should be dependent on relative demand for speed versus accuracy. Under strict accuracy instructions, color naming was especially slowed if that response matched the distractor word presented in the immediately preceding trial, as had been found by Neill (1977, Experiment 1). However, under lax accuracy instructions, the effect was reversed, consistent with the results of Neill (1977, Experiment 2). As such, the results directly parallel the effects of speed versus accuracy demands on intracategory inhibition in character matching (Neill, 1979).

After a selected response has been emitted, there is no further need to inhibit the distracting information. It is logically necessary that the distracting information must recover from the inhibition; if not, then all possible responses would be equally inhibited after a few trials. Whereas Lowe (1985) found distractor suppression to increase with SOA, it would be predicted to decrease with time after response to a conflict stimulus. Accordingly, Experiment 2 manipulated RSI between trials. Suppression increased slightly from RSI = 20 ms to RSI = 520 ms, perhaps reflecting inertia of the inhibitory mechanism, but it subsequently decreased over longer RSIs. The suppression effect appears to persist for at least a second, but to dissipate completely by two seconds.

Several investigators (Allport et al., 1985; Lowe, 1985; Tipper & Cranston, 1985) have assumed that the failure of Neill (1977, Experiment 2) to find the distractor-suppression effect was due to the use of manual responding, rather than vocal naming, in that experiment. Note in this regard that both of the present experiments found distractor suppression with manual keypress responding, as had previous unpublished experiments (Neill, 1982; Westberry, 1983). Note also that the suppression effect generalized to neutral, nonconflict stimuli in both experiments; neither experiment gave any indication of an interaction of stimulus type (conflict vs. nonconflict) with relatedness. If the suppression effect does in fact reflect response inhibition carried over from the preceding trial, then that response would be expected to be slowed regardless of the present stimulus type.

A finding that has been problematic for a simple model of response inhibition is Lowe's (1979, Experiment 4) finding of facilitation for naming color patches when the response matched the preceding distractor word, even when suppression occurred for naming the color of a Stroop conflict word. Because subjects could not predict whether a color patch or conflict word would be shown, processing of the preceding distractor word must be assumed to be the same in either case.[4] On the other hand, Lowe did find suppression to generalize to nonconflict stimuli composed of random letters, just as it did to colored zeroes in the present experiments. Hence, it is clear that the probe stimulus need not itself generate competing responses in order to manifest suppression of the preceding distractor response. Why should color patches be exceptional?

Lowe (1985) suggested that the distractor-suppression effect may be due to difficulty in coordinating activated responses with the appropriate perceptual information. For example, if the response "green" has just been activated by a distractor word on the preceding trial, and is now appropriate to naming the color, the subject may have difficulty determining that "green" is in fact appropriate to the relevant attribute (color) on the current trial. Presumably the subject does not experience this confusion if the irrelevant attribute is sufficiently discriminable from a color word, as in the case of a color patch. Allport et al. (1985) proposed a similar "code-coordination" account for distractor suppression in their letter-naming and picture-naming tasks.

The code-coordination hypothesis would seem inadequate to explain Lowe's (1985) own finding that distractor suppression increases with SOA. That is, confusability between stimuli would be expected to diminish with increasing delay between them. Moreover, Tipper and Cranston (1985, Experiment 4) directly tested code coordination against response inhibition in a variation of the letter-naming task used by them and Allport et al. (1985). Subjects were required to identify the red letter and ignore the green letter in an initial (priming) display; then, in a succeeding (probe) display, they were required to identify the green letter and ignore the red. Suppose the letter *A* were presented in green in both displays. According to the code-coordination hypothesis, the response "A" should be doubly associated to the relevant selection attribute for the probe display and should show facilitation relative to a new target letter. According to the response inhibition hypothesis, the response "A" has been suppressed in attending to the priming display, and so should be hampered despite the change in relevant selection attribute for the probe display. Tipper and Cranston found the latter result, and so concluded in favor of response inhibition.

A modified version of the selective inhibition model, proposed by Tipper (1985; Tipper & Cranston, 1985), seems to most effectively account for Lowe's (1979) results, while still accommodating the accumulated evidence for response inhibition. According to this modified model, inhibition and facilitation reflect different stages of processing. If the internal representation or categorization of the distractor were itself inhibited (as assumed by Neill, 1979, in particular), it would be difficult to explain facilitation of naming a related color patch. Tipper suggests instead that what is inhibited is access to mechanisms for overt response. In other words, the representations of ignored objects are "isolated from the control of action", without inhibiting the activated representations themselves. When nonconflict stimuli are sufficiently wordlike (e.g., random letters or

[4] This assumption cannot be made for similar demonstrations by Allport et al. (1985) and Tipper and Cranston (1985). Allport et al. found distractor suppression when a to-be-ignored picture was subsequently presented as a target superimposed on a new distractor ("duplex probe") but facilitation when it was presented as a target without a distractor ("simplex probe"). Tipper and Cranston found similar results when stimuli were letters to be named or ignored. However, in both experiments the type of probe was manipulated between groups of subjects. Thus, it is possible that subjects anticipating the easier task simply attended less to the priming stimulus, and so did not inhibit the distractor response. It should also be noted that probe stimuli used by Tipper (1985), which did show distractor suppression, were drawings superimposed on nonsense figures and, as such, were not conflict stimuli per se.

strings of zeroes) to be perceived as potentially interfering, subjects continue to deny recently ignored activations access to response mechanisms, producing the distractor suppression effect. However, when nonconflict stimuli are clearly different from conflict stimuli (e.g., color patches), this "selection state" is discontinued, and responding will benefit from the prior activation of previously distracting representations.

The difficulty of responding to stimuli related to previously ignored stimuli has been demonstrated using a variety of materials: Stroop color words (Lowe, 1979, 1985; Neill, 1977, 1982; and the present experiments), superimposed or separated pictures (Allport et al., 1985; Tipper, 1985) and superimposed or separated letters (Allport et al., 1985; Tipper & Cranston, 1985). Experiments on color naming (Lowe, 1979, Experiment 4), picture naming (Allport et al., 1985, Experiment 9), and letter naming (Tipper & Cranston, 1985, Experiment 3) all have suggested a release from inhibition when probe stimuli are clearly distinguishable from stimuli leading to conflicting responses. That the distractor-suppression effect is affected by speed versus accuracy emphasis in much the same manner as intracategory inhibition in character matching (Neill, 1979) further supports a commonality of mechanism across selective attention tasks. Such results are incompatible with theories that view attention only as the selective allocation of certain processing resources to the relevant information, and especially with theories that assert that irrelevant information is simply not "picked up" by perceptual systems (e.g., Neisser, 1976). Rather, the relative unavailability of recently ignored information implies that such information has been selectively suppressed.

Although inhibitory processes have played a prominent theoretical role in other fields of psychology (e.g., classical conditioning, neuropsychology, psychoanalytic theory), references to inhibitory mechanisms are scarce in information-processing theories, and references to selective inhibitory mechanisms in particular, practically nonexistent. Nevertheless, as observed by Roediger and Neely (1982), there are numerous phenomena which cannot be explained by only facilitatory relations between associated mental representations. Recent research by Marcel (1980, 1983) and Allport et al. (1985) suggests that selective inhibition of irrelevant information is directly associated with conscious awareness of relevant information. Westberry (1983; Westberry, Anker, & Neill, 1986) has found distractor suppression to be related to individual differences in "attentional style." Taken together with the variety of tasks in which inhibitory effects occur, such findings warrant the speculation that selective inhibition serves a pervasive function in the control of cognitive processing.

References

Allport, D. A., Tipper, S. P., & Chmiel, N. R. J. (1985). Perceptual integration and postcategorical filtering. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 107–132). Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1976). *Language, memory, and thought.* Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory.* Hillsdale, NJ: Erlbaum.

Broadbent, D. E. (1958). *Perception and communication.* London: Pergamon.

Broadbent, D. E. (1982). Task combination and selective intake of information. *Acta Psychologica, 50,* 253–290.

Brown, A. S. (1979). Priming effects in semantic memory retrieval processes. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 65–77.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82,* 407–428.

Dalrymple-Alford, E. C., & Budayr, B. (1966). Examination of some aspects of the Stroop color-word test. *Perceptual and Motor Skills, 23,* 1211–1214.

Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review, 70,* 80–90.

Dyer, F. N. (1973). Interference and facilitation for color naming with separate bilateral presentations of the word and color. *Journal of Experimental Psychology, 99,* 314–317.

Eriksen, C. W., & Schultz, D. W. (1979). Information processing in visual search: A continuous flow conception and experimental results. *Perception & Psychophysics, 25,* 249–263.

Estes, W. K. (1972). An associative basis for coding and organization in memory. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory* (pp. 161–190). Washington, DC: Winston.

Fischler, I. (1977). Associative facilitation without expectancy in a lexical decision task. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 18–26.

Kahneman, D. (1973). *Attention and effort.* Englewood Cliffs, NJ: Prentice-Hall.

Keele, S. W. (1973). *Attention and human performance.* Pacific Palisades, CA: Goodyear.

Keele, S. W., & Neill, W. T. (1978). Mechanisms of attention. In E. C. Carterette & M. P. Friedman (Eds.), *Handbook of perception* (Vol. 9, pp. 3–47). New York: Academic Press.

Lowe, D. G. (1979). Strategies, context, and the mechanism of response inhibition. *Memory & Cognition, 7,* 382–389.

Lowe, D. G. (1985). Further investigations of inhibitory mechanisms in attention. *Memory & Cognition, 13,* 74–80.

Marcel, A. J. (1980). Conscious and preconscious recognition of polysemous words: Locating the selective effects of prior verbal context. In R. S. Nickerson (Ed.), *Attention and performance VIII* (pp. 435–457). Hillsdale, NJ: Erlbaum.

Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15,* 197–237.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88,* 375–407.

Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory & Cognition, 4,* 648–654.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106,* 226–254.

Neill, W. T. (1977). Inhibitory and facilitatory processes in attention. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 444–450.

Neill, W. T. (1978a). Attention: The coordination of internal codes (Doctoral dissertation, University of Oregon, 1977). *Dissertation Abstracts International, 38,* 5067B.

Neill, W. T. (1978b). Decision processes in selective attention: Response priming in the Stroop color-word task. *Perception & Psychophysics, 23,* 80–84.

Neill, W. T. (1979). Switching attention within and between categories: Evidence for intracategory inhibition. *Memory & Cognition, 7,* 283–290.

Neill, W. T. (1982, March). *The suppression of distracting information* *in selective attention.* Paper presented at the meeting of the Southeastern Psychological Association, New Orleans, LA.

Neisser, U. (1976). *Cognition and reality.* San Francisco: Freeman.

Norman, D. A. (1968). Toward a theory of memory and attention. *Psychological Review, 75,* 522–536.

Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology, 7,* 44–64.

Posner, M. I. (1978). *Chronometric explorations of mind.* Hillsdale, NJ: Erlbaum.

Posner, M. I. (1982). Cumulative development of attention theory. *American Psychologist, 37,* 168–179.

Posner, M. I., Klein, R., Summers, J., & Buggie, S. (1973). On the selection of signals. *Memory & Cognition, 1,* 2–12.

Posner, M. I., & Snyder, C. R. R. (1975a). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55–85). Hillsdale, NJ: Erlbaum.

Posner, M. I., & Snyder, C. R. R. (1975b). Facilitation and inhibition in the processing of signals. In P. M. A. Rabbitt & S. Dornic (Eds.), *Attention and performance V* (pp. 669–682). New York: Academic Press.

Roediger, H. L., & Neely, J. H. (1982). Retrieval blocks in episodic and semantic memory. *Canadian Journal of Psychology, 36,* 213–243.

Schvaneveldt, R. W., & Meyer, D. E. (1973). Retrieval and comparison processes in semantic memory. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 395–409). New York: Academic Press.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127–190.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18,* 643–662.

Tipper, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *Quarterly Journal of Experimental Psychology, 37A,* 571–590.

Tipper, S. P., & Cranston, M. (1985). Selective attention and priming: Inhibitory and facilitatory effects of ignored primes. *Quarterly Journal of Experimental Psychology, 37A,* 591–611.

Treisman, A. M. (1964). The effect of irrelevant material on the efficiency of selective listening. *American Journal of Psychology, 77,* 533–546.

Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological Review, 76,* 282–299.

Walley, R. E., & Weiden, T. D. (1973). Lateral inhibition and cognitive masking: A neuropsychological theory of attention. *Psychological Review, 80,* 284–302.

Westberry, R. L. (1983). *The nature of attentional control as a personality dimension.* Unpublished doctoral dissertation, University of South Florida, Tampa.

Westberry, R. L., Anker, J. M., & Neill, W. T. (1986). *Attentional control as a personality dimension.* Manuscript in preparation.

Winer, B. J. (1971). *Statistical principles in experimental design* (2nd ed.). New York: McGraw-Hill.

Received August 5, 1985
Revision received June 30, 1986 ■