

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Adding Distractors Improves Performance by Boosting Top-Down Control

Permalink

<https://escholarship.org/uc/item/0z53h0x1>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

ISSN

1069-7977

Authors

Juvina, Ion
Taatgen, Niels

Publication Date

2009

Peer reviewed

Adding Distractors Improves Performance by Boosting Top-Down Control

Ion Juvina (ijuvina@cmu.edu) & Niels A. Taatgen (taatgen@cmu.edu)

Department of Psychology, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213

Abstract

The effect of adding to-be-ignored extra stimuli to a modified Stroop task is investigated. Adding extra stimuli of the same kind as the distractor causes a temporary improvement in performance (Stroop dilution), whereas adding extra stimuli of the same kind as the target causes an improvement in performance that is only detectable when the extra stimuli are removed (post-treatment). An attempt is made to explain these different outcomes in light of the existing theoretical accounts of the Stroop dilution effect. A computational model that accounts for the observed data is proposed. Results suggest that a top-down control mechanism compensates for lateral inhibition effects, particularly when they have a potentially disruptive influence on performance. This boost of control seems to last longer than needed, causing performance improvements in a post-treatment condition. A further implication of these results is that the top-down control function is trainable.

Keywords: Stroop dilution; lateral inhibition; top-down control.

Background

One of the typical functions of our cognitive control system is interference resolution by protecting the execution of task-relevant sequences of actions against interference and distraction. The Stroop task is a landmark task for studying cognitive control and interference resolution (MacLeod, 1991; Miyake, Friedman, Emerson, Witzki, Howerter, Wager, 2000). In this task the participant is presented with a sequence of words written in various ink colors and instructed to name the ink color of each word. There are three conditions: (1) *congruent*, when the word meaning and the ink color are congruent, for example, the word "red" is written in red ink; (2) *incongruent*, when the word meaning and the ink color are incongruent, for example, the word "red" is written in green ink; and (3) *neutral*, when the word meaning does not refer to a color name, for example, the word "desk" written in any color. Typically, it takes more time to name the color of an incongruent word and less time to name the color of a congruent word than to name the color of a neutral word.

A common augmentation of the Stroop paradigm is obtained by adding to-be-ignored extra stimuli to the classical Stroop trial. Typically, these extra stimuli are words and they determine an increase in performance on the Stroop task, a phenomenon frequently referred to as *Stroop dilution* (Brown, Roos-Gilbert, & Carr, 1995; Cho, Lien, & Proctor, 2006; Kahneman & Chajczyk, 1983; MacLeod & Bors, 2002; Mitterer, La Heij, & Van der Heijden, 2003). Two main theoretical accounts have been proposed for the Stroop dilution effect:

- The attentional capture account (Kahneman & Chajczyk, 1983; Mitterer et al., 2003) states that the extra stimulus captures attention on some of the trials, thus causing the Stroop distractor to interfere less often with the Stroop target.
- The visual interference account (Brown et al., 1995; MacLeod & Bors, 2002) proposes that recognition of two or more words occurs in parallel. The extra words cause degradation of the early visual percept of the Stroop distractor impairing its recognition, thus making it less able to interfere with the Stroop target.

Both accounts postulate bottom-up mechanisms (attentional capture or visual interference) to be responsible for the Stroop dilution effect. The improvement in performance seems to be caused by deteriorating the automatic reading of the word feature of the Stroop stimulus. In other words, the extra stimuli make the distractor less interfering. However, the Stroop task is known to have a top-down component as well (Herd, Banich, & O'Reilly, 2006). None of the studies with which we are familiar was concerned with the impact of the extra stimuli on the top-down component involved in the Stroop task, which is selecting and processing the target.

Experiment

In order to study the impact of the extra stimuli on the top-down component of the task, we added a separate condition with extra stimuli of the same kind as the target (extra colors). The hypothesis is that the extra words would cause an increase in performance (Stroop dilution), whereas the extra colors would cause a decrease in performance. The reasoning behind the latter is that the extra colors would disrupt target identification while leaving the distractor with its full interfering potential.

We are also interested to know what impact these changes in performance have on a post-treatment condition when the extra stimuli are removed. We hypothesize that, if the changes are solely due to bottom-up mechanisms (attentional capture or visual interference), performance at post-treatment should return to its pre-treatment levels. If post-treatment performance is differentially or similarly affected by the two treatments (extra colors vs. extra words), then other mediating or modulating processes might be involved.

Method

This experiment is part of a larger project aimed at investigating the cognitive control aspects of multitasking. We are interested in interference control in tasks that

involve perceptual, cognitive, and motor components; the vocal component of the Stroop task does not interest us in this study. For this reason we have considered using a manual version of the Stroop task. However, the typical manual Stroop task, in which each color is mapped on a unique manual response, has been shown to produce reduced levels of interference and fast decrease in interference with practice (see MacLeod, 1991, for a review). The reduced interference is probably caused by the direct association that is formed with practice between the perception of colors and the associated manual responses. Thus the mapped key presses lose their dimensional overlap with color concepts (Kornblum, 1994) because the retrieval of a color name is bypassed. When memory retrievals are bypassed, the main source of interference in the Stroop task, that is reading and retrieving color names, no longer exists.

We asked participants to select the right answer from two options given on the screen, thus reintroducing words as source of interference. This way, naming a color involves going through a verbal step. Having to select names of colors presented on screen makes the manual Stroop task more compatible with the standard (vocal) Stroop task, by bringing back its semantic and linguistic components. Interference arises from the possibility to retrieve an incorrect color name as in the vocal variant of the task. Each response option has an equal probability to appear on the left or right sides of the stimulus, thus preventing the selection process from becoming automated.

Participants Sixty-three participants were recruited from Carnegie Mellon University's community via a website advertisement. Participant age ranged from 18 to 47 years with an average of 22. There were 22 women and 41 men.

Design The three classical trial types of the Stroop task – incongruent, congruent and neutral – were administered within participant, in equal proportions, and randomly mixed. The participants were assigned at random to two experimental groups: the extra colors group and the extra words group. The total of 602 trials were divided in three blocks as follows:

- The pre-treatment block was composed of 231 trials without extra stimuli.
- The treatment block was composed of 232 trials. In this block, the participants received extra colors or extra words depending on which experimental group they were assigned to.
- The post-treatment block was composed of 139 trials of the same kind as the pre-treatment trials.

The three blocks were administered within participant, that is, all participants went through all three blocks of trials. The pre- and post-treatment control blocks were intended to measure whether the changes in performance caused by the added extra stimuli lasted when the extra stimuli were removed.

Stimuli Stimuli were color names (red, blue, orange and green) and neutral words colored with one of the four colors denoted by the mentioned color names. The neutral words were 10 common English words unrelated semantically or phonologically to any of the color names. Stimuli were presented one at a time in the center of the screen. Two response options were also displayed flanking the stimulus on its left and right sides. Response options were non-colored (i.e., in black) color names. One response option contained the correct answer and the other one an incorrect answer. In the incongruent condition the incorrect answer was identical with the distractor word.

Figures 1 and 2 show a typical set of stimuli from a trial in the extra colors condition and a trial in the extra words condition, respectively. In the extra colors condition, three color-patches were presented above the stimulus. The extra colors were drawn at random from a set of eleven colors that included the four target colors.




Figure 1: A typical set of stimuli in the extra colors condition.

In the extra words condition, three words were presented above the stimulus. The extra words were drawn at random from a set of eleven words that included the four names of the target colors. The location of stimuli on the screen was kept constant.



Figure 2: A typical set of stimuli in the extra words condition.

Procedure Instead of verbally naming the color of the stimulus as in the classical Stroop task, participants were instructed to select as fast as possible the response option that matched the color of the stimulus from the two options presented on the left and right sides of the stimulus by pressing a key for each option. The session started with a short computer-guided tutorial that emphasized the correct response. Before the treatment block, participants were

instructed that the extra stimuli are to be ignored since they might interfere with the execution of the main task. During the task no feedback was provided. The task was self-paced.

Results

The data of one participant were excluded from analysis, because the latencies exceeded 2000 ms on average (this criterion had previously been used to exclude data from analysis in Miyake et al., 2000; Cho et al. 2006). As a consequence of this exclusion, each group was composed of 31 participants. A number of trials (4.52%) were excluded from analysis because they had very low (lower than 300 ms) or very high (higher than 2000 ms) latencies.

A first analysis was intended to check if our modified version of the manual Stroop task produced levels of interference comparable with the classical vocal Stroop task. Table 1 shows the mean latencies and accuracies obtained in our study. They are in the same range as the values obtained in studies that used the vocal version of the Stroop task (Cho et al., 2006; see also MacLeod, 1991, for a review).

Table 1: Accuracies and latencies by trial type.

| | Incongruent | Congruent | Neutral |
|-------------|-------------|-----------|---------|
| Accuracy | 0.92 | 0.99 | 0.98 |
| Latency (s) | 1.012 | 0.866 | 0.909 |

For the following analyses, only the correct trials were retained (96.09%). The data were submitted to a Linear Mixed Effects (LME) analysis with *latency* as a dependent variable, *group* (extra colors and extra words) and *block* (pre-treatment, treatment and post-treatment) as fixed factors, and a hierarchical grouping structure with *condition* (incongruent, congruent and neutral) nested in *participant*. This type of analysis is superior to the classical analysis of variance (ANOVA) in that it supports analysis of hierarchical effects and repeated measures, thus adequately handling data where observations are not independent (Garson, n.d.).

Figure 3 shows the mean latencies for the two groups across the three blocks. The LME analysis reveals a significant main effect of block ($t=-11.77, p=0.00$ and $t=-23.77, p=0.00$, respectively), that is, latency decreases from block 1 (pre-treatment) to block 2 (treatment) and further to block 3 (post-treatment). There are also two significant interactions between group and block ($t=-4.95, p=0.00$; and $t=2.35, p=0.02$, respectively). The main effect of group is not significant, although it is qualified by the two significant interactions between group and block. Thus, separate LME analyses show that, although both groups improve performance from block 1 to block 2, the group receiving extra words has a larger improvement ($t=-19.08, p=0.00$) than the group receiving extra colors ($t=-11.25, p=0.00$). Analogously, although both groups improve performance from block 2 to block 3, the extra colors group has a larger improvement ($t=-13.54, p=0.00$) than the extra words group ($t=-4.34, p=0.00$).

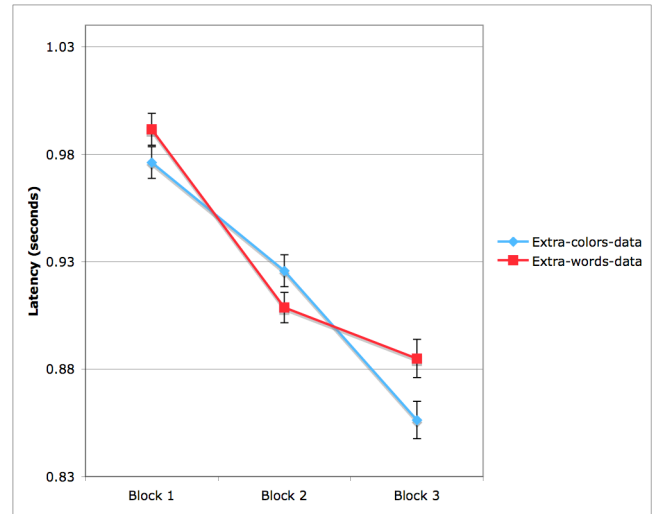


Figure 3: Task performance (mean latency) for each group across the three blocks. The vertical bars around the mean values depict 95% confidence intervals.

The pattern of results shown in Figure 1, particularly the interaction between block and group, suggests that the treatment in Block 2 (extra colors versus extra words) does have an influence on performance. However, in order to be properly interpreted, the effect of our manipulation must be separated from the main effect of block, that is, the general improvement in performance across the three blocks for both groups. Most likely, this is a practice effect that is independent of our manipulation and, thus, can be subtracted out.

Subtracting out the general practice effect was accomplished through the following procedure:

- A regression line was estimated for each participant with the aid of simple linear regression with latency as dependent measure and trial number as predictor. Thus, the value of the regression coefficient of the predictor (the slope) indicated the size of the practice effect in each participant, with larger negative values indicating steeper learning.
- For each participant and for each trial, an adjusted value of the latency was calculated as the difference between the actual latency and the product of trial number with the value of the learning slope of that participant.
- As a result of this procedure, each data point was discounted with an amount that corresponded to the general improvement in performance due to practice as estimated for each participant. The outcome of this procedure is depicted in Figure 4.

As shown in Figure 4, subtracting out the general practice effect has preserved the relative positions of the data points. Repeating the LME analyses presented above reveals similar t and p values for all effects except for the main effect of block, which can be taken as indication that the general practice effect has successfully been removed from the data.

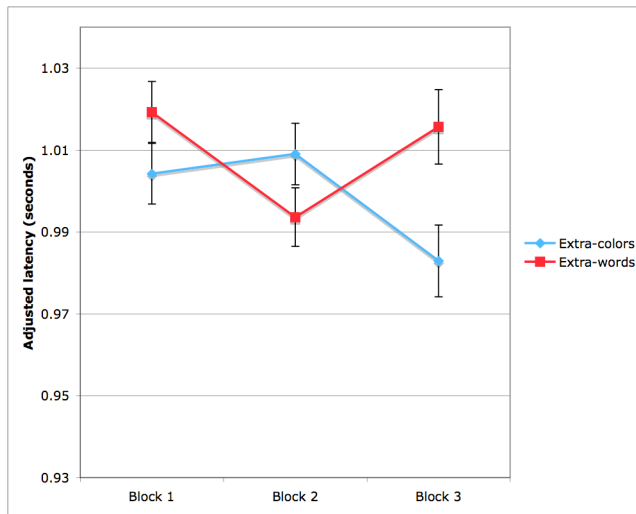


Figure 4: The result of removing the general practice effect. The remaining differences and interactions are likely to be determined solely by the treatment. Thus, the extra words group improves performance in block 2 (treatment) as compared with block 1 (pre-treatment), but this improvement does not extend to block 3 (post-treatment). The extra colors group improves performance only in block 3 (post-treatment). The vertical bars around the mean values depict 95% confidence intervals.

The adjusted latencies allow us to better understand the effect of our manipulation. The transformed data were submitted to separate Linear Mixed Effects (LME) analyses for each group with *adjusted latency* as a dependent variable, *block* (pairs of pre-treatment, treatment and post-treatment) as fixed factor, and a hierarchical grouping structure with *condition* (incongruent, congruent and neutral) nested in *participant*.

For the extra words group, there is a clear improvement in performance (i.e., decrease in latency) when extra words are added ($t=-6.49, p=0.00$). When extra words are removed (block 3) performance returns to its pre-treatment value (block 1) and this change is also significant ($t=4.93, p=0.00$).

For the extra colors group, there is no significant change in performance when extra colors are added (block 2), but there is a significant improvement in performance (i.e., decrease in latency) when the extra colors are removed (block 3). Latency in block 3 (post-treatment) is significantly lower than in block 1 ($t=-4.69, p=0.00$) and block 2 ($t=-5.01, p=0.00$).

Discussion

Adding extra stimuli to a task that already has a high level of interference would be expected to decrease performance by increasing information load and interference. This expectation has been disconfirmed by the Stroop dilution effect, which has been documented by several authors and replicated in our experiment. According to the existing

theoretical accounts, Stroop dilution seems to be a direct consequence of the fact that the extra stimuli are of the same kind as the distractor (i.e., words). They compete with the distractor for perceptual or attentional resources. As a consequence, the distractor is degraded and interferes less with the target. According to this account, if the extra stimuli were of the same kind as the target (i.e., colors), one would expect a decrease in performance, that is, a reverse of the dilution effect (i.e., increase in latency). This was our hypothesis for the extra colors group, and it was not supported by the data. The increase in latency in the treatment condition for the extra colors group was insignificant.

What is the reason for this unexpected result? Assuming that the existing theoretical accounts of the Stroop dilution effect were valid, the color dimension of the Stroop stimulus would compete with the extra colors and become harder to identify. Since we have no good reasons to question the existing theoretical accounts, we assume that the target has indeed become harder to identify but the cognitive system has reacted by allocating more resources to the target identification process. Since the target was already under the focus of top-down control (as the Stroop task requires), we assume that the system reacts by boosting its top-down control signal. It is a known property of our top-down control system to be able to modulate other perceptual and cognitive processes (Gazzaley, Cooney, McEvoy, Knight, & D'Esposito, 2005). The net effect of these two influences is zero, that is, performance does not change as compared to the pre-treatment condition.

More convincing evidence that a control boost has indeed occurred comes from the post-treatment block. Performance of the two groups goes in opposite directions. For the extra words group, performance goes back to its level at pre-treatment. This is in line with the existing accounts of the Stroop dilution effect, proving that it is indeed a pure bottom-up effect – it lasts as long as the extra words are displayed. For the extra colors group, performance goes to its highest level, that is, latency significantly decreases below its pre-treatment level. Given that the general learning effect had been removed, there would be no reason for such a significant change in performance at post-treatment unless something happened at treatment, presumably a boost in top-down control.

Computational Model

Building a computational model of the top-down modulation of the Stroop dilution effect is a useful endeavor; it helps us refine the theories on cognitive control and it generates questions and predictions for further research. The model we present here was developed with the aid of the latest version of the ACT-R¹ cognitive

¹ Adaptive Control of Thought – Rational. The ACT-R6 modeling software is available at <http://act-r.psy.cmu.edu/>. The Lisp code of the model presented here is available for download at <http://www.andrew.cmu.edu/user/ijuvina/Publications.htm>.

architecture (Anderson, 2007). ACT-R is a hybrid (symbolic and sub-symbolic) cognitive architecture used to develop cognitive models of various tasks. The architecture is composed of specialized modules (vision, memory, motor, etc.) coordinated by productions rules. The symbolic elements of the architecture (procedural rules, declarative memories) have associated sub-symbolic quantities (activations, utilities) that govern their availability and their manifestation in model's behavior.

The focus of our modeling efforts is on the top-down modulation of Stroop dilution. For this reason, other aspects of the Stroop task are treated very briefly. Whenever possible, we have used modeling ideas and techniques that have been proposed by other authors (Altmann & Davidson, 2001; Cohen, Dunbar, & McClelland, 1990; Herd, Banich, & O'Reilly, 2006; Lovett, 2005; Roelofs, 2003).

The following is a verbal trace of the model behavior at pre-treatment. The model perceives the stimulus and temporarily represents its color and word dimensions. These representations spread activation toward their associated words and concepts in memory. For example, if the stimulus is the word "red" in blue color (incongruent condition), the word "red" and the concept of blueness will be activated in memory. In order to reflect the fact that humans have higher practice at reading words than at naming colors, word representations have higher strengths of association with their corresponding memory elements than color representations. A consequence of this in our example is that the word "red" gets retrieved as a potential response. At this point, if memory retrieval were sufficient for performing an action, the model would commit an error, responding *red* instead of *blue*. However, the behavior of an ACT-R model is guided not only by perception and memory retrievals but also by firing of production rules of the kind "if condition, then action." In this case, a production rule detects the wrong retrieval and requests a new retrieval directed at the right color concept. The same mechanism of detecting a wrong retrieval is implemented in other models of the Stroop task (Altmann & Davidson, 2001; Lovett, 2005). Since memory retrievals take time, responses to incongruent stimuli take longer than responses to neutral stimuli. In the congruent condition, both representations spread activation toward the same element in memory, thus increasing its activation and speeding up its retrieval. In addition, for congruent stimuli, the first retrieval is sufficient for generating a correct response, even when it is guided solely by the word dimension of the stimulus.

When the extra stimuli are added (treatment block), the model perceives and represents them in a short-term storage structure from where they can spread inhibition toward memory elements of the same kind. When the extra stimuli are words, they inhibit the words in memory, including the word corresponding to the distractor. As a result, the probability that the distractor word will be retrieved decreases. Since the target will be retrieved more often, performance increases. When the extra stimuli are colors they inhibit the color concepts in memory, including the

concept corresponding to the target color. An inhibited color concept would be retrieved more slowly, potentially causing a decrease in performance. However, the retrieval of the target color concept is not slowed down by the extra colors because a task control unit boosts its activation. Thus, the task control unit compensates for the decrease in activation of the target color concept caused by lateral inhibition from the extra colors.

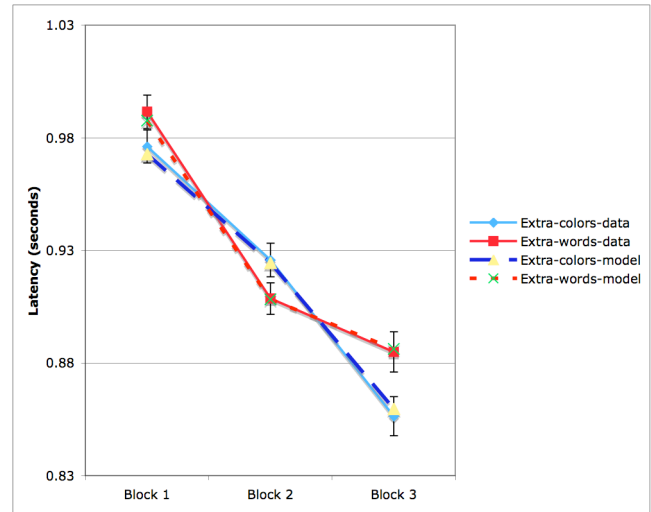


Figure 5: Fit of the computational model to the data. The vertical bars around the mean values depict 95% confidence intervals for the data.

When the extra stimuli are removed (post-treatment), their inhibitory effect on memory elements is discontinued. In the extra words condition, the model goes back to its pre-treatment behavior. In the extra colors condition, the model preserves the boost in activation of the target color concept by the task control unit that was initiated in the treatment block. Figure 5 shows the fit of the model to the empirical data. In order to simulate the general learning effect that was observed in the data, parameters controlling the speed of perceptual encoding and memory retrieval have been varied across blocks and separately for each group.

Although there is no space here to show it graphically, the model accounts reasonably well for the latencies of incongruent, congruent and neutral conditions across blocks for each group ($R^2=0.82$, mean deviation=0.038 seconds for the extra colors group; $R^2=0.85$, mean deviation=0.029 seconds for the extra words group).

General Discussion and Conclusion

Although they differ on details, the two existing accounts of the Stroop dilution effect (visual interference and attentional capture) postulate the same general principle: stimuli of the same type tend to compete against each other allowing a single stimulus of a different type to be identified faster. This principle is known in the visual search field as the "pop-out effect" (Desimone & Duncan, 1995). We have modeled the Stroop dilution effect as a lateral inhibition

process that implements the general principle of competition among stimuli of the same type. Different implementations of the same principle are found in ACT-R models of visual salience (Tamborello & Byrne, 2007) and memory retrieval (Van Maanen & Van Rijn, 2007).

Lateral inhibition would predict a reverse Stroop dilution effect when color patches instead of words are displayed as extra stimuli. Our data did not support this prediction. In order to reconcile the dilution from extra words with the lack of inverse dilution from extra colors, we postulated a modulating effect of top-down control. The essential difference between the extra words and extra colors conditions is that the former affects the distractor and the latter affects the target. The target is under the focus of top-down control, and it is presumably protected against lateral inhibition. Moreover, evidence from the post-treatment block suggests that a boost in top-down control has occurred at treatment. We model this boost in top-down control by an increase in activation spread from a task control unit that favors the target dimension against the distractor dimension of the stimulus.

What triggers such a boost in top-down control? We do not have yet a good theoretical and computational solution for this problem. Perhaps the control boost is triggered by the negative impact that the extra stimuli have on performance. Our model should be able to assess its performance on-the-fly and adjust its parameters accordingly. Why does the boost last after the harmful extra stimuli have been removed? How long does it last? These questions are being addressed in our current research and modeling efforts.

Acknowledgments

This research was supported by Office of Naval Research grant N00014-08-10541 and Air Force Office of Scientific Research grant FA95500710359. We thank Shahla Marie Hosseini for her comments on an earlier draft of this paper.

References

- Altmann, E. M., & Davidson, D. J. (2001). An integrative approach to Stroop: Combining a language model and a unified cognitive theory. *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, (pp. 21-26). Hillsdale, NJ.
- Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* New York: Oxford University Press.
- Brown, T. L., Roos-Gilbert, L., & Carr, T. H. (1995). Automaticity and word perception: Evidence from Stroop and Stroop Dilution Effects. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 21, 1395-1411.
- Cho, Y. S., Lien, M.-C., & Proctor, R. W. (2006). Stroop Dilution Depends on the Nature of the Color Carrier but Not on Its Location. *Journal of Experimental Psychology: Human Perception and Performance*, 32(4), 826-839.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97, 332-361.
- Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, 18, 193-222.
- Garson, G. D. (n.d.). Linear Mixed Models. In *Statnotes: Topics in Multivariate Analysis*: Retrieved 05/21/2008 from <http://www2.chass.ncsu.edu/garson/pa765/statnote.htm>.
- Gazzaley, A., Cooney, J. W., McEvoy, K., Knight, R. T., & D'Esposito, M. (2005). Top-down enhancement and suppression of the magnitude and speed of neural activity. *Journal of Cognitive Neuroscience*, 17(3), 507-517.
- Herd, S. A., Banich, M. T., & O'Reilly, R. C. (2006). Neural Mechanisms of Cognitive Control: An Integrative Model of Stroop Task Performance and fMRI Data. *J. Cognitive Neuroscience*, 18(1), 22-32.
- Kahneman, D., & Chajczyk, D. (1983). Tests of the automaticity of reading: Dilution of Stroop effects by color-irrelevant stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, 9(4), 497-509.
- Kornblum, S. (1994). The way irrelevant dimensions are processed depends on what they overlap with: The case of Stroop- and Simon-like stimuli. *Psychological Research*, 56(3), 130-135.
- Lovett, M. C. (2005). A Strategy-Based Interpretation of Stroop. *Cognitive Science*(29), 493-524.
- MacLeod, C. M. (1991). Half a Century of Research on the Stroop Effect: An Integrative Review. *Psychological Bulletin*, 109(2), 163-203.
- MacLeod, C. M., & Bors, D. A. (2002). Presenting two color words on a single Stroop trial: Evidence for joint influence, not capture. *Memory and Cognition*, 30(5), 789-797.
- Mitterer, H., La Heij, W., & Van der Heijden, A. H. C. (2003). Stroop dilution but not word-processing dilution: Evidence for attention capture. *Psychological Research*, 67, 30-42.
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology*, 41, 49-100.
- Roelofs, A. P. A. (2003). Goal-referenced selection of verbal action: Modeling attentional control in the Stroop task *Psychological review*, 110, 88-124.
- Tamborello, F. P., & Byrne, M. D. (2007). Adaptive but non-optimal visual search behavior in highlighted displays. *Journal of Cognitive Systems Research*, 8, 182-191.
- Van Maanen, L., & Van Rijn, H. (2007). An Accumulator Model of Semantic Interference. *Cognitive Systems Research*, 8(3), 174-181.