Altmann, E. M. & Davidson, D. J. (2001). *Proceedings of the 23rd annual meeting of the Cognitive Science Society* (pp. 21-26), Hillsdale, NJ: Erlbaum.

An Integrative Approach to Stroop: Combining a Language Model and a Unified Cognitive Theory

Erik M. Altmann (ema@msu.edu)
Department of Psychology
Michigan State University

Douglas J. Davidson (doug@eyelab.msu.edu)
Department of Psychology
University of Illinois at Urbana-Champaign

Abstract

The rich empirical puzzle of the Stroop effect has traditionally been approached with narrowly focused and somewhat atheoretical models. A recent exception is a simulation model based on the WEAVER++ language theory. The present model, WACT, combines components of WEAVER++ with the memory and control processes of the ACT-R cognitive theory. WACT accounts for the time course of inhibition from incongruent word distractors, facilitation from congruent word distractors, the lack of effect of color distractors, and the semantic gradient in inhibition. WACT goes beyond WEAVER++ to account for Stroop performance errors as well as latencies, and its implementation in a unified cognitive theory opens doors to broader coverage of Stroop phenomena than standalone models are likely to attain. Documented and executable code for WACT is available for inspection and comment at **www.msu.edu/~ema/stroop**.

Introduction

The Stroop effect is the mental confusion (and its behavioral consequences) induced when a word such as green is printed in a color such as red and the task is to name the color (red, in this case). Word meaning (green, in this case) seems to be processed automatically, in some sense, causing it to interfere with the color-naming task. Thus, the system may think green even though it sees red, because it can't stop itself from reading the word.

The rich pool of data on the Stroop effect (see MacLeod, 1991) has to date been approached with relatively lean cognitive theory. For example, the dominant simulation models remain the connectionist models of Cohen, McClelland, and Dunbar (1990) and Phaf, Van der Heijden, and Hudson (1990). The former model shows that Stroop phenomena can be simulated with simple information-processing units appropriately wired together. However, it makes no obvious contact with other cognitive theory – there are no identifiable linguistic or perceptual constraints, for example. Also, the model fails to capture the time course of inhibition, in which inhibition falls off gradually as the distractor occurs further ahead of the target (Glaser & Glaser, 1982; Glaser & Glaser, 1989; Sugg & McDonald, 1994). Indeed, simulated interference increases monotonically with temporal separation (Cohen et al., 1990, Fig. 7), suggesting basic flaws in the model's representation. The SLAM model (Phaf et al., 1990) is embedded in a theory of visual attention, but says little about the role of memory and executive control, and fails to capture the time course of inhibition (their Fig. 14a) and the asymmetry of reading and naming (their Fig. 14b).

Our approach to modeling Stroop effects is to integrate existing theory from other cognitive domains. Our model adopts mechanisms of the WEAVER++ language theory (Roelofs, 2000c), which explains Stroop phenomena in terms of competing lemmas (syntactic properties of words; Roelofs, 2000a, 2000b). These linguistic mechanisms are integrated into the ACT-R cognitive theory (Anderson & Lebiere, 1998), which specifies memory and executive-control mechanisms. The resulting model, which we refer to as WACT, goes beyond WEAVER++ to account for errors as well as latencies, and benefits from its embedding in ACT-R in terms of potential extensions to other phenomena. ACT-R suggests how automaticity of the dominant Stroop task might develop (MacLeod & Dunbar, 1988), and implements a theory of perceptual, motor, and cognitive constraints (Byrne & Anderson, in press) that could integrate a diverse range of Stroop effects into one model.

We begin by describing the effect to be explained – the time course of Stroop inhibition, in which latency and errors increase as distractor onset approaches target onset (Glaser & Glaser, 1989). We then describe WACT and its account of these effects, as well as its account of Stroop facilitation, a semantic gradient, and the non-effect of color distractors. In the discussion, we examine WACT's limitations and some possible extensions suggested by ACT-R.

The Time Course of Stroop Inhibition

Figure 1 illustrates the Stroop effect of primary interest here. The empirical data (solid lines) are from Experiment 1 of Glaser and Glaser (1989), in which a word and a color are shown with some temporal separation. Of interest here is the case in which the word (appearing first) is the distractor and the color patch (appearing second) is the target (the stimulus to which the participant responds). Thus, the word green might precede the color red by 100 msec. This temporal difference is the stimulus onset asynchrony (SOA). By convention, SOA is negative when the distractor precedes the target.

The latency difference measure in Figure 1 is derived by subtracting neutral latencies from incongruent and congruent latencies. On neutral trials, the distractor is a stimulus that consists of letters but is not a color word (e.g., xxxx). On incongruent trials, the distractor is a color word whose meaning conflicts with the color patch (e.g, green and red). On congruent

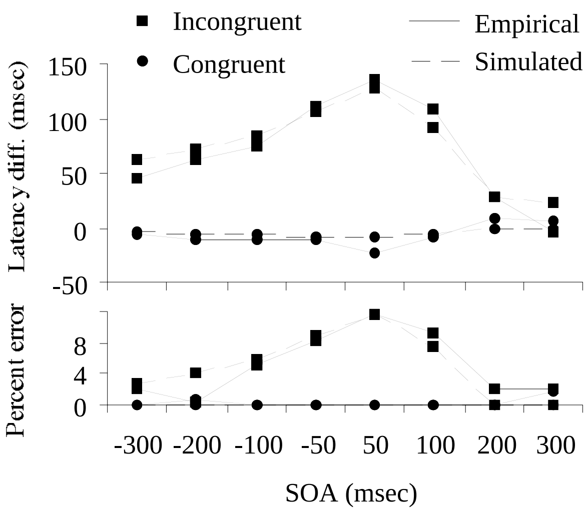

Figure 1: Stroop inhibition and facilitation. Latency difference is Incongruent/Congruent minus a neutral condition (see text). Stimulus onset asynchrony (SOA) is target onset minus distractor onset. Empirical data are from Glaser and Glaser (1989), Exp. 1, and simulated data are from WACT.

trials, the distractor is a color word whose meaning matches the color patch (e.g., red and red). In all three kinds of trials, the target stimulus is the color patch.

In Figure 1, the upper curves in each panel (square markers) are from incongruent trials, in which the distractor interferes with the target. As is typical, interference is greatest (i.e., the latency difference is greatest) when SOA is near zero – when target and distractor appear at roughly the same time. Note that even if the distractor occurs slightly after the target (e.g., an SOA of 50 msec), it still causes substantial interference. The lower curves in each panel (round markers) are from congruent trials, in which the "distractor" actually slightly facilitates performance.

The error measure in Figure 1 is the raw percentage of substitution errors, or trials on which the wrong response word was given. (With no detectable facilitation from congruence, there is no need for a difference measure.) Only incongruent naming is particularly error prone, and there, as with latency, interference is greatest at near-zero SOA.

The WACT Model

Long-term lexical knowledge in WACT is organized in a multi-layer declarative network, as shown in Figure 2. The top level of this network contains semantic nodes, or concepts. Below this is the lemma layer, which contains syntactic information (lemmas) crucial for fitting a word into the grammatical organization of a phrase or sentence. Below the lemma layer is the form layer, which contains the information necessary to produce an individual word.

In WEAVER++ and WACT, interference and facilitation occur at the lemma layer. Word stimuli have direct access to their lemmas, whereas non-verbal stimuli like colors gain access only indirectly, via concepts (Figure 3). The direct link from a word stimulus to its lemma is the route by which words trigger automatic language processing. The benefit of this automaticity is efficiency, helping the system to meet immediacy constraints on comprehension (Just & Carpenter, 1987). The cost of this automaticity, on our view, is that it leaves behind traces of information that can interfere with subsequent tasks like color naming.

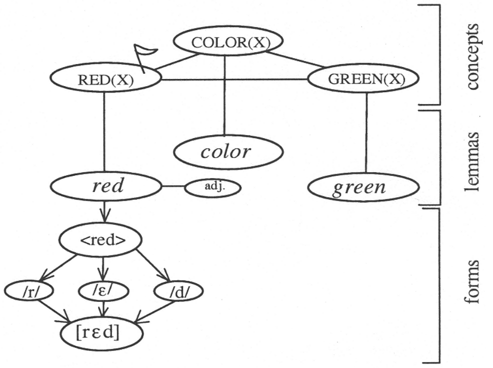

Figure 2: WEAVER++ long-term lexical knowledge (from Roelofs, 2000a). WACT represents concepts, lemmas, and the top layer of forms (e.g., <red>).

A word stimulus automatically activates the corresponding lemma. We assume that the purpose of this activation, relative to language comprehension, is to facilitate parsing of subsequent tokens. For example, the stimulus "and" might establish an expectation for a subsequent conjunct, by virtue of causing the "and" lemma to be active when the conjunct arrives. Stroop interference (and facilitation) are caused by the activation of a lemma for the distractor word. If this distractor lemma is incongruent with the target lemma, it produces a form of response competition when the system tries to retrieve the target lemma. On the other hand, if the distractor lemma is congruent with the target lemma, then the system benefits from intrusions of the distractor lemma.

In WEAVER++ and WACT, latency to retrieve a target depends on the target's activation relative to distractors – the more active the target is relative to its distractors, the quicker it is retrieved. Relative activation is a common way to formalize interference (Baddeley & Hitch, 1993; Luce, 1959; Murdock, 1985; Neath, 1993). In ACT-R the formulation is

$$P_i = \frac{e^{A_i/s}}{\sum_j e^{A_j/s}} \qquad \text{Equation (1)}$$

where P_i is the probability of retrieving item i on a given attempt given the j items in memory at the time. A_i is the activation of i, and s is system noise.

Importantly, WACT (unlike WEAVER++) specifies the processing consequences of retrieving the wrong item on a given attempt. WEAVER++ predicts latency simply by scaling relative activation. WACT actually uses the retrieved lemma to decide how to respond, so

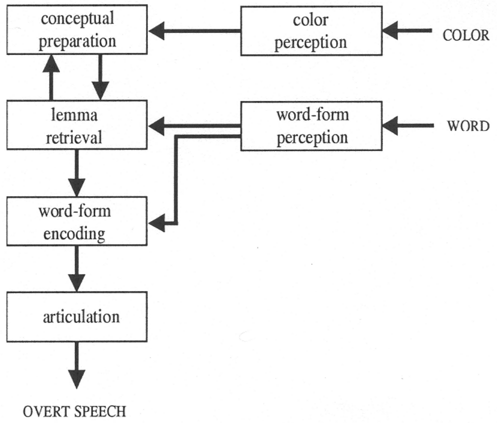

Figure 3: Stages of linguistic processing in WEAVER (from Roelofs, 2000a) and in WACT.

there is the possibility of an incorrect lemma retrieval causing a substitution error.

The lemma-retrieval process is shown in Figure 4. The figure is taken from Murdock (1974), but retrieve-decide models like this are common (e.g., Anderson & Bower, 1972; Kintsch, 1970; Watkins & Gardiner, 1979) and map naturally to ACT-R memory-retrieval productions. Such a process also explains tip-of-the-tongue effects, in which subjects appear to monitor correctness of retrievals (e.g., Levelt, 1983). Probed with a word stimulus, the system tries to retrieve the corresponding lemma. Any retrieved lemma is evaluated (in the decision process) by comparing the current concept to conceptual cues retrieved with the lemma. In case of a mismatch, the system tries again. Eventually the system retrieves a lemma it considers correct, or runs out of time. Either way, the last lemma retrieved is the basis for form retrieval (the next stage of language production; Figure 3). If the last lemma retrieved is incorrect, then form processing begins with the wrong input, likely causing a performance error. Thus, the distractor lemma interferes with the target lemma by affecting the duration and potentially the output of the lemma-retrieval process.

In WACT, the amount of interference caused by a distractor lemma depends on its activation, which in turn depends on the time elapsed since the distractor stimulus was presented. Activation in ACT-R is

$$A = \ln\left(\frac{2n}{\sqrt{T}}\right) \qquad \text{Equation (2)}$$

where n is the number of times the item has been retrieved and T is the length of the item's lifetime. For a distractor lemma, $T=n=1$ when the distractor stimulus is presented. After that n remains essentially constant, but T increases throughout the trial, causing A to decrease (decay). Thus, the more time elapses between distractor and target, the more the distractor lemma decays and the less it intrudes on the target lemma.

To illustrate, Figure 5 shows activation values from Monte Carlo simulations of naming trials at various SOAs. The top curve is the activation of long-term lemma representations. For these representations, n and T are both large, so activation is stable over short intervals. The bottom curve is activation of the distractor lemma. At large negative SOAs, the distractor lemma decays by the time the target appears, but at near-zero SOAs its activation is close to that of long-term lemmas. Syntactic representations are known to decay rapidly (e.g., Potter & Lombardi, 1990; Sachs, 1967), and here this rapid decay explains the time course of Stroop interference. On an incongruent trial, the system must retrieve a long-term lemma in order to process the target correctly, and this retrieval is faster and more accurate if the distractor lemma has decayed.

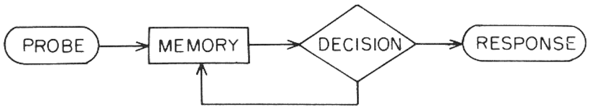

Figure 4: The WACT retrieve-decide process for lexical retrieval (from Murdock, 1974).

Two other comments on Figure 5 are in order. First, the activation values reflect the sum of two sources: base-level activation (A in Equation 2) and associative activation from cues like the current stimulus; these are the two possible sources of activation in ACT-R (Anderson & Lebiere, 1998). The sum of these two sources is the activation factor in the item's likelihood of being retrieved (A in Equation 1). That said, the scale on the ordinate of Figure 5 is arbitrary, because relative activation, not absolute activation, is what governs retrieval probability in WACT. A second point is that the lower curve ends at SOA 100. This means only that at SOA 200 and 300, the target lemma was always retrieved before the distractor appeared. When the target is retrieved in time to avoid interference from the distractor, the model implementation simply skips the step of activating the distractor lemma, as the trial is functionally over by then.

Comparing WACT to Data

WACT behavioral data, from the same simulations that produced the activations in Figure 5, are presented in Figure 1 (dashed lines). The fits are quite respectable: $r^2=.98$ and RMSD=11.0 for latencies, $r^2=.94$ and RMSD=1.4 for errors. The model clearly captures the peak in inhibition near zero SOA and the gradual falling-off (leftward) as the distractor word is presented

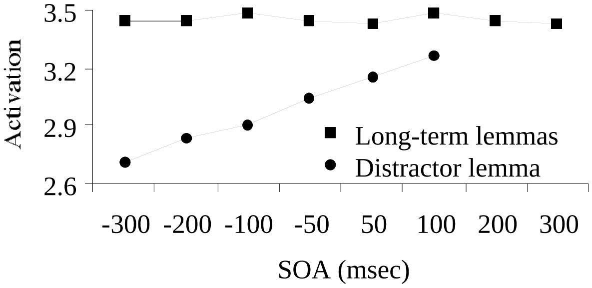

Figure 5: The time course of activation in WACT, showing decay of distractor lemma at negative SOAs.

further ahead of the target color. Inhibition also falls off sharply for positive SOAs (rightward), where the target is usually fully processed and the response formulated by the time the distractor appears.

WACT also captures Stroop facilitation, a common though relatively small effect. In the model, facilitation arises when a congruent distractor lemma is correct, in which case the "distractor" lemma is indistinguishable from the target. Functionally, activating a congruent distractor lemma is equivalent to a slight increase in activation of the target lemma.

An important "non-effect" captured by WACT is that a color distractor has no effect on word reading, either inhibitory or facilitative. In the model, color distractors have no effect because color stimuli are not processed ballistically, as words are. In response to a color stimulus, the system does not automatically activate a lemma; in terms of cognitive economy, there is no reason to process an arbitrary stimulus verbally unless the task requires it. Thus, a color distractor leaves behind no activated lemmas to interfere with the subsequent target word. This account is a point of distinction between WACT and WEAVER++, which accounts for the non-effect of colors not in terms of ballistic processing but in terms of levels of processing (Roelofs, 2000b; personal communication).

The Semantic Gradient

Beyond time course effects, another important effect captured by WACT is the semantic gradient, in which a distractor like "lawn" inhibits the naming of a target color like red. This inhibition arises presumably because lawn primes some representation of green, which then conflicts with the response to red. This effect is important because it is one of a class of effects involving higher-level semantic relations among categories (Glaser & Dungelhoff, 1984; Roelofs, 2000c). It is also an opportunity to compare WACT and WEAVER++, and to address a point of contention in how best to model gradient effects in ACT-R.

In WACT, the semantic gradient arises from operation of the retrieve-decide process (Figure 4) at the concept level. Earlier, we described this process operating at the lemma level, but in fact the process operates at each of the levels of speech production (Figure 3), as befits a general process for detecting and correcting memory errors. Thus, just as with lemma retrieval, when the model needs a concept, interference from incorrect concepts will degrade performance.

To illustrate how the semantic gradient arises, suppose the distractor word is lawn and the target color is red. When WACT sees "lawn", the lawn concept is activated as a side effect of processing the lawn lemma. The corresponding assumption in WEAVER++ is that activation spreads from lemma to concept in parallel as it spreads from lemma to form (Figure 3). In WACT, the lawn concept cues related knowledge through semantic priming. Among the concepts related to lawn is green – which also belongs to the response set for the current task (i.e., sometimes the target color is green). The combination of priming from the lawn concept and priming from the task environment is enough to cause the green concept to intrude occasionally on the concept for the actual target color. That is, having processed lawn, the model may think green, even if it sees red. Relevant data (from Exp. 5 of Glaser & Glaser, 1989) appear in Figure 6. The semantic gradient is represented by the small, positive latency difference across SOAs, reflecting modest interference from distractors like lawn. (The small peak at SOA 50, which Glaser & Glaser, 1989, attribute to random variation, is unrelated to the semantic gradient.) WACT again follows the trend, with distractors like lawn causing some interference but not as much as distractors like green.

The WACT account of semantic gradients may be another point of distinction relative to WEAVER++. In ACT-R, activation spreads only one link from whatever cues are in the focus of attention. Thus, spreading activation over a distance of multiple links requires a sequential process of chained retrievals in which each retrieval brings the next cue into the focus of attention. In WEAVER++, by contrast, activation spreads uncontrolled throughout the lexical network. Though attenuated by distance (number of intervening links) from the activation source, this uncontrolled spreading seems to make WEAVER++ quite sensitive to representational assumptions. For example, current reports (Roelofs, 2000a, 2000b) suggest that activation from the lawn word would reach the green lemma (via the concepts lawn and green), causing conflict with the red lemma. The same reports suggest that activation from the lawn word would also reach the red lemma (via the concepts lawn, green, and red), compensating for the activation reaching the green lemma. Thus, the word lawn could produce inhibition, facilitation, or neither with naming the color red, depending on the relative strengths of the various associations involved.

The WACT account of semantic gradients is also important because it shows that such effects can be accommodated by ACT-R's core theoretical premises. ACT-R assumes that performance (including memory performance) adapts to the statistical structure of the environment (Anderson & Lebiere, 1998; Anderson &

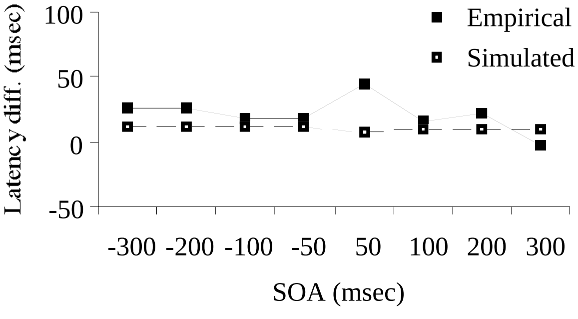

Figure 6: Inhibition from distractors like "lawn". Empirical data are from Glaser and Glaser (1989), Exp. 5, Cond. 2. Simulated data are from WACT.

Milson, 1989). Thus, WACT assumes that lawn and green concepts are associated in memory because they co-occur in the environment. Associative mechanisms also account for temporal gradients in order memory (Altmann, 2000). Nonetheless, the adequacy of such representations has been questioned, and ACT-R has come to incorporate a "partial matching" mechanism for fitting gradient data (e.g., Anderson & Matessa, 1997). The current work suggests that this mechanism, which has no clear motivation in terms of independent theoretical constraints, is best viewed as a simplifying assumption and not as a part of ACT-R theory proper.

Model Parameters

The parameters used to fit the data in Figures 1 and 6 were as follows. Activation noise (set at 0.33), or *s* in Equation 1, causes some retrieval attempts to produce the incorrect target. Encoding noise (0.0205) causes some stimuli to be encoded out of order at small SOAs. (Activation and encoding noise both index logistic variance; see Anderson & Lebiere, 1998.) The limit on retrieval attempts (3) affects how soon the retrieve-decide process gives up and outputs its last retrieval. Other parameters affect associative activation spreading from a cue. High strength (8.9 units of activation) applies to perceptual cues and medium strength (6.9 units) applies to mental cues. Low strength (4.9 units) applies to lawn as a cue for green, so is relevant only to the fit in Figure 6.

Discussion

Stroop effects are complex and diverse and it seems clear that broad coverage will elude us as long as we continue to approach them with standalone models. A comprehensive theory is required, in which the interactions of various cognitive subsystems can be simulated to investigate whether particular Stroop phenomena emerge as a natural consequence. Several directions indicated by the marriage of WEAVER++ and ACT-R in WACT are discussed below.

First, there is the question of Stroop development – how a process like reading becomes automatic enough to interfere with other tasks like color naming. MacLeod and Dunbar (1988, Experiment 3) demonstrated that this development can be induced through training. Their participants received extensive practice (daily, for a month) on associations between arbitrary shapes and color names. These associations eventually became automatic enough to interfere with color naming, which, before training, had been the more automatic task. Would ACT-R, as a unified theory that integrates learning mechanisms, allow WACT to be extended to develop automaticity?

A WACT account of automaticity would likely be grounded in ACT-R's utility-based theory of procedural skill acquisition. Procedural skills in ACT-R are represented as production rules that govern retrieval from declarative memory. Skill acquisition itself is represented in part as the acquisition of cost-benefit knowledge about individual production rules – the more a rule succeeds, the more it is preferred when the system has a choice. This mechanism has been used to account for set effects (Lovett, 1998) on the view that these are driven by frequency of rule use. In WACT, automaticity is represented by fixed settings of these cost-benefit parameters. That is, productions that read a word stimulus have high utility and thus are preferred to productions for less-used skills like color naming. It seems feasible and useful to extend WACT to simulate training data sets like that of MacLeod and Dunbar (1988). With such an extension in hand, one could assess its predictive value by manipulating system parameters like the rate at which skill acquisition takes place.

A second question concerns the relationship between Stroop phenomena and task switching. Stroop interference and executive control interact, in that switching to controlled tasks like color naming is easy, whereas switching to automatic tasks like reading is hard (Allport, Styles, & Hsieh, 1994). However, Stroop conflict is robust to task uncertainty – inhibition effects are largely unchanged when the task is determined dynamically on each trial by stimulus order (Glaser & Glaser, 1989). These disparate effects of task may inherit explanations from studies of task switching conducted within ACT-R (Altmann & Gray, in press; Sohn, Ursu, Anderson, Stenger, & Carter, 2000).

A third question concerns the relationship between Stroop phenomena and the psychological refractory period paradigm used to investigate perceptual, motor, and cognitive bottlenecks. The data in Figure 7, from which the latency-difference measure in Figure 1 was computed, hint at a bottleneck in Stroop processing. The slowing near SOA zero suggests a "jamming" of some kind (Meyer & Kieras, 1997) when stimuli appear close together in time. WEAVER++ accounts for this effect (Roelofs, personal communication) in terms of its activation dynamics. In contrast, ACT-R incorporates a structural and processing theory of bottleneck effects generally (Byrne & Anderson, in press). Extending WACT in this direction would take it well beyond current Stroop models, by integrating perception, action, language, memory, and executive control in one running model.

Finally, we hope to extend WACT to phrase

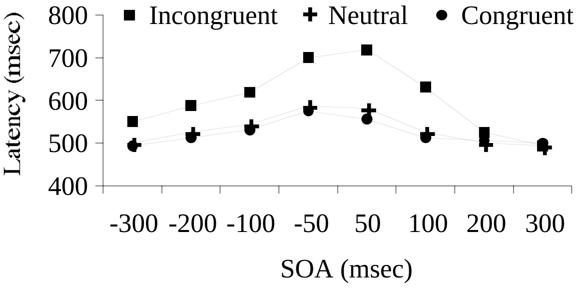

Figure 7: Latencies from which the difference scores in Figure 1 were computed, showing that processing slows near zero SOA, regardless of condition.

production, in particular the production of conjunctive phrases. When the task is to name the color and read the word, utterance onset depends on which response is to be given first. In particular, utterance onset is delayed when the color is to be named first (and the word meaning is incongruent). These data help to characterize planning in speech production, but also offer an opportunity to integrate Stroop phenomena more broadly with psycholinguistic theory.

Acknowledgements

Thanks to Tom Carr, Ardi Roelofs, and the conference reviewers for comments on the original draft.

References

Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In Umilta & Moscovitch (Eds.), *Attention and performance IV* (421-452). Cambridge: MIT Press.

Altmann, E. M. (2000). Memory in chains: A dual-code associative model of positional uncertainty. *Proc. 3rd international conference on cognitive modeling* (9-16). Veenendaal, NL: Universal Press.

Altmann, E. M., & Gray, W. D. (in press). Forgetting to remember: The functional relationship of decay and interference. *Psychological Science*.

Anderson, J. R. & Bower, G. (1972). Recognition and retrieval processes in free recall. *Psych. Rev., 79*, 97-123.

Anderson, J. R., & Lebiere, C. (Eds.). (1998). *The atomic components of thought*. Hillsdale: Erlbaum.

Anderson, J. R., & Matessa, M. (1997). A production system theory of serial memory. *Psych. Rev., 104*, 728-748.

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psych. Review, 96*, 703-719.

Baddeley, A. D., & Hitch, G. (1993). The recency effect: Implicit learning with explicit retrieval? *Memory & Cognition, 21*, 146-155.

Byrne, M. D., & Anderson, J. R. (in press). Serial modules in parallel: The psychology refractory period and perfect time-sharing. *Psychological Review*.

Cohen, J. D., McClelland, J. L., & Dunbar, K. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review, 97*, 332-361.

Glaser, M. O., & Glaser, W. R. (1982). Time course analysis of the Stroop phenomenon. *Journal of Experimental Psychology: HPP, 8*, 875-894.

Glaser, W. R., & Dungelhoff, F.-J. (1984). The time course of picture-word interference. *Journal of Experimental Psychology: HPP, 10*, 640-654.

Glaser, W. R., & Glaser, M. O. (1989). Context effects in Stroop-like word and picture processing. *Journal of Experimental Psychology: General, 118*(1), 13-42.

Just, M. A., & Carpenter, P. A. (1987). *The psychology of reading and language comprehension*. Boston: Allyn and Bacon.

Kintsch, W. (1970). Models for free recall and recognition. In D. A. Norman (Ed.), *Models of Human Memory*. New York: Academic Press.

Levelt, W. J. M. (1983). Monitoring and self-repair in speech. *Cognition*, *14*, 41-104.

Lovett, M. (1998). Choice. In Anderson & Lebiere (Eds.), *Atomic components of thought* (255-296). Hillsdale, NJ: Erlbaum.

Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.

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., & Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. *J. Exp. Psych.: LMC, 14*, 126-135.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance. *Psych. Rev., 104*, 3-65.

Murdock, B. B. (1974). *Human memory: Theory and data*. New York: Wiley.

Murdock, B. B. (1985). An analysis of the strength-latency relationship. *Mem. & Cog., 13*, 511-521.

Neath, I. (1993). Distinctiveness and serial position effects in recognition. *Mem. & Cog., 21*, 689-698.

Phaf, R. H., Van der Heijden, A. H. C., & Hudson, P. T. W. (1990). SLAM: A connectionist model for attention in visual selection tasks. *Cognitive Psychology, 22*, 273-341.

Potter, M. C., & Lombardi, L. (1990). Regeneration in the short-term recall of sentences. *Journal of Memory and Language, 29*, 633-654.

Roelofs, A. (2000a). Attention to action: Securing task-relevant control in spoken word production, *Proc. 22nd annual meeting of the Cognitive Science Society* (411-416). Mahwah, NJ: Erlbaum.

Roelofs, A. (2000b). Control of language: A computational account of the Stroop asymmetry. *Proc. 3rd int'l conference on cognitive modeling* (234-241). Veenendaal, NL: Universal Press.

Roelofs, A. (2000c). WEAVER++ and other computational models of lemma retrieval and word-form encoding. In L. Wheeldon (Ed.), *Aspects of language production*. Philadelphia: Psychology Press.

Sachs, J. S. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception and Psychophysics, 2*, 437-442.

Sohn, M.-H., Ursu, S., Anderson, J. R., Stenger, V. A., & Carter, C. S. (2000). The role of prefrontal cortex and posterior parietal cortex in task switching. *Proc. National Academy of Sciences, 97*(24), 13448-13453.

Sugg, M. J., & McDonald, J. E. (1994). Time course of inhibition in color-response and word-response versions of the Stroop task. *Journal of Experimental Psychology: HPP, 20*, 647-675.

Watkins, M. & Gardiner, J. (1979). An appreciation of generate-recognize theory of recall. *Journal of Verbal Learning and Verbal Behavior, 18*, 687-704.