

Attention to Action: Securing Task-Relevant Control in Spoken Word Production

Ardi Roelofs (ardi@mpi.nl)

Max Planck Institute for Psycholinguistics
Wundtlaan 1, 6525 XD, Nijmegen, The Netherlands

Abstract

The Stroop phenomenon is the finding that color naming is inhibited by incongruent color words but word reading not by incongruent colors. When the stimulus onset asynchrony (SOA) is manipulated, maximal inhibition of incongruent words on color naming is obtained when the words are presented within 100 msec of the colors, whereas facilitation of preexposed congruent words is constant. These findings are obtained both with and without task certainty. Whereas existing models explain the basic Stroop effects, they fail to account for the time course findings and for performance under task uncertainty. In this paper, I extend and apply the WEAVER++ model of spoken word production (Roelofs, 1992, 1993, 1997c; Levelt, Roelofs, & Meyer, 1999) to performance on the Stroop task and show that the model accounts for the key findings.

Introduction

Performance on the Stroop task is of direct relevance to theories of language production and comprehension. The basic modes of language use, namely speaking, listening, reading, and writing, all seem to make use of overlapping sets of basic processing components (e.g., Caplan, 1992; Levelt, 1989; Shallice, 1988). Whereas language perception occurs automatically, hearing or reading a word does not automatically lead to its production but this is under the control of a language user. Similarly, seeing an object does not automatically lead to the naming of it. Furthermore, words do not occur in isolation, but are typically part of a spoken discourse, a text on a page, or appear on objects in the real world. This points to the need to deal with the issue of selectivity. It is generally assumed that performance on the Stroop task can provide evidence on how language is controlled, that is, how a speaker secures task-relevant control over the basic language processes underlying naming and oral reading (e.g., Allport, 1993).

Since Stroop's (1935) experiments in the 1930s, over 700 articles have appeared using his task (reviewed by MacLeod, 1991), which established the following basic empirical picture. Color naming is inhibited by incongruent color words, but word reading not by incongruent color patches. For example, saying "red" to a red color patch on which the word "green" is superimposed proceeds slower than saying "red" in a control condition consisting of a string of Xs. The stimulus onset asynchrony (SOA) between color patch and word has an important effect. Maximal inhibition of incongruent words on color naming is obtained when the words are presented within 100 msec of the colors (e.g., Glaser & Glaser, 1982). Preexposed congruent words yield facilitation, which is constant over SOAs. Whereas color naming is affected by incongruent words,

reading aloud words (e.g., "red") is not influenced by incongruent colors (e.g., green). The asymmetry in effect is not due to a difference in processing speed between words and colors (i.e., reading is some 200 msec faster than color naming), as evident from manipulating the SOA. When a color patch is presented 300 or 400 msec before the word to be read, still no effect of incongruent colors is obtained (e.g., Glaser & Glaser, 1982).




Figure 1: Time course of the Stroop effects (relative to control) in color naming under task uncertainty as measured by Glaser and Glaser (1989): ■ = incongruent, ▲ = congruent

In a standard Stroop experiment, participants are certain about what task to perform. Typically, one group of participants is asked to name the color and to ignore the word, and another group of participants is asked to read aloud the word and to ignore the color patch. Thus, trials are blocked by task. However, in examining the effect of task uncertainty, Glaser and Glaser (1989) asked a single group of participants to perform both tasks. They instructed their participants to respond to the second stimulus component (in the condition with negative SOAs) or the first component (in the condition with positive SOAs). Words had to be read aloud and colors had to be named. Participants can perform this task up to differences in presentation time of 50 msec. The experiment was run with SOAs of -300, -200, -100, -50, 50, 100, 200, and 300 msec (a minus sign indicates preexposure of the irrelevant stimulus). With task uncertainty, Glaser and Glaser obtained the normal patterns of inhibition and facilitation observed with task certainty, for example, with the instruction to name the color and ignore the word. Inhibition increased when the SOA be-

came less negative, peaking at SOAs between -100 and 100 msec. And there was a flat pattern of facilitation at the negative SOAs. Figure 1 shows the SOA curves.

The Stroop phenomenon is not restricted to naming colors and reading color words but appears in many other verbal domains. For example, numerals interfere with the naming of numerosity (e.g., saying “two” to two 6s), but there is no reverse effect (e.g., Flowers, Warner, & Polansky, 1979). Alternating between tasks that exhibit the Stroop conflict (i.e., between color naming and numerosity naming trials) does not yield a greater task switch cost than alternating between tasks that do not yield the Stroop conflict (i.e., between word reading and numeral reading trials), as observed by Allport, Styles, and Hsieh (1994).

The Challenge Posed by Task Uncertainty

The Stroop Conflict as Task Conflict

Recent proposals in the literature (e.g., Rogers & Monsell, 1995) have suggested that the essence of the Stroop conflict is competition between tasks *per se* (i.e., word reading and color naming). The Stroop conflict is explained as inhibition of the “weaker” color naming task by the supposedly “stronger” reading task, while the reverse does not occur. However, task competition fails to explain why the congruent condition (where the same tasks compete) yields facilitation. Furthermore, the fact that task uncertainty has no influence on the SOA patterns poses a challenge. Also, task competition would predict a greater cost for switching between conflict trials with color and numerosity naming than between nonconflict trials with word and numeral reading, but Allport et al. (1994) found no cost difference.

The findings on the time course of the effects and on task uncertainty also pose a challenge to models that do not conceive of the Stroop conflict as a task conflict *per se*. Whereas existing models account for the basic Stroop effects obtained with SOA = 0 msec, they fail to explain the time course findings and they cannot cope with task uncertainty.

The Model of Cohen, Dunbar, and McClelland (1990)

Among the most influential models of the Stroop phenomenon is the connectionist model of Cohen, Dunbar, and McClelland (1990). The model assumes a feedforward network with parallel reading and color naming pathways, which differ in strength. Task relevant control is achieved in the model by task nodes for color naming and reading. These task nodes provide extra input to the color and reading pathway, depending on the task. Each response node in the network is connected with an evidence accumulator. Before the beginning of a simulated Stroop trial, all evidence accumulators are set to zero. A task node is activated and the model is run until the activation of all nodes stabilizes. This allows the system to settle into a “ready state” for the task. Next, the components of a Stroop stimulus are presented with the appropriate SOA. A response is selected when one of the accumulators exceeds a fixed response threshold.

The model of Cohen et al. does well in accounting for the basic Stroop effects obtained with SOA = 0 msec, but there are two major problems. First, as simulations by Cohen et al. (1990, p. 344) showed, the amount of evidence accumulated for the irrelevant stimulus is a positive function of its preexposure time. That is, more evidence is collected at more negative SOAs. Thus, the inhibition in the incongruent condition peaks at the most negative SOA and decreases when the SOA becomes less negative. Similarly, the amount of facilitation in the congruent condition peaks at the most negative SOA and decreases when the SOA becomes less negative. The problem is that these patterns are exactly contrary to the empirical results, where maximal impact of incongruent words is observed when the words appear within 100 msec of the colors and facilitation from preexposed congruent words is constant. The model also predicts a small Stroop effect at negative SOAs in reading aloud, contrary to the real data. The second major problem with the model is that it cannot handle task uncertainty. Before the beginning of a trial, a task node is activated and the model is run until the activation of all nodes stabilizes, which allows the system to settle into a ready state for the task. But with task uncertainty, the task is not known beforehand so that such task-dependent settling of activation is not possible.

The Model of Phaf, Van der Heijden, and Hudson (1990)

Another influential model of the Stroop phenomenon is the connectionist model of Phaf, Van der Heijden, and Hudson (1990), called SLAM (for SeLective Attention Model), which has been developed within the framework of Van der Heijden’s (1992) general theory of attention. The model assumes an interactive-activation network. Input nodes for colors are connected to corresponding hidden nodes for colors, which in their turn are linked to word output nodes. Input nodes for words are directly connected to these output nodes. Thus, unlike the model of Cohen et al. (1990), the model assumes asymmetrical pathways for reading and color naming. Processing occurs through activation spreading from color input via hidden to output nodes, and directly from word input to output nodes, whereby nodes change their activation with time in a continuous, nonlinear manner. There are excitatory links between nodes representing compatible information and there are inhibitory links between nodes standing for incompatible information. All nodes of a particular type within a layer inhibit each other. Selective attention to the color naming and reading tasks is achieved by adding extra external activation to all hidden color nodes for color naming and all output nodes for word reading. The task activation is given from trial onset onward. On each simulated trial, word and color input is given to the network and activation cycles around from one unstable pattern to another until a stable pattern of activation is reached. The excitatory and inhibitory connections push activation of the response nodes into one stable state depending on the inputs provided to the layer (e.g., color and task input). To choose one response or another, activation of the response layer is

input to a sampling and recovery procedure that stochastically favors the most highly activated response node.

The Phaf et al. (1990) model successfully accounts for the basic Stroop effects with SOA = 0 msec, but, again, there are the same two major problems. First, the model does not adequately account for the time course of the Stroop phenomenon that has been observed by Glaser and Glaser (1982) and others. As simulations by Phaf et al. (1990, p. 324) showed, the model predicts that the amount of inhibition of words in color naming does not vary with SOA but remains constant for negative SOAs, contrary to the empirical findings. The reason for predicting a constant SOA effect in color naming is that after perceiving the word, the system quickly settles into a stable state of activation for the response corresponding to the word. By definition, the stable state does not vary with time, and hence making the SOA more or less negative has no effect, until an SOA is used that is too short for the distractor to reach an attractor basin. Consequently, the amount of time it takes for the color name to overcome the inhibition from the word is constant. The second major problem is that the model cannot cope with task uncertainty.

In line with Phaf et al. (1990), Cohen and Huston (1994) discuss an attractor version of the model proposed by Cohen et al. (1990). The behavior of this model is similar to that of Phaf et al. (1990). The amount of inhibition at negative SOAs is constant (see Figure 18.11 of Cohen and Huston, 1994), contrary to the real data. And the new version of the model also cannot cope with task uncertainty.

An alternative to connectionist task control is provided by “production rule system” models (e.g., Anderson, 1983; Anderson & Lebiere, 1998). Below, I show that the WEAVER++ model of word production (Levelt et al., 1999; Roelofs, 1992, 1993, 1997c), which falls into this general class of model, accounts for the findings on task uncertainty. The relevant features of the model are: (1) words are retrieved by spreading activation and (2) task-relevant control is achieved by production rule application.

Control in the WEAVER++ Model

Planning Stages

In WEAVER++, naming a perceptual entity such as a color involves a number of processing stages, illustrated in Figure 2. First, there is the conceptual identification of the color based on perceptual input (e.g., red) and its designation as goal concept (i.e., RED(X)). Second, the lemma of the corresponding word is retrieved (i.e., *red*), in the Stroop literature often referred to as response selection (except that it involves here lemmas, which is new). A lemma is a representation of the syntactic properties of a word, crucial for its use in sentences (cf. Roelofs, Meyer, & Levelt, 1998). Third, the form of the word is encoded (i.e., [rɛd]), called response programming. Lemma retrieval and word-form encoding are discrete processes in that only the form of a *selected* lemma becomes activated and selected (Levelt, Schriefers, Vorberg, Pechmann, Meyer, & Haviga, 1991). And finally, the name is articulated, called response execution.

A perceived written word activates its lemma and its

output form in parallel. Oral reading is achieved by a shallow form-to-form route (e.g., from the orthographic form *red* to [rɛd]) or may involve an extra step of lemma retrieval (i.e., from *red* via *red* to [rɛd]), roughly corresponding to what is traditionally called the “semantic” route (e.g., Caplan, 1992; Shallice, 1988). I refer to Levelt et al. (1999) and Roelofs, Meyer, and Levelt (1996) for an extensive discussion.

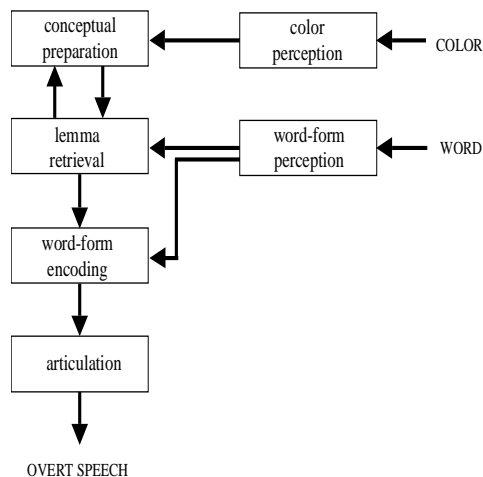


Figure 2: Stages of spoken word planning in WEAVER++

Network Structure

The model assumes that the mental lexicon is a huge network with information about words, a small fragment of which is illustrated in Figure 3.




Figure 3: Fragment of the lexical network of WEAVER++

The network comprises three major strata: a conceptual, a syntactic, and a form stratum. The conceptual stratum represents concepts as nodes in a semantic network, following Collins and Loftus (1975), and many others. For example, the concept RED is represented by the node RED(X). The syntactic stratum contains lemma nodes, such as *red*, which are connected to nodes for their syntactic

class (e.g., adjective). And finally, the form stratum contains nodes representing morphemes, segments, and motor programs. For an extensive discussion of the theoretical and empirical motivation of these assumptions, I refer to Levelt (1989), Levelt et al. (1999), Roelofs (1992, 1993, 1996a,b,c, 1997a,b,c, 1998, 1999, 2000, submitted), and Roelofs and Meyer (1998).

Spreading Activation and Production Rule Application

Information is retrieved from the network by the spreading of activation. For example, a perceived color (e.g., red) activates the corresponding concept node (i.e., RED(X)) in the network. Activation then spreads through the network following a linear activation rule with a decay factor. Each node sends a proportion of its activation to the nodes it is connected to. For example, RED(X) sends activation to other concepts such as GREEN(X) and also to its lemma node *red*. Selection of nodes is accomplished by production rules (i.e., condition-action pairs). A rule is triggered when its nodes become active. A lemma retrieval production rule selects a lemma if the connected concept is flagged as goal concept. For example, *red* is selected for RED(X) in case it is the goal concept and *red* has reached a critical difference in activation compared to other lemmas. The actual moment in time of the firing of a production rule whose condition is satisfied is determined by the ratio of activation of the relevant lemma node and the sum of all the others. Thus, how fast a node is selected depends on how active the other nodes are.

Performing the Stroop Task

In color naming, a production rule like P1 controls general aspects of the task and a rule like P2 achieves the actual lemma selection (and sets a subgoal to encode the word's form, which is omitted here). Word reading is accomplished by a task rule like P3 that maps the orthographic code of a word onto the corresponding articulatory program. Earlier (Roelofs, 1992) I proposed an "intersection" mechanism to achieve selective attention in response selection, which has recently been dropped and replaced by the task production rules (see Roelofs, 2000, submitted).

- (P1) IF the goal is to say the name of the color
and the concept is the color of the stimulus
THEN select the concept
and flag the concept as goal concept
and enhance its activation
- (P2) IF RED(X) is flagged as goal concept
and the activation of *red* exceeds threshold
THEN select *red*
- (P3) IF the goal is to say the name of the word
and the morpheme is the name of the stimulus
THEN select the morpheme
and flag the morpheme as goal morpheme

With task uncertainty, the task itself has to be set during

each trial. This is achieved by production rules like P4 and P5 (in the negative SOA condition).

- (P4) IF the first stimulus is a color
THEN the goal is to name the word
- (P5) IF the first stimulus is a word
THEN the goal is to name the color

To assess the Stroop performance of the model, computer simulations were run. The simulations employed a basic set of eight parameters, whose values were the same as in all earlier simulations (e.g., Levelt et al., 1999; Roelofs, 1992, 1993, 1996a, 1997c) except for two parameter values, which were changed slightly to fine-tune the fit of the model to the data. The "distractor duration" was set to 100 msec and the response threshold to 1.6. The distractor duration determines the gain of the distractor input relative to the target input. Roelofs (submitted) gives all the details of the simulations and applies the model to the key findings from over half a century of Stroop research (e.g., reviewed by MacLeod, 1991).

Illustration of a Simulated Trial

Assume that the task is to name the second stimulus, which may be a color patch or a word. Assume that on a particular trial a red color patch is presented on which the word "green" is superimposed, with the word presented 100 msec before the color patch (i.e., the SOA is -100 msec). The simulation starts with the lemma node of "green" receiving external activation (for 100 msec, the distractor duration). This triggers production rule P5, which sets the goal to naming the color. Activation spreads through the network, with the node *green* sending a proportion of its activation to GREEN(X), and this node in its turn spreads activation to the other concept nodes. After the number of time steps that is the equivalent of 100 msec (the SOA), the concept node RED(X) receives external input from the color. Next, production rule P1 fires, RED(X) becomes flagged as goal concept, and its activation level is selectively enhanced. After the response threshold of the lemma *red* is exceeded, production rule P2 fires and *red* is selected as response.

Simulation Results

The key finding to account for in this paper is that color naming is similarly affected by color words under task certainty and task uncertainty. Maximal inhibition of incongruent words on color naming is obtained when the words are presented within 100 msec of the colors, whereas the facilitation of preexposed congruent words is constant. Whereas color naming is affected by words, reading aloud is not affected by colors. Again, this holds both for task certainty and for task uncertainty.

Figure 4 shows how WEAVER++ performs. The figure shows the SOA curves of the Stroop effects for color naming under task uncertainty. The curves for task certainty (not shown) exhibit the same patterns. Maximal impact of incongruent words occurs in the model when the words are

presented within 100 msec of the color patches, exactly as empirically observed. For reading aloud in the model, no inhibition and facilitation is obtained at any SOA (also not shown), both for task certainty and task uncertainty, as empirically observed. Thus, WEAVER++ accounts for the time course findings and for the effect of the task certainty manipulation.




Figure 4: Time course of the Stroop effects (relative to control) in color naming under task uncertainty from WEAVER++ simulations: \square = incongruent, \triangle = congruent

Why is there no inhibition from colors on word reading? In WEAVER++, lemma retrieval and word-form encoding are discrete. Only the form of a selected lemma becomes activated and selected. Thus, activation does not spread automatically from lemmas to forms but this is under task control. Furthermore, color naming requires both lemma retrieval and form encoding, whereas word reading requires form encoding only. In reading “red” superimposed on a green color patch, the lemma but not the form of “green” becomes active: Because the task is reading and not color naming, the lemma of “green” (corresponding to the color) is not selected and the form of “green” does not become active. The task rule P3 for reading achieves direct selection of the morpheme <red> from the orthographic form *red* rather than indirect selection of <red> by first selecting the lemma *red* and next selecting <red> via the lemma. Thus, selecting <red> from the orthographic form *red* controls the response. Since the form of “green” is not active, planning the form of “red” is unaffected by the color patch.

Independent empirical support for the assumption of a discreteness of stages comes from double-task experiments. Levelt et al. (1991) asked participants to name pictured objects. On one third of the trials (the critical ones), a spoken probe was presented, and participants had to perform a lexical decision on this probe. Peterson and Savoy (1998) also asked participants to name pictures, but on the critical trials in their study written words were presented, which had to be read aloud. The lexical decision and reading latencies showed that in naming a perceptually given entity,

there is no form activation for non-synonymous semantic relatives (i.e., fellow category members) of the target. For example, in naming a cat, there is lemma activation for “cat” and “dog” and word form activation for “cat”, but the word form of “dog” is not activated. By extrapolation, in naming a red color patch, the lemmas of “red” and “green” become active and this also holds for the word form of “red”, but the word form of “green” is not activated. Only the form of a *selected* lemma becomes activated.

O’Seaghdha (1999) argues that the form of “dog” is activated during the naming of a cat, but that the experiments of Levelt et al. (1991) and Peterson and Savoy (1998) were insufficiently powerful to measure this. In support, he refers to a study by O’Seaghdha and Marin, who ran six experiments using word reading with word-word stimuli and an SOA of -500 msec. The effects in the experiments ranged from -2 to +5 msec and were not significant. However, by pooling the observations from the 248 participants in all six experiments, an overall effect of +3 msec was obtained, which reached significance by participants but not by items. By standard criteria, however, such an effect is nonsignificant. Moreover, with large negative SOAs (i.e., -500 msec) expectancy-based priming cannot be excluded (e.g., Neely, 1991). Thus, the findings of O’Seaghdha and Marin do not challenge the discreteness assumption.

Summary and Conclusions

I have argued that performance on the Stroop task provides evidence on how speech production is controlled, that is, how a speaker exerts task-relevant control over the basic language processes underlying naming and oral reading. Color naming is inhibited by incongruent color words but word reading not by incongruent colors. Maximal impact of incongruent words on color naming is obtained when the words are presented within 100 msec of the colors, whereas the effect of preexposed congruent words is constant. The key observation for the current paper is that these findings are obtained both *with* and *without* task certainty. Whereas existing models (e.g., Cohen et al., 1990; Phaf et al., 1990) explain the basic Stroop effects, they fail to account for time course of the findings and for performance under task uncertainty. In this paper, I have extended and applied the WEAVER++ model of word production to performance on the Stroop task, and I have shown that the model accounts for the findings on the time course as well as for the performance under task uncertainty.

References

- Allport, A. (1993). Attention and control: Have we been asking the wrong questions? A critical review of 25 years. In D. Meyer & S. Kornblum (Eds.), *Attention and Performance XIV: A silver jubilee* (pp. 183-218). Cambridge, MA: MIT Press.
- Allport, D. A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and Performance XV: Conscious and nonconscious information processing* (pp. 421-452). Cambridge, MA: MIT

- Press.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. London: Erlbaum.
- Caplan, D. (1992). *Language: Structure, processing, and disorders*. Cambridge, MA: MIT Press.
- Cohen, J., Dunbar, K., & McClelland, J. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97, 332-361.
- Cohen, J. D., & Huston, T. A. (1994). Progress in the use of interactive models for understanding attention and performance. In C. Umiltà & M. Moscovitch (Eds.), *Attention and Performance XV: Conscious and nonconscious information processing* (pp. 453-476). Cambridge, MA: MIT Press.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407-428.
- Flowers, J. H., Warner, J. L., & Polansky, M. L. (1979). Response and encoding factors in "ignoring" irrelevant information. *Memory and Cognition*, 7, 86-94.
- Glaser, M. O., & Glaser, W. R. (1982). Time course analysis of the Stroop phenomenon. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 875-894.
- Glaser, W. R., & Glaser, M. O. (1989). Context effects in Stroop-like word and picture processing. *Journal of Experimental Psychology: General*, 118, 13-42.
- Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.
- Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, 22, 1-38.
- Levelt, W. J. M., Schriefers, H., Vorberg, D., Meyer, A. S., Pechmann, Th. & Havinga, J. (1991). The time course of lexical access in speech production: A study of picture naming. *Psychological Review*, 98, 122-142.
- MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, 109, 163-203.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264-336). Hillsdale, NJ: Erlbaum.
- O'Seaghdha, P. G. (1999). Parsimonious feedback. *Behavioral and Brain Sciences*, 21, 51-52.
- Peterson, R. R. & Savoy, P. (1998). Lexical selection and phonological encoding during language production: Evidence for cascaded processing. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 24, 539-557.
- 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.
- Roelofs, A. (1992). A spreading-activation theory of lemma retrieval in speaking. *Cognition*, 42, 107-142.
- Roelofs, A. (1993). Testing a non-decompositional theory of lemma retrieval in speaking: Retrieval of verbs. *Cognition*, 47, 59-87.
- Roelofs, A. (1996a). Computational models of lemma retrieval. In T. Dijkstra, & K. De Smedt (Eds.), *Computational psycholinguistics: AI and connectionist models of human language processing* (pp. 308-327). London: Taylor & Francis.
- Roelofs, A. (1996b). Serial order in planning the production of successive morphemes of a word. *Journal of Memory and Language*, 35, 854-876.
- Roelofs, A. (1996c). Morpheme frequency in speech production: Testing WEAVER. In G. E. Booij and J. van Marle (Eds.), *Yearbook of Morphology 1996* (pp. 135-154). Dordrecht: Kluwer Academic Publishers.
- Roelofs, A. (1997a). A case for nondecomposition in conceptually driven word retrieval. *Journal of Psycholinguistic Research*, 26, 33-67.
- Roelofs, A. (1997b). Syllabification in speech production: Evaluation of WEAVER. *Language and Cognitive Processes*, 12, 657-693.
- Roelofs, A. (1997c). The WEAVER model of word-form encoding in speech production. *Cognition*, 64, 249-284.
- Roelofs, A. (1998). Rightward incrementality in encoding simple phrasal forms in speech production: Verb-particle combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 904-921.
- Roelofs, A. (1999). Phonological segments and features as planning units in speech production. *Language and Cognitive Processes*, 14, 173-200.
- Roelofs, A. (2000). Control of language: A computational account of the Stroop asymmetry. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the Third International Conference on Cognitive Modeling* (pp. 234-241). Veenendaal, The Netherlands: Universal Press.
- Roelofs, A. (submitted). *Attentional control of verbal action: Stroop phenomena and their time course*.
- Roelofs, A., & Meyer, A. S. (1998). Metrical structure in planning the production of spoken words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 922-939.
- Roelofs, A., Meyer, A.S., & Levelt, W.J.M. (1996). Interaction between semantic and orthographic factors in conceptually driven naming: Comment on Starreveld and La Heij (1995). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 246-251.
- Roelofs, A., Meyer, A.S., & Levelt, W.J.M. (1998). A case for the lemma-lexeme distinction in models of speaking: Comment on Caramazza and Miozzo (1997). *Cognition*, 69, 219-230.
- Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, 124, 207-231.
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge, UK: Cambridge University Press.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643-662.
- Van der Heijden, A. H. C. (1992). *Selective attention in vision*. London: Routledge.