Understanding Word and Sentence
G.B. Simpson (Editor)
Elsevier Science Publishers B.V. (North-Holland), 1991

Chapter 5

The Role of Suppression in Sentence Comprehension

Morton Ann Gernsbacher and Mark Faust
University of Oregon
Eugene, Oregon U.S.A.

Nearly twenty years ago, Gough (1971) wrote: "The problem of when and how a sentence is understood is, in my view, the central problem of experimental psycholinguistics. Its solution in the form of a machine which could understand sentences would, at the least, earn its inventor an invaluable patent. But while a machine which could understand sentences would be something to marvel at, a person who could do only that would not even make good company." (p. 64).

Two decades later, we continue to share Gough's appreciation, amazement, and curiosity. How *do* people comprehend sentences? We have approached this question by tracing the cognitive processes and mechanisms that underlie sentence comprehension (and more generally, language comprehension). We have identified a few of those cognitive processes and mechanisms in a framework we call the Structure Building Framework (Gernsbacher, in press-a; Gernsbacher, in press-b).

According to the Structure Building Framework, the goal of comprehension is to build a coherent mental representation or *structure*. These structures represent sentences, paragraphs, passages, and any other meaningful unit. For instance, comprehending a sentence requires building a mental structure to represent that sentence. The building blocks of mental structures are memory cells. Memory cells represent previously stored memory traces. Their representation might be in either the traditional sense of an individual cell representing an individual trace, or the distributed sense of a group of cells representing an individual trace.

Memory cells are automatically activated by incoming stimuli. Once activated, the information they represent can be used by cognitive processes. Furthermore, according to the Structure Building Framework, once activated, memory cells transmit processing signals. These processing signals either suppress or enhance the activation of other memory cells. So, once memory cells are activated, two mechanisms modulate their level of activation: They are suppression and enhancement.

Suppression decreases or dampens the activation of memory cells when the information they represent is no longer as necessary for the structure being built. Enhancement increases or boosts the activation of memory cells when the information they represent is relevant to the structure being built. By modulating the activation of memory cells, suppression and enhancement contribute to comprehension.

The notion that incoming stimuli activate memory representations is familiar. What is novel about the Structure Building Framework is its proposal that activated memory cells transmit processing signals. This additional proposal more fully captures the analogy of neural activity — an analogy that inspires many models of cognition. The familiar notion that incoming stimuli activate memory representations captures only one aspect of the analogy, the electrical transmission of information (along axons). But the novel proposal that activated memory cells transmit processing signals (for suppression and enhancement) continues the analogy by paralleling the chemical transmission of information (across synapses, via neurotransmitters).

We propose that the mechanisms of suppression and enhancement are general cognitive mechanisms. They are not dedicated to language; they play vital roles in many nonlinguistic phenomena, too. Yet, they are crucial to language comprehension.

In this chapter, we focus on the mechanism of suppression and the vital role that suppression plays in sentence comprehension. In the first half of the chapter, we illustrate the vital role that suppression plays in sentence comprehension by demonstrating how suppression fine tunes the meanings of words. In the second half of the chapter, we illustrate the vital role that suppression plays in sentence comprehension by documenting that less-skilled comprehenders suffer from less-efficient suppression mechanisms.

The Role of Suppression in Fine Tuning the Meanings of Words

According to many models of word understanding, when comprehenders first hear or read a word, information provided by that word activates various potential meanings. Then, constraints provided by lexical, semantic, syntactic, and other sources of information alter those meanings' levels of activation. Eventually, one meaning becomes most strongly activated. That meaning is what comprehenders access and incorporate into their developing mental structures (these ideas are culled from the models of (Becker, 1976; Kintsch, 1988; Marslen-Wilson, Tyler, & Seidenberg, 1978; McClelland & Kawamoto, 1986; Norris, 1986).

What the Structure Building Framework adds to these ideas is the proposal that the mechanisms of suppression and enhancement modulate the different meanings' levels of activation. In particular, according to the Structure Building Framework, the mechanism of suppression helps fine tune the meanings of words by suppressing the less likely meanings. An excellent arena for demonstrating the vital role that suppression plays in fine tuning word meaning is provided by ambiguous words — for instance, words like *bug* that have at least two diverse meanings.

Ambiguous words have clearly distinct meanings, and in sentence contexts one meaning is usually more appropriate. But contrary to intuition, immediately after comprehenders hear or read an ambiguous word in context, multiple meanings are often activated, even when only one meaning is suggested by the context.

For example, immediately after comprehenders hear the word *bug*, both the "insect" meaning and the "covert microphone" meaning are activated (Swinney, 1979). Both meanings are activated even when the context is biased toward the "insect" meaning, as in

(1) The man was not surprised when he found several spiders, roaches, and other *bugs* ....

This immediate activation of multiple meanings, regardless of context, is demonstrated by the following experimental task: Subjects listen to a series of sentences. At a critical point during each sentence, the subjects see a test word. The subjects must decide rapidly whether that test word is an English word.

For example, if sentence (1) was presented in such an experiment, then immediately after subjects heard the word *bug*, they might see the test word *ANT*. That test word is related to the contextually appropriate meaning of *bug* (the meaning implied by the context). In another condition of the same experiment, subjects might see the test word *SPY* immediately after they hear *bug*. The test word *SPY* is related to a contextually inappropriate meaning of *bug* (a meaning not implied by the context). In a third condition, the subjects might see the test word *SEW*. That test word is unrelated to any meaning of *bug* and serves as a control.

If subjects are tested immediately after they hear the word *bug*, they respond just as rapidly to *SPY* as they respond to *ANT*. And they respond to both *SPY* and *ANT* more rapidly than they respond to the unrelated test word *SEW*. In other words, subjects respond to test words that are related to the contextually inappropriate meanings just as rapidly as they respond to test words that are related to the contextually appropriate meanings. This result suggests that immediately after comprehenders hear ambiguous words, both appropriate and inappropriate meanings are activated — and both meanings are more activated than unrelated concepts.

But this is only what happens when activation is measured immediately after comprehenders hear ambiguous words. Comprehenders do not keep multiple meanings activated forever. If they did, they would never unambiguously understand any utterance or passage. Instead, multiple meanings are activated only momentarily.

For instance, when subjects continue listening to sentence (1) and are tested only four syllables after hearing the word *bug*, they still respond rapidly to *ANT*. But they respond no more rapidly to *SPY* than they respond to *SEW*. In other words, they respond no more rapidly to test words that are related to contextually inappropriate meanings than they respond to test words that are unrelated to any meaning. This finding suggests that after one and a half syllables, the inappropriate meanings have decreased in activation. Other experiments have demonstrated that inappropriate meanings decrease in activation even more quickly, often within only 200 ms. That is probably why comprehenders are typically aware of only one meaning — the contextually appropriate one.

This phenomenon, immediate activation of multiple meanings but continued activation of only appropriate meanings, is also demonstrated with other laboratory tasks. It is demonstrated when subjects read sentences one word at a time, and occasionally, instead of seeing the next word of a sentence, they see the test words. They decide rapidly whether each test word is an English word (Kintsch & Mross, 1985; Till, Mross, & Kintsch, 1988).

The phenomenon is also demonstrated when subjects listen to sentences and are visually presented with test words. But instead of rapidly deciding whether each test word is an English word, they simply pronounce each test word as rapidly as possible. Or they simply name the color of ink in which each test word is printed (Conrad, 1974; Lucas, 1987; Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Tanenhaus, Leiman, & Seidenberg, 1979).

Each of these laboratory tasks demonstrates that multiple meanings of ambiguous words are often immediately activated — regardless of semantic context. But only contextually appropriate meanings remain activated a short while later. This phenomenon occurs even when one meaning is a noun and the other is a verb (Seidenberg et al., 1982). For example, *watch* refers to both an object, *a timepiece*, and an action, *looking*. Sentence (2) implies the noun meaning of *watch*, while sentence (3) implies the verb meaning.

(2) I like *the* watch.
(3) I like *to* watch.

Why multiple meanings are immediately activated without regard to context intrigues researchers, perhaps because the phenomenon challenges introspection. Many laboratory investigations have searched for its boundary conditions (Blutner & Sommer, 1988; Burgess, Tanenhaus, & Seidenberg, 1989; Glucksberg, Kreuz, & Rho, 1986; Tabossi, 1988; Tabossi, Colombo, & Job, 1987; Van Petten & Kutas, 1987; Williams, 1988).

But equally intriguing are the following questions: What happens to the inappropriate meanings? How do they become less activated? Unfortunately, scant empirical attention has been directed toward answering these questions.

According to the Structure Building Framework, inappropriate meanings become less activated via the mechanism of suppression. The memory cells representing the semantic or syntactic context transmit processing signals. These processing signals suppress the contextually inappropriate meanings. Dampening the activation of inappropriate meanings could be one of the most important roles that the mechanism of suppression plays in sentence comprehension.

But other theories assume that inappropriate meanings become less activated via other mechanisms. For instance, according to some theories, inappropriate meanings are inhibited by appropriate meanings, and according to other theories, inappropriate meanings simply decay. Unfortunately neither assumption has been tested empirically. That was the purpose of the experiments we shall describe next.

Are Inappropriate Meanings Mutually Inhibited?

Some theories propose that inappropriate meanings become less activated through a mechanism we shall call *compensatory inhibition* (McClelland & Kawamoto, 1986; Waltz & Pollack, 1985). These theories assume that all concepts compete for a fixed amount of activation. So when multiple meanings of ambiguous words are immediately activated, they are sharing this fixed sum. Later, inappropriate meanings must decrease in activation presumably because appropriate meanings have increased. Like a seesaw, when one meaning becomes more activated, the other must become less activated.

But if reaction times reflect activation, which is what many reaction time researchers assume (Posner, 1978), the behavioral data do not demonstrate this compensatory pattern. Simply put: The appropriate meanings do not increase in activation when the inappropriate meanings decrease. For instance, in neither Swinney's (1979) nor Seidenberg et al.'s (1982) data did the appropriate meanings increase in activation from the immediate to the delayed test point. But in both sets of data, the inappropriate meanings decreased. This is the pattern typically observed in these experiments.

Perhaps appropriate meanings do not observably increase in activation because during the delay they are competing with other concepts for the fixed sum of activation. By definition, Swinney's (1979) four-syllable delay introduced new syllables (four, to be precise). Perhaps during these four syllables, new concepts were introduced. For example, sentence (1) continued,

(4) The man was not surprised when he found several spiders, roaches, and other bugs *in the corner of the room.*

We need some way to introduce a delay without introducing new concepts. In the following experiment we did just that. We selected 48 ambiguous words that were just as likely to be thought of as verbs as nouns, according to ambiguity norms (Cramer, 1970; Kausler & Kollasch, 1970; Nelson, McEvoy, Walling, & Wheeler, 1980). For each ambiguous word, we constructed two experimental sentences. The two sentences were identical until after the ambiguous word occurred, with the following exception: In one sentence, the ambiguous word was preceded with the infinitive marker *to*, whereas in the other sentence, the ambiguous word was preceded with the definite article *the*. For example,

(5) Jack tried *to* punch ....
(6) Jack tried *the* punch ....

For each ambiguous word, we selected two test words: One test word was related to the verb meaning, and the other was related to the noun meaning. The two test words for sentences (5) and (6) are illustrated in Table 1.

Table 1

SENTENCES	TEST WORDS	
	HIT	**DRINK**
Jack tried to punch ...	**Related** to APPROPRIATE Meaning	**Related** to INAPPROPRIATE Meaning
Jack tried the punch ...	**Related** to INAPPROPRIATE Meaning	**Related** to APPROPRIATE Meaning
Jack tried to bluff ...	**Unrelated** to Any Meaning	**Unrelated** to Any Meaning
Jack tried the rolls ...	**Unrelated** to Any Meaning	**Unrelated** to Any Meaning

For each ambiguous word, we also constructed two control sentences, which were identical to the two experimental sentences up to the point where the ambiguous words occurred. In the control sentences, the experimental ambiguous words were replaced with other ambiguous words (which they matched

in length and familiarity). For example,

(7) Jack tried to *bluff* ....
(8) Jack tried the *rolls* ....

The control words (e.g., *bluff* or *rolls)* were unrelated to the test words (e.g., *HIT* or *DRINK)*. This relationship is also illustrated in Table 1.

Finally, we constructed 48 "lure" sentences that resembled the experimental and control sentences. The test words for the lure sentences were pronounceable strings of letters that did not form English words (e.g., *HUP*, *DRACK)*. All of the sentences were presented visually, word-by-word in the center of a computer screen. Immediately after the ambiguous word disappeared (e.g., *punch)*, or the control word disappeared (e.g., *bluff)*, a test word appeared. The test words appeared at the top of the screen in capital letters. Subjects decided rapidly whether each test word was an English word.

After the ambiguous or control words occurred, their sentences continued in meaningful but different ways. For example,

(9) Jack tried the punch *but he didn't think it tasted very good.*

However, remember that the test words always appeared immediately after the ambiguous or control words; so, activation was always measured before the sentences diverged.

We measured activation at two test intervals. These test intervals were produced by manipulating the rate at which the words in the sentences appeared. There were two presentation rates: At the faster rate, each word ap-

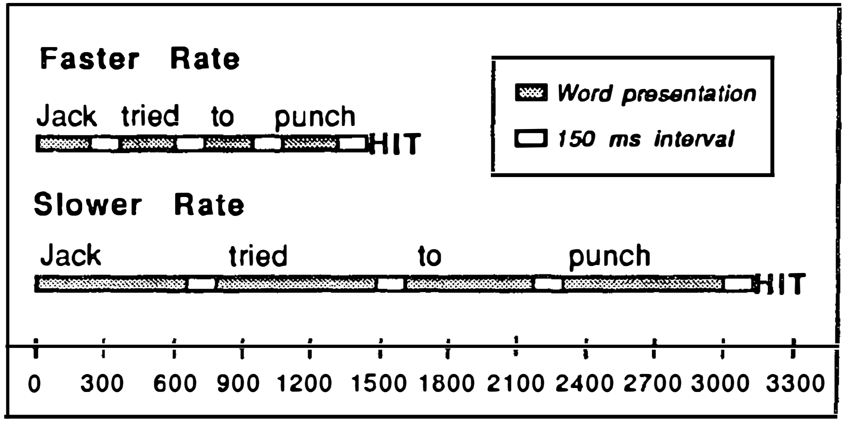

Figure 1. Rates at which the words of the sentences were presented.

peared for 16.667 ms per character, plus a constant 150 ms. At the slower rate, each word appeared for 50 ms per character, plus a constant 450 ms. Figure 1 illustrates these presentation rates.

In both the fast and slow presentation rate, a constant 150 ms intervened between the appearance of each word in a sentence. And in both the fast and slow presentation rate, a constant 150 ms intervened between the ambiguous (or control) word and its test word. The difference in these rates created the difference between the two test points. With the faster rate, a five-letter word (like *punch)* appeared for 233 ms; with the slower rate, the same five-letter word

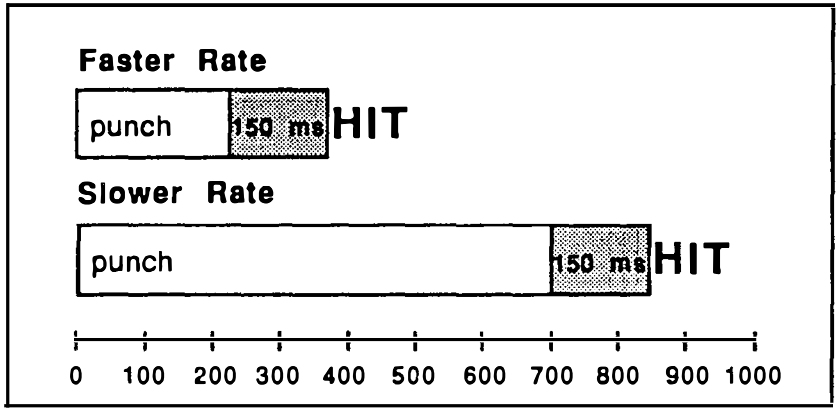

Figure 2. Rates at which the words of the sentences were presented.

appeared for 700 ms. So, the difference between the two test points for five-letter words was 467 ms. Figure 2 illustrates this difference. For continuity with the other experiments we have discussed, we shall call the test point produced by the faster rate *Immediate*, and the test point produced by the slower rate *Delayed*.

Figure 3 displays our 80 subjects' data. We estimated activation by subtracting subjects' latencies to respond to test words that were related to the appropriate or inappropriate meanings of the ambiguous words from their latencies to respond to test words that were unrelated to any meaning of the ambiguous words. For instance, we estimated the activation of the contextually appropriate meaning of *to punch* by subtracting subjects' latencies to respond to *HIT* after reading *Jack tried to punch* from their latencies to respond to *HIT* after reading *Jack tried to bluff*. Similarly, we estimated the activation of the contextually inappropriate meaning of *to punch* by subtracting subjects' latencies to respond to *DRINK* after reading *Jack tried to punch* from their latencies to respond to *DRINK* after reading *Jack tried to bluff*.

First examine what happened at the immediate test point. As Figure 3

illustrates, at the immediate test point (caused by the faster presentation rate), both the appropriate and the inappropriate meanings were reliably more activated than unrelated concepts, *minF'* (1,83) = 6.495, $p < .01$ for appropriate

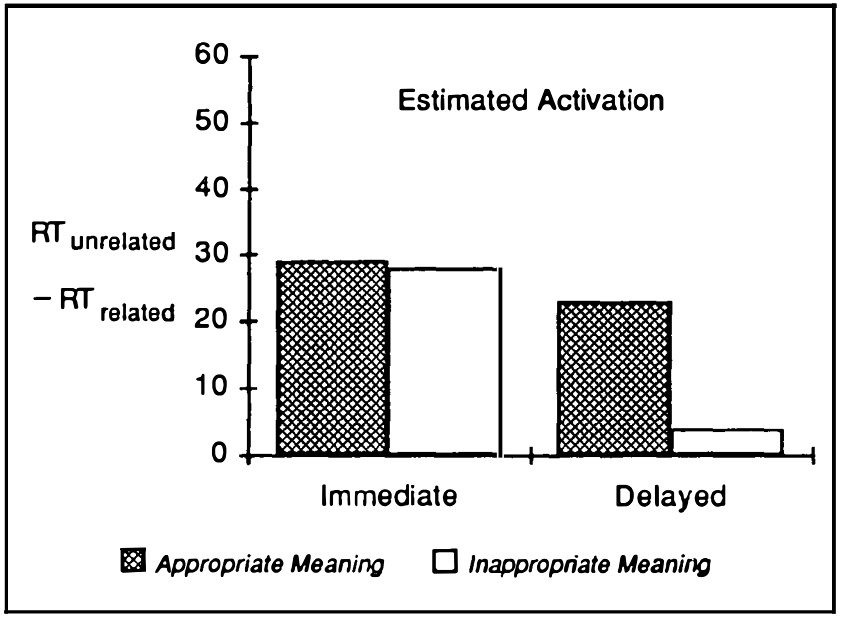

Figure 3. Subjects' average activation scores.

meanings, and *minF'* (1,82) = 10.26, $p < .005$ for inappropriate meanings. Indeed, at this immediate point, the appropriate and inappropriate meanings were activated at the same level (i.e., their activation levels did not differ, both *Fs* < .5).

Now, examine what happed at the delayed test point. As Figure 3 illustrates, after the delay (caused by the slower presentation rate), only the appropriate meanings were reliably activated, *minF'* (1,82) = 4.562, $p < .05$. In contrast, the inappropriate meanings were considerably less activated than the appropriate meanings, *minF'* (1,82) = 3.919, $p < .05$. Indeed, the inappropriate meanings were no more activated than unrelated concepts, both *Fs* < 1.0.

These data replicate those of Swinney (1979) and Seidenberg et al. (1982). They also demonstrate that when inappropriate meanings decrease in activation, appropriate meanings do not increase; in other words, there is no compensation. If reaction times reflect activation levels, then there is no evidence that inappropriate meanings lose activation because appropriate meanings take a larger share (of a fixed sum). In other words, there is no evidence to support the compensatory inhibition explanation for why inappropriate meanings lose activation, especially when no new concepts are introduced during the test delay.

Do Inappropriate Meanings Simply Decay?

Another explanation for why inappropriate meanings become less activated is that they *decay*. In many models of cognition, mental representations automatically decay when they are not continuously stimulated (Anderson, 1983). Inappropriate meanings might therefore decay because they do not continuously receive stimulation from a biasing semantic or syntactic context. We empirically tested this decay explanation in the following experiment.

We selected 48 ambiguous words that were just as likely to be thought of as one noun as another (according to ambiguity norms). For example, the word *quack* is just as likely to be interpreted as "an incompetent doctor" as "the sound a duck makes."

For each of the 48 ambiguous words, we constructed three experimental sentences. One experimental sentence was biased toward one meaning of the ambiguous word, for example,

(10) Pam *was diagnosed by* a quack ....

A second experimental sentence was biased toward another meaning of the ambiguous word, for example,

(11) Pam *heard a sound like* a quack ....

But the third experimental sentence was neutral: Neither its semantic nor its syntactic context was biased toward either meaning of the ambiguous word, for example,

(12) Pam *was annoyed by* the quack ....

To ensure that our sentences were effectively biased or neutral, we had 50 subjects read the beginnings of the sentences (e.g., *Pam was annoyed by the quack ....*). These subjects decided which meaning was intended. We used biased sentences only if 95% of these subjects agreed with the meaning we intended, and we used neutral sentences only if these subjects were roughly split over which of the two meanings we intended.

For each of the 48 ambiguous words, we selected two test words. One was related to one of the biased meanings (e.g., *DOCTOR*), and the other was related to the other biased meaning (e.g., *DUCK*). The test words and experimental sentences are illustrated in Table 2.

For each of the 48 ambiguous words, we also constructed a control sentence. The control sentences were identical to the neutral experimental sentences to the point where the ambiguous words occurred. In the control sentences, the experimental ambiguous words were replaced with unrelated ambiguous words (which matched the experimental words in length and familiarity). For example,

(13) Pam was annoyed by *the pupil* ....

The ambiguous words in the control sentences were unrelated to the test words. This relationship is also illustrated in Table 2.

Table 2

SENTENCES	TEST WORDS	
	DOCTOR	**DUCK**
Pam was diagnosed by a quack ...	**Related** to APPROPRIATE Meaning	**Related** to INAPPROPRIATE Meaning
Pam heard a sound like a quack ...	**Related** to INAPPROPRIATE Meaning	**Related** to APPROPRIATE Meaning
Pam was annoyed by the quack ...	**Neutral**	**Neutral**
Pam was annoyed by the pupil ...	**Unrelated** to Any Meaning	**Unrelated** to Any Meaning

We also constructed 48 lure sentences that resembled the experimental and control sentences, but the test words for the lure sentences were pronounceable strings of letters that did not form English words. All the sentences were presented visually, as in the experiment we described before. And as in the experiment we described before, the sentences continued in meaningful but different ways after the ambiguous or control words. For example,

(14) Pam heard a sound like a quack *but couldn't imagine where it was coming from.*

However, it was before the sentences diverged that we measured activation. We again manipulated the presentation rate (as illustrated in Figures 1 and 2), so that we could measure activation at two test points without introducing new concepts.

To summarize, there were three experimental sentences. One was biased toward one meaning of the ambiguous words; one was biased toward another meaning; and the third was neutral — there was no semantic or syntactic bias. While subjects read these experimental sentences, we measured how activated the multiple meanings were. And we made this measurement at two test points.

The decay explanation and the suppression explanation make identical predictions about the biased sentences; these sentences should replicate earlier experiments: At the immediate test point, both appropriate and inappropriate meanings should be activated, but at the delayed test point, the inappropriate meanings should be less activated (in relation to the unrelated control condition).

Where the decay and the suppression explanations differ is their predictions about the neutral sentences. According to the decay explanation, inappropriate meanings become less activated because they automatically decay. And they decay because they lack stimulation from a semantic or syntactic context. Because neutral sentences also lack stimulation from a semantic or syntactic context, multiple meanings of ambiguous words should also decay. In other words, the decay explanation predicts that with neutral sentences, both meanings should be less activated after the delay than they are immediately. This is because neither meaning receives stimulation from a semantic or syntactic context.

In contrast, according to the suppression explanation, inappropriate meanings become less activated because the memory cells representing semantic or syntactic contexts transmit processing signals; these processing signals suppress the inappropriate meanings' activation. So, the suppression explanation predicts that only the inappropriate meanings of the biased sentences should become less activated after the delay; the multiple meanings of the neutral sentences should be just as activated after the delay as they are immediately. This is because there are no bases from which suppression signals can be transmitted.

So, the decay explanation predicts that with the neutral sentences, both meanings should be less activated after the delay than they are immediately. But the suppression explanation predicts that both meanings should be just as activated after the delay as they are immediately.

Figure 4 displays our 80 subjects' data. We estimated activation by subtracting subjects' latencies to respond to test words that were related (to the appropriate, inappropriate, or both meanings of the neutral sentences) from their latencies to respond to test words that were unrelated to any meaning.

First, examine what happened at the immediate test point. As Figure 4

illustrates, with the biased sentences, both the appropriate and inappropriate meanings were reliably activated as were both meanings with the neutral sentences, *minF'* $(1,73) = 4.503$, $p < .05$ for appropriate meanings, *minF'* $(1,83) = 4.773$, $p < .05$ for inappropriate meanings, and *minF'* $(1,83) = 3.711$, $p < .05$ for both meanings with the neutral sentences.

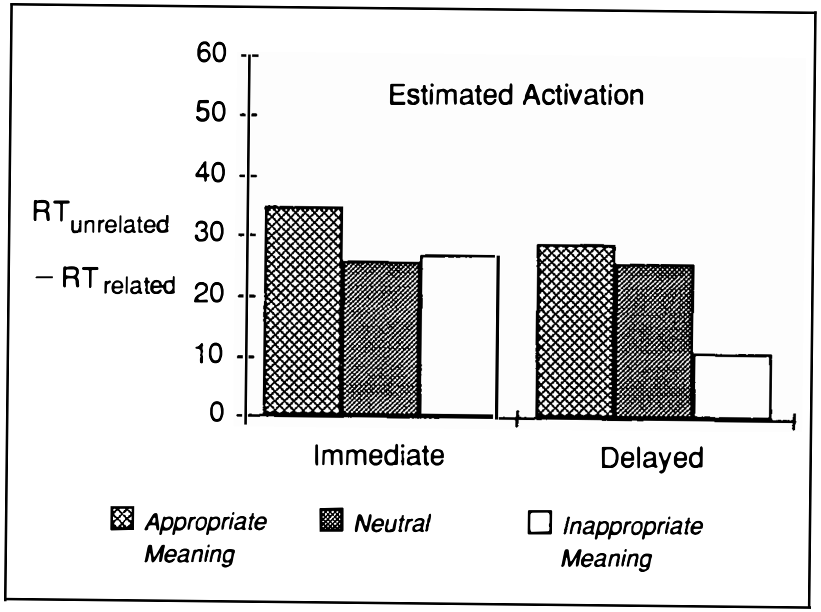

Figure 4. Subjects' average activation scores

Now, examine what happened after the delay. As Figure 4 illustrates, after the delay, the inappropriate meanings of the biased sentences were less activated; indeed, they were (statistically) no more activated than unrelated concepts, *minF'* < 1.0. In contrast, with the neutral sentences, both meanings were still reliably more activated than unrelated concepts, *minF'* $(1,83) = 3.846$, $p < .05$. The same was true of the appropriate meanings (with the biased sentences), *minF'* $(1,83) = 4.702$, $p < .05$.

Indeed, as Figure 4 illustrates, with the neutral sentences, the ambiguous words' multiple meanings were just as activated after the delay as they were immediately. These results confirm the prediction made by the suppression explanation, not the decay explanation. The suppression explanation, drawn from the Structure Building Framework, predicts that inappropriate meanings become less activated because the memory cells representing semantic or syntactic contexts transmit processing signals; these processing signals suppress the inappropriate meanings' activation. With a neutral context, multiple meanings remain activated because there are no bases from which suppression signals can be transmitted.

These two experiments demonstrate that the inappropriate meanings of ambiguous words do not decrease in activation because they are mutually inhibited; neither do they decrease in activation because they decay. Rather, we suggest that they are suppressed. In this way, the mechanism of suppression plays a vital role in sentence comprehension: It fine tunes the meanings of words.

The Role of Suppression in Comprehension Skill

There are many situations in which irrelevant or inappropriate information is automatically activated, unconsciously retrieved, or naturally perceived. But for successful comprehension, this irrelevant or inappropriate information must not affect ongoing processes. According to the Structure Building Framework, irrelevant or inappropriate information is suppressed.But what if a comprehender's suppression mechanism was faulty? Irrelevant or inappropriate information would remain activated. Surely that would affect the comprehender's success. Perhaps that is one of the reasons why some comprehenders are less successful: They have less-efficient suppression mechanisms. In several experiments, we have investigated whether less-skilled comprehenders are indeed characterized by less-efficient suppression mechanisms.

Are Less-skilled Comprehenders Less Efficient at Suppressing the Inappropriate Meanings of Ambiguous Words?

We have suggested that successful comprehension requires suppressing the contextually inappropriate meanings of ambiguous words. For example, successfully comprehending sentence (15) requires suppressing the playing card meaning of the word *spade.*

(15) He dug with the *spade.*

If less-skilled comprehenders are plagued by less-efficient suppression mechanisms, then they should be less able to suppress these contextually inappropriate meanings.

We tested this hypothesis in Gernsbacher, Varner, and Faust (1990). We selected two samples of more- versus less-skilled comprehenders from a distribution of 270 University of Oregon students. All 270 students had previously been tested on our Multi-Media Comprehension Battery (Gernsbacher & Varner, 1988). The more-skilled comprehenders' scores were from the upper third of the distribution of Comprehension Battery scores; the less-skilled comprehenders were from the lower third of the distribution. When these more- and less-skilled comprehenders returned to the lab, they performed the following task: They read short sentences, and after each sentence, they saw a test word. Their task was to verify whether the test word fit the meaning of the sentence they just read. On 80 trials, the test word did indeed fit the sentence, but we were more interested in the 80 trials in which the test word did *not* fit the sentence.

On half of those trials, the last word of the sentence was an ambiguous word, for example,

(15) He dug with the *spade.*

The test word on these trials was a meaning of the ambiguous word that was inappropriate to the context, for example, *ACE* We measured how long subjects took to reject a test word like *ACE* after reading a sentence like (15). And we compared that latency with how long subjects took to reject *ACE* after reading the same sentence but with the last word replaced by an unambiguous word, for example,

(16) He dug with the *shovel.*

This comparison showed us the activation level of the inappropriate meanings; the more time subjects took to reject *ACE* after the *spade-* versus the *shovel-*sentence, the more activated the inappropriate meaning must have been.

We presented the test words at two points: immediately (100 ms) after subjects finished reading each sentence, and after an 850 ms delay. We predicted that at the immediate test point, both the more- and less-skilled comprehenders would take longer to reject test words after reading the ambiguous words as opposed to the unambiguous words. For example, both groups would take longer to reject *ACE* after reading the *spade* sentence than after reading the *shovel* sentence. This prediction was based on the studies we described earlier which demonstrate that immediately after ambiguous words are read, contextually inappropriate meanings are often activated. We particularly expected the inappropriate meanings to be activated because our task required comprehenders to focus their attention on a subsequent word and try to integrate that word into the previous context (Glucksberg et al., 1986; Van Petten & Kutas, 1987).

Our novel predictions concerned what would happen after the 850 ms delay. We predicted that by that point the more-skilled comprehenders would not take longer to reject test words following ambiguous words. This is because more-skilled comprehenders should be able to successfully suppress the inappropriate meanings. But we made a different prediction for our less-skilled comprehenders. If less-skilled comprehenders are plagued by less-efficient suppression mechanisms, then even after the delay, the inappropriate meanings should still be activated.

Figure 5 displays our 64 subjects' data. We estimated activation by subtracting subjects' latencies to reject test words like *ACE* after reading ambiguous words like *spade* from their latencies to reject test words like *ACE* after reading unambiguous words like *shovel*. The more-skilled comprehenders are represented by hashed lines, and the less-skilled comprehenders are represented by unfilled bars.

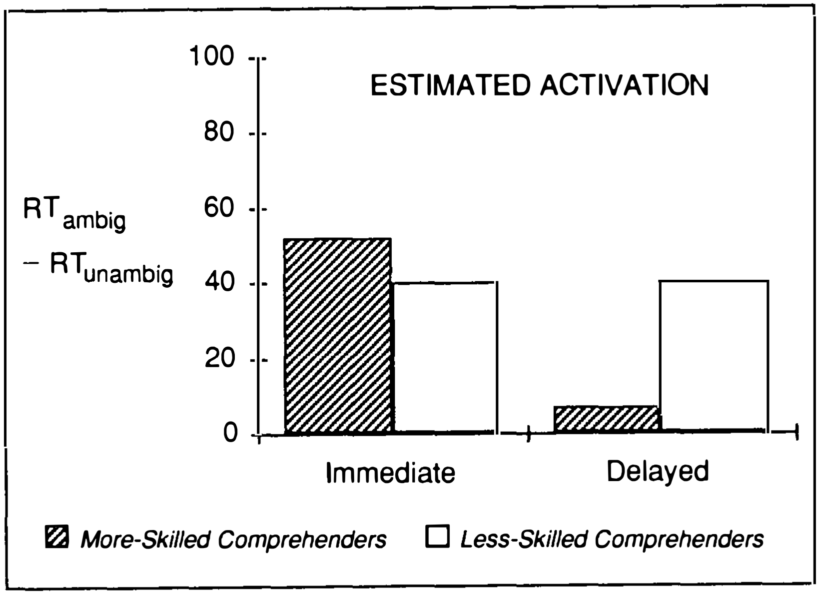

Figure 5. Subjects' average activation scores (from Gernsbacher et al., 1990)

First, examine what happened at the immediate test point. As Figure 5 illustrates, immediately after both the more- and less-skilled comprehenders read the ambiguous words, the inappropriate meanings were highly activated. Now, examine what happened after the delay. As Figure 5 illustrates, 850 ms after the more-skilled comprehenders read the ambiguous words, the inappropriate meanings were no longer reliably activated; by this time, the more-skilled comprehenders had successfully suppressed them. But the less-skilled comprehenders were less fortunate: As Figure 5 illustrates, even after the 850 ms delay, the inappropriate meanings were still highly activated. In fact, they were as highly activated after the delay as they were immediately. So, almost a second after the less-skilled comprehenders read the ambiguous words, they were unable to suppress the inappropriate meanings. These results support the hypothesis that less-skilled comprehenders are plagued by less rapid (and therefore less-efficient) suppression mechanisms.

Are Less-skilled Comprehenders Less Efficient at Suppressing the Incorrect Forms of Homophones?

Reading a string of letters activates an array of information. Virtually always reading a letter string activates orthographic information — information about the individual letters in the string and their relative position to one another. Often, reading a letter string activates semantic information, lexical information, and phonological information. In fact, semantic, lexical, and phonological information is often activated even when the string does not compose an English word (Coltheart, Davelaar, Jonasson, & Besner, 1977; Rosson, 1985).

Automatic activation of phonological information was the focus of our next experiment. By automatic activation of phonological information we meant the phenomenon in which reading the letter string *rows* activates the phonological sequence */roz/*. In fact, reading *rows* can activate */roz/*, which can activate *rose*. In other words, reading a homophone *(rows)* can activate a phonological sequence *(/roz/)*, which can then activate another form of the homophone *(rose)*. How do we know that a letter string often activates phonological information, which in turn activates other forms of homophones? Consider the following finding: Comprehenders have difficulty quickly rejecting the word *rows* as not being an exemplar of the category *FLOWER* (van Orden, 1987; van Orden, Johnston, & Hale, 1988).

But to successfully comprehend a written passage, these incorrect forms cannot remain activated. According to the Structure Building Framework, sentence comprehension involves the mechanism of suppression. The same cognitive mechanism that suppresses the inappropriate meanings of ambiguous words, could also suppress the incorrect forms of homophones. If this is the same mechanism, and if this general suppression mechanism is less efficient in less-skilled comprehenders, then less-skilled comprehenders should also less efficiently suppress the incorrect forms of homophones.

Related evidence already supports this prediction. Consider the sentence:

(17) She *blue* up the balloon.

Six-year olds are more likely to accept that sentence than are 10-year olds — even when they clearly know the difference between *blue* and *blew* (Coltheart, Laxon, Rickard, & Elton, 1988; Doctor & Coltheart, 1980). If we assume that 6-year olds are less skilled than 10-year olds at comprehension, this finding suggests that less-skilled comprehenders are less able to suppress the incorrect forms of homophones that are often automatically activated.

In Gernsbacher and Faust (in press), we tested this hypothesis more directly, with adult subjects whom we knew differed in comprehension skill. Our subjects were US Air Force recruits who were drawn from a sample of 455 subjects whom we had previously tested with the Multi-Media Comprehension

Battery.[1] We drew 48 subjects from the top third of the distribution (those who scored the highest) and 48 subjects from the bottom third of the distribution (those who scored the lowest).

When these more- versus less-skilled comprehenders returned to the lab, they performed a laboratory task similar to the task we used in Gernsbacher et al. (1990). The subjects read short sentences, and following each sentence, they saw a test word. The subjects' task was to verify whether the test word fit the meaning of the sentence they just read. On 80 trials, the test word did indeed fit the sentence's meaning, but on 80 trials it did not. We were interested in those trials in which the test word did *not* fit the meaning.

On half of those trials, the last word of the sentence was one form of a homophone, for example,

(18) He had lots of *patients.*

On these trials, the test word was related to the homophone's other form; for example, the test word *CALM* is related to *patience.* We compared how long subjects took to reject *CALM* after reading sentence (18) with how long they took to reject *CALM* after reading the same sentence with the last word replaced by a nonhomophone, for example,

(19) He had lots of *students.*

This comparison showed us the activation levels of the incorrect forms; the more time subjects took to reject *CALM* after the *patients*- versus *students*-sentence, the more activated the *patients* form of the homophone must have been.[2]

We presented the test words at two test points: immediately (100 ms) after subjects finished reading each sentence, and after a one-second delay. We predicted that at the immediate test point, both the more- and less-skilled comprehenders would take longer to reject test words following homophones than nonhomophones. For example, both groups would take longer to reject *CALM* after reading the *patients* sentence than after reading the *students* sentence. This result would corroborate the results of van Orden (1987; van Orden et al., 1988). This result would also demonstrate that comprehenders of both skill levels often activate phonological information during reading.

Our novel predictions concerned what would happen after the one-second delay. We predicted that after the one-second delay, the more-skilled comprehenders would not take longer to reject test words following homophones versus nonhomophones; this is because more-skilled comprehenders should be able to successfully suppress incorrect forms. But we made a different prediction for our less-skilled comprehenders. If less-skilled comprehenders are characterized by less-efficient suppression mechanisms, then even after the one-

second delay, the incorrect forms of the homophones should still be highly activated.

Figure 6 illustrates our 96 subjects' data. We estimated activation by sub-

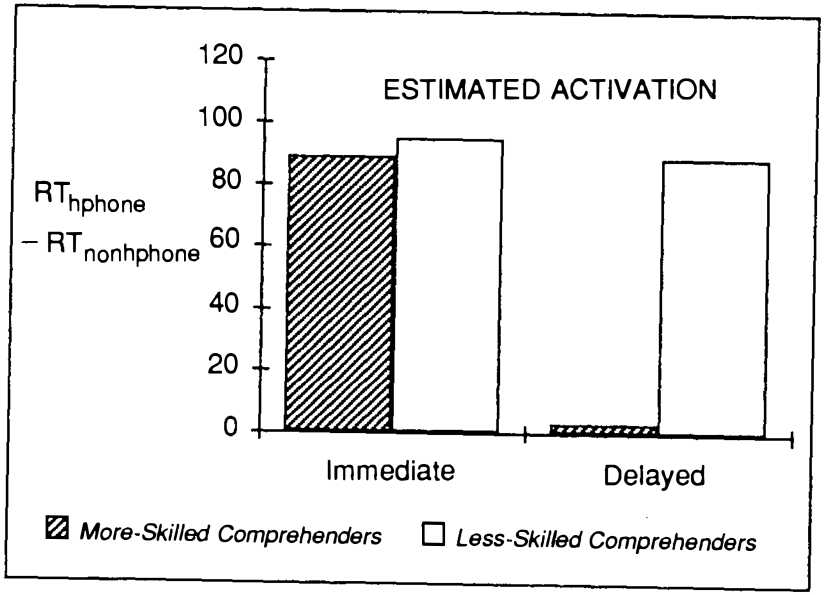

Figure 6. Subjects' average activation scores (from Gernsbacher & Faust, in press)

tracting subjects' latencies to reject test words like *CALM* after reading nonhomophones like *students* from their latencies to reject test words like *CALM* after reading homophones like *patients.* First examine what happened at the immediate test test point. As Figure 6 illustrates, immediately after both the more- and less-skilled comprehenders read the homophones, the inappropriate forms were highly activated; in fact, they were almost equally activated for the more- versus less-skilled comprehenders. So, 100 ms after comprehenders of both skill levels read homophones, other forms are often activated.

Now, examine what happened after the one-second delay. As Figure 6 illustrates, one second after the more-skilled comprehenders read the homophones, the incorrect forms were no longer reliably activated; the more-skilled comprehenders had successfully suppressed them. But as Figure 6 also illustrates, the less-skilled comprehenders were less fortunate: Even after the one-second delay, the incorrect forms were still highly activated; in fact, they were as activated after one second as they were immediately. So, a second after the less-skilled comprehenders read the homophones, they were unable to suppress the incorrect forms. These data support the hypothesis that less-skilled comprehenders are plagued by less-efficient suppression mechanisms.

Are Less-skilled Comprehenders Less Efficient at Suppressing Information Across Modalities?

Comprehension often requires making sense of stimuli that originate from various modalities. We would be severely handicapped if we were skilled at only reading written words, or only listening to spoken words, or only comprehending graphic displays. Information originates from different modalities, often simultaneously. We read while listening to music, and we drive while carrying on a conversation.

Comprehenders often experience interference across modalities. For instance, it is harder to name a pictured object such as an *ashtray* if a letter string such as *INCH* is written across the picture, as illustrated in the upper left panel

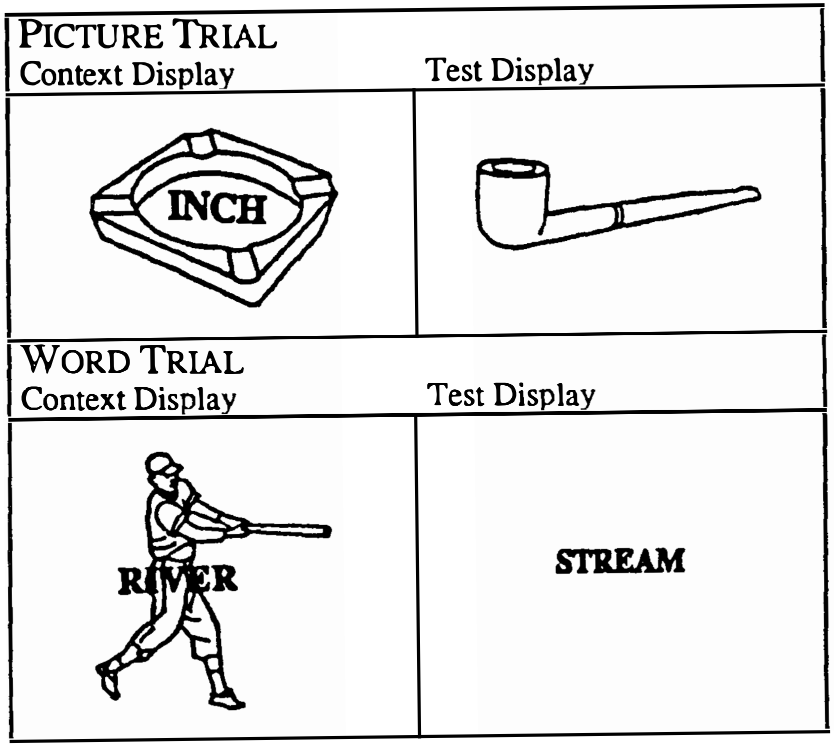

Figure 7. (From Gernsbacher & Faust, in press)

of Figure 7. The opposite is also true: It is harder to read a word such as *RIVER* if it is superimposed on a picture, as illustrated in the bottom left panel of Figure 7 (Smith & McGee, 1980).

Successful comprehension often requires suppressing information across modalities. The same mechanism that suppresses information within modality, could suppress information across modalities. If this is the same mechanism, and if this general suppression mechanism is less efficient in less-skilled comprehenders, then less-skilled comprehenders should also be less efficient in suppressing information across modalities.

We tested this hypothesis in the following way. We modified Tipper and Driver's (1988) experimental task. In our modification, subjects first viewed a context display. Each context display contained a line-drawn picture of a common object and a familiar word. For example, the top panel in Figure 7 illustrates a picture of an *ashtray* with the word *INCH* written across it. The bottom panel of Figure 7 illustrates the word *RIVER* superimposed on a picture of a *baseball player*. All context displays contained both a picture and a word.

After subjects viewed each context display, they were shown a test display. Each test display contained either another picture or another word. Half the time, the test display contained another picture, and we referred to those trials as Picture trials; half the time, the test display contained another word, and we referred to those trials as Word trials. Subjects were told before each trial whether that trial would be a Picture trial or a Word trial.

The top panel of Figure 7 illustrates a Picture trial. On Picture trials, subjects were told to focus on the picture in the context display and ignore the word. For example, for the Picture trial shown in Figure 7, subjects should have focused on the *ashtray* and ignored the word *INCH*. Following each context display, subjects were shown a test display. On the Picture trials, the test display contained another picture. The subjects' task (on Picture trials) was to verify whether the picture shown in the test display was related to the picture shown in the context display. For the Picture trial shown in Figure 7, subjects should have responded "yes," because the picture shown in the test display, the *pipe*, was related to the picture shown in the context display, the *ashtray*.

The bottom panel of Figure 7 illustrates a Word trial. On Word trials, subjects were supposed to focus on the word in the context display and ignore the picture. For example, for the Word trial shown in Figure 7, subjects should have focused on the word *RIVER* and ignored the *baseball player*. The test display on Word trials contained another word. The subjects' task was to verify whether the word written in the test display was related to the word written in the context display. For the Word trial shown in Figure 7, subjects should have responded "yes," because the word written in the test display, *STREAM*, was related to the word written in the context display, *RIVER*.

On 40 Picture trials and 40 Word trials, the test display was related to what the subjects were to focus on in the context display, just as they are in

Figure 7. However, we were more interested in the 80 trials in which the test display was unrelated to what the subjects were supposed to focus on in the context display. On half of those trials, the test display was unrelated to what the subjects were to focus on in the context display, but it was related to what they were supposed to ignore.

For example, the top panel in Figure 8 illustrates an experimental Picture

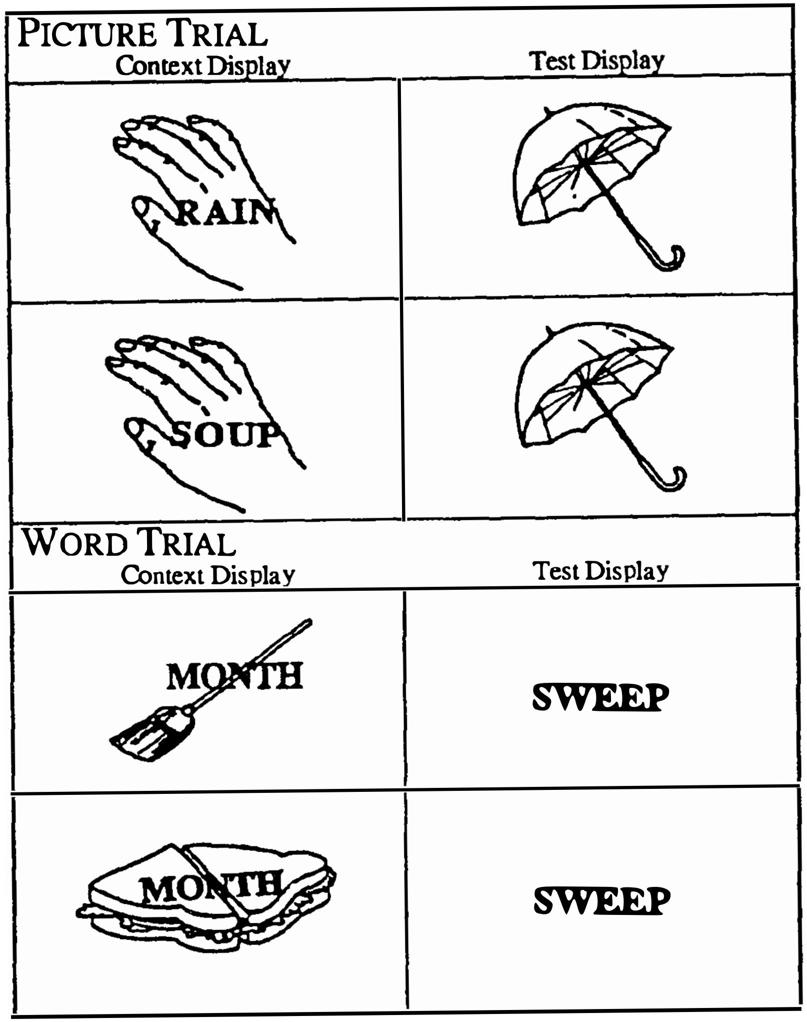

Figure 8. (From Gernsbacher & Faust, in press)

trial. The context display contains a picture of a *hand* with the superimposed word *RAIN*. Because this is a Picture trial, subjects should have focused on the picture (the *hand)* and ignored the word. The test display is a picture of an *umbrella*. So the test display, the *umbrella*, is unrelated to what the subjects were supposed to focus on in the context display, the *hand;* therefore, the subjects should have responded "no." But the test display is related to what the subjects were supposed to ignore, the word *RAIN*. We measured how long subjects took to reject the test display, the picture of the *umbrella*, after viewing the context display, the picture of the *hand* with the superimposed word *RAIN*. And we compared that to how long subjects took to reject the same test display, the picture of the *umbrella*, after viewing the same context display, the picture of the *hand*, but with another word superimposed, *SOUP*. This comparison showed us how quickly comprehenders could suppress information across modalities.

Experimental Word trials worked similarly, as illustrated by the third panel of Figure 8. When reading this context display, subjects should have focused on the word *MONTH* and ignored the surrounding picture of a *broom*. Then, they should have rejected the test display, the word *SWEEP*, because it is unrelated to the word *MONTH*. We compared how long subjects took to reject the word *SWEEP* after reading the word *MONTH* surrounded by the *broom*. And we compared that to how long subjects took to reject *SWEEP* after viewing the same context display with the picture of a *broom* replaced by a picture of a *sandwich* (as illustrated by the bottom panel of Figure 8). This comparison showed us how quickly comprehenders could suppress information across modalities.

As in our other experiments, we presented the test displays at two test points: Immediately (50 ms) after the context-setting display, and after a one-second delay. We predicted that at the immediate test point, both the more- and less-skilled comprehenders would take longer to reject a test display when it was related to the ignored picture or word in the context display. This result would corroborate Tipper and Driver (1988). This result would also demonstrate that for both more-and less-skilled comprehenders, ignored pictures or words are often activated.

Our novel predictions concerned what would happen after the delay. We predicted that after the one-second delay, the more-skilled comprehenders would not take longer to reject test displays when they were related to the ignored pictures or words. After one second, more-skilled comprehenders should be able to successfully suppress information across modalities. We made a different prediction for our less-skilled comprehenders. If less-skilled comprehenders are characterized by less-efficient suppression mechanisms, then even after the one-second delay, the ignored pictures and words should still be highly activated.

Figure 9 displays our 160 subjects' data. We estimated activation by sub-

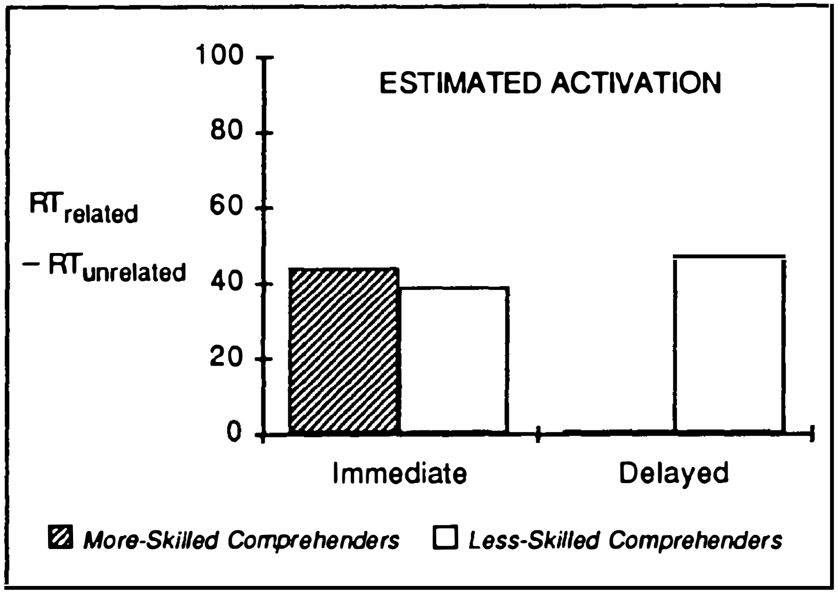

Figure 9. Subjects' average activation scores (from Gernsbacher & Faust, in press)

tracting subjects' latencies to reject test displays that were unrelated to ignored pictures/words from their latencies to reject test displays that were related to ignored pictures/words.[3] First examine what happened at the immediate test test point. As Figure 9 illustrates, immediately after both the more- and less-skilled comprehenders saw the context displays, the ignored pictures/words were highly activated; in fact, they were almost equally activated for the more- versus less-skilled comprehenders. So, 50 ms after viewing pictures with superimposed words or reading words surrounded by pictures, comprehenders of both skill levels have difficulty suppressing related pictures or words, even when they are told explicitly to ignore them.

Now examine what happened after the one-second delay. As Figure 9 illustrates, one second after the more-skilled comprehenders saw the context displays, the ignored pictures/words were no longer reliably activated; the more-skilled comprehenders had successfully suppressed them. But as Figure 9 also illustrates, the less-skilled comprehenders were less fortunate: Even after the one-second delay, the ignored pictures/words were still highly activated; in fact, they were as activated after the delay as they were immediately. So, a second after less-skilled comprehenders view pictures with superimposed words or read words surrounded by pictures, they still have difficulty suppressing the ignored pictures or words. These data support the hypothesis that less-skilled comprehenders are plagued by less-efficient suppression mechanisms.

Are Less-skilled Comprehenders Simply Less Appreciative of Sentence Context?

The three experiments we just described demonstrate that less-skilled comprehenders are less able to reject inappropriate meanings of ambiguous words, incorrect forms of homophones, and ignored pictures and words.We suggest that this inability arises because less-skilled comprehenders are plagued by less-efficient suppression mechanisms.

Another explanation is that less-skilled comprehenders have difficulty rejecting inappropriate information because they less fully appreciate what is contextually appropriate. For instance, by this logic, less-skilled comprehenders have difficulty rejecting *ACE* after reading *He dug with the spade* simply because they less fully appreciate that the context of ***digging*** *with a spade* implies a garden tool, not a playing card.

This explanation seems unlikely given the repeated finding that less-skilled comprehenders are not less appreciative of predictable sentence contexts — just the opposite: Less-skilled comprehenders often benefit from predictable contexts more than more-skilled comprehenders do. For example, the word *dump* is very predictable in the following context:

(20) The garbage men had loaded as much as they could onto the truck. They would have to drop off a load at the garbage *dump*.

In contrast, *dump* is less predictable in the following context:

(21) Albert didn't have the money he needed to buy the part to fix his car. Luckily, he found the part he wanted at the *dump*.

All comprehenders pronounce the word *dump* more rapidly when it occurs in the very predictable context than when it occurs in the less predictable context; in other words, all comprehenders benefit from the predictable contexts. But less-skilled comprehenders benefit even more than more-skilled comprehenders (Perfetti & Roth, 1981).

We evaluated this counter-explanation with adult comprehenders and a task similar to those we used in our previous experiments. Subjects read short sentences, and following each sentence, they saw a test word. As in our other experiments, the subjects' task was to verify whether the test word fit the meaning of the sentence they just read. However, in this experiment we were most interested in the 80 trials in which the test word *did* indeed match the meaning of the sentence (and, therefore, the subjects should have responded "yes").

On half of those trials, the last word of the sentence was an ambiguous word, for example, *spade*, and the verb in the sentence biased one meaning of the ambiguous word, for example,

(22) He *dug with* the spade.

The test word was related to the meaning of the ambiguous word that was biased by the verb, for example, *GARDEN*. In a comparison condition we presented the same sentence, but the biasing verb was replaced with a neutral verb, for example,

(23) He *picked up* the spade.

The *spade* in sentence (23) could be either a garden tool or a playing card.

We measured how rapidly subjects accepted test words after reading sentences with biasing verbs versus neutral verbs.[4] This comparison showed us how fully comprehenders could appreciate the biasing contexts: The faster subjects were to accept *GARDEN* after reading the sentence with the biasing verb phrase *dug with* versus the neutral verb phrase *picked up*, the more fully they appreciated the biasing context.

We presented the test words at two test points: Immediately (100 ms) after subjects finished reading each sentence, and after a one-second delay. We predicted that both the more- and less-skilled comprehenders would benefit from the biasing contexts; that is, both groups of comprehenders would accept test words more rapidly when the sentences contained biasing as opposed to neutral verbs. However, we were especially interested in whether the less-skilled comprehenders would benefit less than the more-skilled comprehenders.

If less-skilled comprehenders are less able to reject contextually inappropriate information (as we found in our previous experiments) because they are less appreciative of context, then the less-skilled comprehenders should have benefitted less from the biasing contexts. In contrast, if less-skilled comprehenders are less able to reject inappropriate information because they have less efficient suppression mechanisms, then the less-skilled comprehenders should have benefitted just as much from the biasing contexts as the more-skilled comprehenders did. Based on previous experiments, we predicted that the less-skilled comprehenders would benefit even more from the biasing contexts than the more-skilled comprehenders did.

Figure 10 displays our 120 subjects' data. We estimated activation by subtracting subjects' latencies to accept test words like *GARDEN* after reading sentences with biasing verbs like *dug with* from their latencies to accept *GARDEN* after reading sentences with unbiased verbs like *picked up*. As Figure 10 illustrates, at both the immediate and the delayed test test points, the biased verbs led to greater activation, and this occurred for both more- and less-skilled comprehenders. Indeed, as Figure 10 also illustrates, at both test test points, the less-skilled comprehenders benefitted from the biasing verbs more than the more-skilled comprehenders benefitted. These data do not support the hypothesis that less-skilled comprehenders are are less able to reject contextually inappropriate information because they are less appreciative of context.

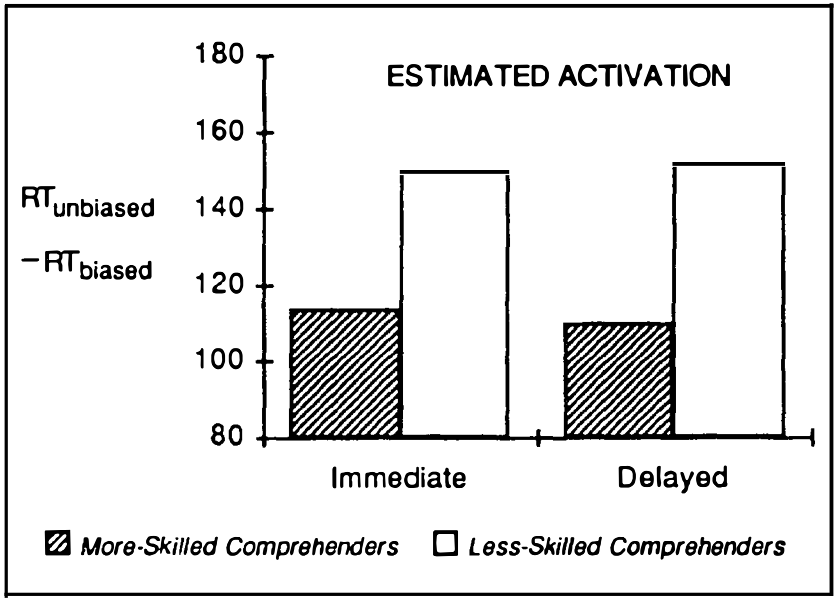

Figure 10. Subjects' average activation scores (from Gernsbacher & Faust, in press)

Conclusions

The experiments we have described here demonstrate the vital role that the mechanism of suppression plays in comprehension. Suppression helps fine tune the meanings of ambiguous words by decreasing the activation of the contextually inappropriate meanings. Indeed, less-skilled comprehenders are less able to suppress contextually inappropriate meanings. Less-skilled comprehenders are also less able to suppress the incorrect forms of homophones, and they are less able to suppress words while viewing pictures or suppress pictures while reading words.

The mechanism of suppression and the mechanism of enhancement play other important roles in sentence comprehension. For instance, the mechanisms of suppression and enhancement are vital to anaphoric reference (Gernsbacher, 1989). Anaphoric reference is the process by which speakers and writers use an anaphor, such as a repeated noun phrase or a pronoun, to refer to a previously mentioned concept (called an antecedent). Many anaphors improve their antecedents' accessibility by enhancing the activation of their antecedents (the concepts they refer to). Many anaphors also improve their antecedents' accessibility by suppression; they suppress the activation of other concepts (the concepts not referred to by the anaphors). When other concepts are suppressed, a rementioned concept can rise to the top of the queue of potential referents.

Anaphors differ in how much suppression and enhancement they trigger: The more explicit the anaphor, the more suppression and enhancement it triggers. Some anaphors are very explicit; for instance, repeated noun phrases match their antecedents exactly, and are, therefore, very explicit (e.g., *The man* went to the store. *The man* bought a quart of milk). Repeated noun phrase anaphors trigger a lot of suppression and enhancement, and they do so immediately. Other anaphors are less explicit; for instance, zero anaphors provide no information about their antecedents (e.g., *The man* went to the store. *The man* bought a quart of milk). Zero anaphors trigger very little suppression and virtually no enhancement (e.g., *The man* went to the store and ø bought a quart of milk). Information from other sources that identifies antecedents, for instance, information from the semantic, syntactic, and pragmatic context, also triggers suppression (although not enhancement).

The mechanisms of suppression and enhancement are also crucial to a process we have called cataphoric access (Gernsbacher & Jescheniak, 1990; Gernsbacher & Shroyer, 1989). Just as there are anaphoric devices which enable access to previously mentioned concepts, we have demonstrated that there are cataphoric devices which improve access to subsequently mentioned concepts. For instance, we have demonstrated that spoken stress operates as a cataphoric device. The unstressed indefinite *this* also operates as a cataphoric device, as in *So this man walks into a bar and says* ... One way that cataphoric devices improve their concepts' accessibility is by enhancing the activation of the concepts they mark. Another way that cataphoric devices improve their concepts' accessibility is by triggering the suppression of other concepts. And a third way that cataphoric devices improve their concepts' accessibility is by making the concepts they mark more resistant to being suppressed by other concepts.

Thus, the general cognitive mechanisms of suppression and enhancement play a vital role in language comprehension. These mechanisms enable us to build mental structures that represent sentences; in other words, these mechanisms enable us to understand sentences, which as Gough (1971) noted, is a paradoxically common but complex behavior.

Acknowledgements

This research was supported by National Science Foundation Grant BNS 85-10096, National Institutes of Health Research Career Development Award KO4 NS-01376, and Air Force Office of Sponsored Research Grants 89-0258 and 89-0111. We are indebted to Rachel R.W. Robertson for her invaluable contribution to our research program.

Notes

[1]Air Force recruits are high school graduates, and typically 20% have completed some college courses. Our subjects' ages ranged from 17 to 23, and approximately 18% were female.

[2]To ensure that the homophones would be familiar to our subjects, 25 students from the University of Oregon judged — without time pressure — whether the test words fit the meanings of our experimental and filler sentences. We only used experimental sentences and test words if 95% of our students agreed that the test words did *not* fit their sentences' meanings, and we only used filler sentences and test words if 95% of our students agreed that the test words *did* fit their sentences' meanings.

[3]Although both more- and less-skilled comprehenders responded more rapidly on Picture trials than Word trials, there were no interactions with modality (Picture vs Word). So, we have collapsed across this variable in our figures.

[4]To ensure that the biased verbs were biased and the neutral verbs were neutral, 25 students at University of Oregon read all of the experimental and comparison sentences and made unspeeded judgments about the meanings of the ambiguous words. We only used biased verbs if 95% of our students selected the meaning of the ambiguous word that we intended, and we only used neutral verbs if our students were roughly split over which meaning we intended (e.g., when given the sentence *He picked up the spade*, approximately 50% chose *GARDEN TOOL* and approximately 50% chose *PLAYING CARD)*.

References

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Becker, C. A. (1976). Semantic context and word frequency effects in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *2*, 556-566.

Blutner, R., & Sommer, R. (1988). Sentence processing and lexical access: The influence of the focus-identifying task. *Journal of Memory and Language*, *27*, 359-367.

Burgess, C., Tanenhaus, M. K., & Seidenberg, M. S. (1989). Context and lexical access: Implications of nonword interference for lexical ambiguity resolution. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 620-632.

Coltheart, M., Davelaar, E., Jonasson, J. T., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and Performance VI* (pp. 535-555). New York: Academic Press.

Coltheart, V., Laxon, V., Rickard, M., & Elton, C. (1988). Phonological recoding in reading for meaning by adults and children. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 14*, 387-397.

Conrad, C. (1974). Context effects in sentence comprehension: A study of the subjective lexicon. *Memory & Cognition, 2*, 130-138.

Cramer, P. (1970). *A study of homographs*. New York: Academic Press.

Doctor, E. A., & Coltheart, M. (1980). Children's use of phonological encoding when reading for meaning. *Memory & Cognition, 8*, 195-209.

Gernsbacher, M. A. (1989). Mechanisms that improve referential access. *Cognition, 32*, 99-156.

Gernsbacher, M. A. (in press-a). Cognitive processes and mechanisms in language comprehension: The structure building framework. In G. H. Bower (Ed.), *The psychology of learning and motivation*. New York: Academic Press.

Gernsbacher, M. A. (in press-b). *Language comprehension as structure building*. Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A., & Faust, M. (in press). The mechanism of suppression: A component of several comprehension skills. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

Gernsbacher, M. A., & Jescheniak, J. D. (1990). *Cataphoric devices in spoken discourse.* Manuscript submitted for publication.

Gernsbacher, M. A., & Shroyer, S. (1989). The cataphoric use of the indefinite *this* in spoken narratives. *Memory & Cognition, 17*, 536-540.

Gernsbacher, M. A., & Varner, K. R. (1988). *The multi-media comprehension battery* (No. 88-04). Institute of Cognitive and Decision Sciences, University of Oregon, Eugene, OR.

Gernsbacher, M. A., Varner, K. R., & Faust, M. (1990). Investigating differences in general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 430-445.

Glucksberg, S., Kreuz, R. J., & Rho, S. H. (1986). Context can constrain lexical access: Implications for models of language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 323-335.

Gough, P. B. (1971). (Almost a decade of) Experimental psycholinguistics. In W. O. Dingwall (Ed.), *A survey of linguistic science* . College Park, MA:

Kausler, D. H., & Kollasch, S. F. (1970). Word associations to homographs. *Journal of Verbal Learning and Verbal Behavior, 9*, 444-449.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review, 95*, 163-182.

Kintsch, W., & Mross, E. F. (1985). Context effects in word identification. *Journal of Memory and Language, 24*, 336-349.

Lucas, M. (1987). Frequency effects on the processing of ambiguous words in sentence context. *Language and Speech, 30*, 25-46.

Marslen-Wilson, W., Tyler, L. K., & Seidenberg, M. (1978). Sentence processing and the clause boundary. In W. J. M. Levelt, & G. B. Flores d'Arcais (Ed.), *Studies in the perception of language* (pp. 219-246). London: Wiley.

McClelland, J. L., & Kawamoto, A. H. (1986). Mechanisms of sentence processing: Assigning roles to constituents of sentences. In J. L. McClelland, & D. E. Rumelhart (Ed.), *Parallel distributed processing: Explorations in the microstructure of cognition* . Cambridge, MA: MIT Press.

Nelson, D., McEvoy, C. L., Walling, J. R., & Wheeler, J. W. (1980). The University of South Florida homograph norms. *Behavior Research Methods & Instrumentation, 12*, 16-37.

Norris, D. (1986). Word recognition: Context effects without priming. *Cognition, 22*, 93-136.

Perfetti, C. A., & Roth, S. (1981). Some of the interactive processes in reading and their role in reading skill. In A. M. Lesgold, & C. A. Perfetti (Ed.), *Interactive processes in reading* (pp. 269-297). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Rosson, M. B. (1985). The interaction of pronunciation rules and lexical representations in reading aloud. *Memory & Cognition, 13*, 90-99.

Seidenberg, M. S., Tanenhaus, M. K., Leiman, J. M., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology, 14*, 489-537.

Smith, M. C., & McGee, L. E. (1980). Tracing the time course of picture-word processing. *Journal of Experimental Psychology: General, 109*, 373-392.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18*, 645-659.

Tabossi, P. (1988). Accessing lexical ambiguity in different types of sentential contexts. *Journal of Memory and Language, 27*, 324-340.

Tabossi, P., Colombo, L., & Job, R. (1987). Accessing lexical ambiguity: Effects of context and dominance. *Psychological Research, 49*, 161-167.

Tanenhaus, M. K., Leiman, J. M., & Seidenberg, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior, 18*, 427-440.

Till, R. E., Mross, E. F., & Kintsch, W. (1988). Time course of priming for associate and inference words in a discourse context. *Memory & Cognition, 16*, 283-299.

Tipper, S. P., & Driver, J. (1988). Negative priming between pictures and words in a selective attention task: Evidence for semantic processing of ignored stimuli. *Memory & Cognition, 16*, 64-70.

van Orden, G. C. (1987). A rows is a rose: Spelling, sound, and reading. *Memory & Cognition, 15*, 181-198.

van Orden, G. C., Johnston, J. C., & Hale, B. L. (1988). Word identification in reading proceeds from spelling to sound to meaning. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 14*, 371-386.

van Petten, C., & Kutas, M. (1987). Ambiguous words in context: An event-related potential analysis of the time course of meaning activation. *Journal of Memory and Language, 26*, 188-208.

Waltz, D. L., & Pollack, J. B. (1985). Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science, 9*, 51-74.

Williams, J. N. (1988). Constraints upon semantic activation during sentence comprehension. *Language and Cognitive Processes, 3*, 165-206.