ATTENTUATING INTERFERENCE DURING COMPREHENSION: The Role of Suppression

Morton Ann Gernsbacher

I. Introduction

The goal of my research is to identify the cognitive processes and mechanisms that underlie language comprehension and comprehension in general. I have identified a few of those processes and mechanisms in a framework I call the Structure Building Framework (Gernsbacher, 1990, 1991a, 1995). According to the Structure Building Framework, the goal of comprehension is to build coherent mental representations or *structures.* These structures represent clauses, sentences, paragraphs, passages, and other meaningful units.

To build these structures, first, comprehenders lay foundations for their mental structures (Carreiras, Gernsbacher, & Villa, 1995; Gernsbacher & Hargreaves, 1988, 1992; Gernsbacher, Hargreaves, & Beeman, 1989). Then comprehenders develop their mental structures by mapping on information, when that incoming information coheres or relates to the previous information (Deaton & Gernsbacher, in press; Gernsbacher, 1996; Gernsbacher & Givón, 1995; Gernsbacher & Robertson, 1992; Haenggi, Gernsbacher, & Bolliger, 1993; Haenggi, Kintsch, & Gernsbacher, 1995). However, if the incoming information is less coherent, comprehenders employ a different process: They shift and initiate a new substructure (Foertsch & Gernsbacher, 1994; Gernsbacher, 1985). So, most mental representations comprise several branching substructures.

THE PSYCHOLOGY OF LEARNING
AND MOTIVATION, VOL. 37

0079-7421/97 $25.00

The building blocks of mental structures are what I refer to as memory nodes. According to the Structure Building Framework, memory nodes are activated by incoming stimuli. Once activated, the information they represent can be used by cognitive processes. Furthermore, according to the Structure Building Framework, activated memory nodes transmit processing signals. These processing signals either suppress or enhance the activation of other memory nodes. In other words, once memory nodes are activated, two mechanisms modulate their level of activation. The two mechanisms are suppression and enhancement.

Suppression decreases or dampens the activation of memory nodes when the information they represent is no longer as necessary for the structure being built. Enhancement increases or boosts the activation of memory nodes when the information they represent is relevant to the structure being built. By modulating the activation of memory nodes, suppression and enhancement contribute greatly to language comprehension. I want to stress, however, that suppression and enhancement are general cognitive mechanisms. They are not dedicated solely to language; they play vital roles in nonlinguistic processes, too. But language comprehension draws heavily on these two mechanisms. This chapter focuses on the mechanism of suppression.

While I believe that most people can appreciate that we need a mechanism that enhances relevant or related information, I have suggested that a mechanism that suppresses inappropriate or irrelevant information is perhaps even more crucial to the goal of comprehension, or in the words of the Structure Building Framework, the goal of building coherent mental representations. We need a mechanism of suppression because whenever we comprehend language, we experience various types of interference. Sometimes this interference during comprehension arises from the external environment, as when we conduct a conversation in a noisy restaurant or listen to a lecture with some clod in the audience whispering beside us.

At other times, interference during comprehension arises internally, as when we have to deal with the competing meanings of a word or phrase, or the alternate references of a pronoun. Indeed, even in a process as seemingly straightforward as reading a string of letters, such as *ROWS*, mental information that is related to that string of letters is often activated in our minds. This mental information might be orthographically related (such as the letter string *BOWS*), or phonologically related (such as the sound/roz/), or even semantically related (such as the concept "rose"). And indeed, laboratory experiments demonstrate that adults have difficulty quickly rejecting the letter string *ROWS* as not being a member of the semantic category, *flower* (van Orden, 1987; van Orden, Johnston, & Hale, 1988).

External information often interferes with our comprehension. For example, laboratory experiments demonstrate that it is harder to read a word when it is written within a line drawing of an object, and it is harder to name a line-drawn object if a word is written within it (Rayner & Posnansky, 1978; Rosinski, Golinkoff, & Kukish, 1975; Smith & McGee, 1980). Thus, successful comprehension involves successfully attenuating or inhibiting interfering information.

I have argued that a particular cognitive mechanism, what I call the cognitive mechanism of suppression, reduces such interference. In my previous research, I have empirically illustrated the crucial role that suppression plays in many comprehension phenomena. These phenomena include the following:

1. Lexical Access: how comprehenders understand or "access" from their memory the meanings of words (Faust, Balota, Duchek, Gernsbacher, & Smith, in press; Faust & Gernsbacher, 1996; Gernsbacher & Faust, 1991b, 1995; Gernsbacher & St. John, in press);
2. Anaphoric Reference: how comprehenders understand to whom or to what anaphors, such as pronouns, refer (Foertsch & Gernsbacher, 1994; Garnham, Traxler, Oakhill, & Gernsbacher, 1996; Gernsbacher, 1989);
3. Cataphoric Reference: how words that are marked by devices, such as spoken stress, gain a privileged status in comprehenders' memories (Gernsbacher & Jescheniak, 1995; Gernsbacher & Shroyer, 1989);
4. Syntactic Parsing: how we decode the grammatical forms of sentences into meaning (Gernsbacher & Robertson, 1996);
5. Surface Information Loss: the finding that seemingly superficial information, such as syntactic form, is often forgotten more rapidly than seemingly more important information, such as thematic content (Gernsbacher, 1985);
6. Metaphor Interpretation: how we understand figurative expressions such as "lawyers are sharks" (Gernsbacher, Keysar, & Robertson, 1995);
7. Inferencing: how comprehenders infer information that is only implied by a text or discourse (Carreiras & Gernsbacher, 1992; Gernsbacher, 1991b, 1994; Gernsbacher, Goldsmith, & Robertson, 1992; Gernsbacher & Robertson, 1992; Oakhill, Garnham, Gernsbacher, & Cain, 1992);
8. General Comprehension Skill: skill at comprehending linguistic and nonlinguistic media (Gernsbacher, 1993; Gernsbacher & Faust, 1991a, 1995; Gernsbacher & Robertson, 1995; Gernsbacher, Varner, & Faust, 1990).

In the remainder of this chapter, I briefly review some of these experiments. They demonstrate that the mechanism of suppression plays a powerful role in many language comprehension phenomena. Indeed, the role is so crucial that persons who are less skilled at comprehension are marked

by less efficiency in suppressing or inhibiting interfering information. Let me begin by describing the role of suppression in lexical access.

II. Attenuating Interference during Lexical Access

During lexical access, the cognitive mechanism of suppression attenuates the interference caused by other lexical information that is activated when a printed word is read, or a spoken word is heard. This information might be the meanings of a word that are not relevant to the immediate context, for example, the saloon meaning of *bar* in the pun *Two men walk into a bar and a third man ducks.* Or the interfering information might be other words or phrases that are related to the sound pattern of a spoken word or phrase, as in the classic *new display* often erroneously interpreted as *nudist play.*

Most models of lexical access propose that multiple types of information are activated when we read or hear a word; however, my research demonstrates that the mechanism of suppression dampens the activation of the unnecessary information. To examine the role of suppression in lexical access, I have capitalized on a phenomenon that I believe is a quintessential demonstration of the activation of superfluous information during lexical access. The phenomenon involves the comprehension of homonyms—words that share the same lexical form but differ in their meaning or origin, for example, *chest, tire, bowl, match, organ, head, plot, ring, nail.* All languages have homonyms (e.g., *cola, sal, calle, bota, trompa* in Spanish). Indeed, they are usually the most frequently occurring words in a language. Dahlgren has suggested that the average English word has three different meanings, and Britton has estimated that homonyms comprise about 40 percent of our most common open class words. Indeed, the more common the word, the more likely it is to have multiple meanings.

The homonym phenomenon that I have empirically studied is this: Immediately after we hear or read a homonym such as *match* or *duck*, multiple meanings are activated. And more intriguingly, this activation of multiple meanings occurs regardless of the semantic or syntactic context in which the homonym occurs. For example, immediately after we hear or read the homonym *match* in the sentence *He lit the match,* both the "fire stick" and the "competition" meanings are activated. Immediately after we hear or read the homonym *duck* in the sentence *He needed to duck,* both the "crouching" and the University of Oregon mascot meaning are activated.

Swinney (1979) and Tanenhaus, Leiman, and Seidenberg (1979) were among the first researchers to demonstrate this nonintuitive phenomenon, more than 15 years ago; it has been replicated numerous times since. These researchers also demonstrated—in line with our introspections—that the contextually inappropriate meanings of homonyms do not remain activated forever. What happens to these contextually inappropriate meanings? How do they become less activated? I have proposed that the cognitive mechanism of suppression dampens their activation. More specifically, I have hypothesized that memory nodes that represent a higher level structure—in this case the sentence-level structure—transmit processing signals to suppress the activation of the inappropriate lexical-level meanings. My research has provided several sources of converging evidence to support this proposal.

For example, in a series of laboratory experiments, postdoctoral fellow Mark Faust and I empirically demonstrated that suppression and not decay reduces the activation of inappropriate meanings (Gernsbacher & Faust, 1991b). That is, inappropriate meanings do not lose activation over time simply because their activation fades with time. We also empirically ruled out a mental "winner takes all" explanation: When inappropriate meanings become less activated it is not because the more appropriate meanings have become more activated. In other words, the inappropriate and appropriate meanings are not "slugging it out *mano a mano*"; rather, the source of the activation reduction comes from a higher level.

Indeed, using a parallel distributed processing network, postdoctoral fellow Mark St. John and I computationally demonstrated how sentence-level suppression can dampen the activation of contextually inappropriate word meanings (Gernsbacher & St. John, in press). In our connectionist network, suppression driven by a sentence-level representation, what St. John refers to as a gestalt level of representation, was the only type of top-down feedback we allowed, and that alone allowed us to perfectly simulate the behavioral data.

Further demonstrating that suppression and not simply decay is the mechanism responsible for decreasing the activation of the inappropriate meanings of homonyms, laboratory coordinator Rachel Robertson and I empirically demonstrated that suppression carries costs (Gernsbacher & Robertson, 1994). After subjects read a sentence such as *He lit the match,* they were considerably slower and considerably less accurate at simply verifying that the sentence *He won the match* made sense. If, after reading the sentence *He lit the match,* the inappropriate meaning of match simply decayed, that is, the competition meaning of match simply returned to base line, that meaning should not have been harder to activate in order to comprehend the subsequent sentence. Indeed, Gernsbacher and Faust (1995) created a laboratory condition in which it behooved subjects to suppress the inappropriate meanings of homonyms, and we discovered that

subjects employed suppression more rapidly in this condition than they did in a condition in which the need for suppression occurred only rarely.

Furthermore, as I describe at the end of this chapter, I have conducted an extensive series of experiments demonstrating that individuals who are less efficient at suppressing many types of information, for example, the color of ink in a Stroop color-naming task, hold onto inappropriate meanings considerably longer than do individuals who are more efficient in suppressing extraneous information. And most recently, postdoctoral fellow Faust and I discovered a right-visual-field, left-cerebral-hemisphere advantage for suppressing the inappropriate meanings of homonyms (Faust & Gernsbacher, 1996). When we presented homonyms to the left visual field (thereby hypothetically stimulating the right hemisphere prior to the left hemisphere), resolution of homonym meanings was slightly delayed. Although we still have miles to go before being able to stake our explorers' flag atop the cerebral location of our putative suppression mechanism, we find it less plausible that a decay mechanism would be similarly lateralized.

From all of these findings, I conclude that the mechanism of suppression, which enables the attenuation of interfering mental activation, such as the inappropriate meanings of homonyms, plays a crucial role in lexical access. I turn now to a review of the research I have conducted that investigates the role of suppression in anaphoric reference.

III. Attenuating Interference during Anaphoric Reference

Anaphoric reference is the process by which readers or listeners understand to whom or to what an anaphor, such as a pronoun, refers. In a series of experiments, I discovered that suppressison enables anaphoric reference by attenuating the interference caused by the activation of other referents. By other referents I mean the people or things to whom or to which an anaphoric expression does not refer. For example, consider the sentence *Ann predicted that Pam would lose the track race, but she came in first very easily.* In this sentence, the pronoun *she* is an anaphoric device, which most people interpret to refer to the referent *Pam.* I discovered that correctly interpreting such anaphoric devices is not so much a matter of activating one of the two possible referents: Both are highly activated because they were just mentioned in the first clause. Rather, understanding to whom the pronoun *she* in the second clause refers depends on how quickly comprehenders can reduce the activation of the referent to whom the pronoun *she* does not refer (i.e., *Ann* in the example sentence).

In my experiments, subjects read sentences word by word. The first clause of each sentence introduced two participants, for example, *Ann* and *Pam,*

as in *Ann predicted that Pam would lose the track race.* In the second clause, one of those two participants was referred to anaphorically, using either a very explicit repeated name anaphor, such as *Pam,* or a less explicit pronominal anaphor, such as *she,* as in *but she came in first very easily.* I measured activation of the anaphors' referents (e.g., *Pam*) and what I called the nonreferents (e.g., *Ann*) using the probe verification task. Subjects were shown a test name, for example, "Pam" or "Ann," or a name that had not occurred in the sentence, and their task was to verify whether the test name had occurred in the sentence. Presumably, the faster subjects respond to the test name, the more activated the participant represented by that test name is. In half the experimental sentences the referent was the first-mentioned participant, and in half the referent was the second-mentioned participant, as *Pam* was in the example sentence. In my first experiment I measured activation immediately before versus immediately after the name versus pronoun anaphors occurred. The first test point served as a baseline.

I observed that immediately after the very explicit name anaphors were read, the referents were considerably more activated than they were before; that is, reaction times decreased. More intriguingly, immediately after the very explicit name anaphors were read, the nonreferents were considerably less activated than they were before; that is, reaction times increased. By rementioning one participant, the other participant decreased in activation. However, this pattern occurred only for the very explicit name anaphors. For the pronouns, neither the referents nor the nonreferents changed in the activation.

This pattern has been replicated in English (MacDonald & MacWhinney, 1990), Spanish (Carreiras, 1997), Korean (Lee, 1992) and American Sign Language (Emmorey, in press). These data suggest that very explicit repeated name anaphors immediately lead to the suppression of nonreferents. In contrast, less explicit—and indeed momentarily ambiguous—pronoun anaphors do not immediately lead to suppression.

In a further experiment (Gernsbacher, 1989), I measured activation immediately before repeated-name versus pronoun anaphors, as I did before, and, again, this before-the-anaphor test point served as a baseline. However, in this experiment my comparison test point was at the end of the sentence, after the semantic/pragmatic information (which could disambiguate the syntactically ambiguous pronouns) had occurred. For example, activation was measured at the two test points indicated by asterisks in the following example sentence: *Ann predicted that Pam would lose the track race, but *Pam/she came in first very easily.** By the end of the sentence, even the gender-ambiguous pronoun anaphors had led to a reliable amount of suppression of the nonreferents.

In yet a further experiment (Gernsbacher, 1989), I placed the contextual information before the anaphors, as in *Ann lost a track race to Pam. Enjoying the victory, Pam/she headed toward the shower,* or *Ann lost a track race to Pam. Accepting the defeat, Ann/she headed toward the shower.* Despite the context preceding the anaphors, the less explicit pronoun anaphors still did not lead to a reliable amount of suppression until the end of the sentence. Thus, information from outside an anaphor can also trigger suppression, although it does so more slowly and less powerfully. This is good, because with zero anaphors, as in *Ann lost a tennis match to Pam and 0 cried all the way home,* the anaphor provides no information about its referent. All the information is provided by the semantic, pragmatic, and syntactic context. Therefore, zero anaphors should be the least effective at triggering suppression, a prediction confirmed by Corbett and Chang (1983). Together, the experiments I have described here demonstrate the role of suppression in enabling anaphoric reference: Suppression enables anaphoric reference by attenuating the interference caused by other referents.

IV. Attenuating Interference during Cataphoric Reference

Just as anaphoric devices enable reference to previously mentioned concepts, cataphoric devices enable reference to subsequently mentioned concepts. Cataphoric devices include such overt markers as vocally stressing a word in spoken discourse, or boldfacing a word in printed text. Presumably, speakers and writers mark certain concepts with cataphoric devices because those concepts will play a key role in the text or discourse. Thus, it would behoove listeners and readers if those key concepts had a privileged status in their mental structures.

Master's student Suzanne Shroyer and I (Gernsbacher & Shroyer, 1989) demonstrated that in spoken English, the unstressed, indefinite article *this,* as in *So this man walks into a bar,* as opposed to *So a man walks into this bar,* operates as a cataphoric device. The indefinite *this* is a relative newcomer to English; Wald (1983) dates its use back to the late 1930s. It occurs almost exclusively in informal spoken dialects rather than in formal or written ones, although I have observed personally that many of my e-mail pen pals use the indefinite *this* in written e-mail.

Because it is an indefinite article, *this* is used to introduce new concepts into discourse. Indeed, of the 243 occurrences of the indefinite *this* that Prince (1981) observed in Stud Terkel's book *Working,* 242 introduced a distinctly new concept. More interestingly—particularly with regard to my conjecture that the indefinite *this* operates as a cataphoric device to enable subsequent reference—in 209 of the 243 occurrences of the indefinite *this,*

the concept introduced with the indefinite *this* was referred to again. Similarly, when Wright and Givón (1987) recorded 8- and 10-year-old children telling one another stories and jokes, they found that when the children introduced concepts with the indefinite *this,* they referred to those concepts an average of 5.32 times in their next 10 clauses. When the children introduced concepts with the indefinite *a/an,* they referred to those concepts only .68 times in their next 10 clauses. These descriptive data suggest that speakers use the indefinite *this* to introduce key concepts. We (Gernsbacher & Shroyer, 1989) tested this proposal experimentally.

We presented spoken narratives to college students, telling them that at some point in each narrative the original narrator would stop talking; when that happened, it was their job to continue. For example, subjects heard the following narrative: *I swear, my friend Vicki, every time we go to a garage sale, she just 'uh, she just goes crazy. I mean like last Saturday we went to one near campus, 'n she just had to buy* **this/an** *ashtray, n'man . . . .* As this example illustrates, the last clause of the part of each narrative that subjects heard introduced a new concept, for example, *ashtray.* We manipulated whether this concept was introduced with the indefinite *this* (*this ashtray*) or the indefinite *a/an* (*an ashtray*). When we introduced the concepts with the indefinite *this,* subjects mentioned those concepts considerably more frequently, virtually always within the first clauses that they produced, and usually with less explicit anaphors such as pronouns. In contrast, when we introduced the concepts with the indefinite *a/an,* subjects mentioned the concepts less frequently, and typically with more explicit anaphors such as repeated noun phrases. (Through cross-splicing we ensured that the acoustic properties of the matched narratives and their critical concepts were otherwise identical.) These data demonstrate that concepts marked by cataphoric devices, such as the indefinite this, are more salient in listeners' mental representations.

Postdoctoral fellow Joerg Jescheniak and I (Gernsbacher & Jescheniak, 1995) discovered the role that the cognitive mechanism of suppression plays in enabling this privileged status: Suppression enables cataphoric reference by attenuating the interference caused by the introduction of other concepts. In this way, a cataphorically marked concept gains that privileged status in comprehenders' mental representations, so that it can be referred to more easily. Consider as an analogy a call for volunteers during which the entire line of candidates steps back, save one. The one candidate who did not step back—who was not suppressed—becomes most accessible for selection.

Subjects in our (Gernsbacher & Jescheniak, 1995) experiments also heard narratives, like the "Vicki going to a garage sale" narrative. We manipulated the indefinite *this* in some experiments, and in other experiments we manip-

ulated a seemingly more powerful cataphoric device, contrastive intonational stress. Using a verification task, we measured activation of the experimental concepts. In addition to cataphorically marked concepts being more activated—in other words, enhanced—we also found that cataphorically marked concepts are very resilient to being suppressed by a subsequently introduced concept. For example, in the "Vicki going to the garage sale" narrative, after we introduced the experimental concept (i.e., *this/an ashtray*), we introduced a new concept, *vase* (i.e., *'n she just had to buy* **this/an** *ashtray, 'n man, then she saw a vase. . .*). We observed that the previously mentioned concept, *ashtray,* was greatly attenuated in its activation when it was not marked by a cataphoric device. However, when the previously mentioned concept was marked by a cataphoric device, it was just as activated after we introduced a new concept as it was immediately after it was introduced.

Thus, cataphoric devices—the indefinite *this* and contrastive, intonational stress—attenuates the interference caused by introducing other concepts. By attenuating the interference from other concepts, cataphoric devices lead to a privileged status. Furthermore, the two cataphoric devices differ in how powerfully they lead to this privileged status; suppression is more powerfully triggered by contrastive stress than by the indefinite *this.* This difference makes sense: Contrastive stress is considerably more marked; it is a very iconic way of emphasizing a word in spoken discourse, similar to boldfacing a word in written text. The indefinite *this* is considerably more subtle; many of us are unaware of our informal use of it. Indeed, our undergraduate research assistants, whom we typically keep blind to our experimental manipulations, were stymied in their attempts to figure out what we were doing in the indefinite *this* experiments. So, like anaphoric devices, the strength of the suppression signals triggered by cataphoric devices is a function of the cataphoric devices' markedness.

All the experiments that I have described so far demonstrate the role that suppression plays in attenuating lexical- or concept-level interference. I have also examined the role of suppression in attenuating sentence-level interference.

V. Attenuating Interference during Syntactic Parsing

Motivated by the adage *Time flies like an arrow; fruit flies like a banana,* often attributed to Groucho Marx, laboratory coordinator Rachel Robertson and I hypothesized a role that the mechanism of suppression might play in syntactic parsing (Gernsbacher & Robertson, 1996). We proposed that suppression attenuates the interference caused by parsing a previous syntactic form. As the *time flies/fruit flies* example demonstrates, once we have parsed the phrase *time flies* as a noun and verb, it is difficult not to parse the phrase *fruit flies* in the same way. We (Gernsbacher & Robertson, 1996) examined a more stringent type of interference by using phrases such as *visiting in-laws,* which can be interpreted either as a plural noun phrase (i.e., people who are related to one's spouse and come to visit) or as a gerundive nominal (i.e., the act of visiting people who are related to one's spouse).

In our experiments (Gernsbacher & Robertson, 1996), we preceded sentences containing phrases such as *visiting in-laws* with sentences that required a similar or conflicting syntactic parse. For example, subjects first read *Washing dishes is a drag,* and then read *Visiting in-laws are, too.* Or subjects first read *Whining students are a drag,* and then read *Visiting in-laws is, too.* The subjects' task was to read each sentence and simply decide whether it was grammatical. We found that subjects were considerably slower and frighteningly less accurate to say that a sentence such as *Visiting in-laws are, too* was grammatical after they read the sentence *Washing dishes is a drag.* Similarly, subjects were considerably slower and less accurate to say that the sentence *Visiting in-laws is, too* was grammatical after they read the sentence *Whining students are a drag.* We interpreted these data as suggesting that correctly responding to the second sentence requires attenuating, or suppressing, the interference caused by the syntactic form in the first sentence.

We observed the same effect when we made the second sentences less syntactically dependent on the first sentence, by omitting the ellipses. For example, subjects were again slower and less accurate to say that the sentence *Visiting in-laws are a drag, too* was grammatical after they read the *Washing dishes* sentence. And, subjects were considerably slower and less accurate to say that the sentence *Visiting in-laws is a drag, too* was grammatical after they read the *Whining students* sentence. Furthermore, we observed the same effect when we made the second sentences syntactically independent of the first sentence, and the verb in the first sentence was not even marked for number. For example, subjects were still slower and still less accurate to say that the sentence *Visiting in-laws are a drag* was grammatical after they read the sentence *Washing dishes can be a bother,* and vice versa after they read the sentence *Whining students can be a bother.* This phenomenon underscores the need for suppression to attenuate the interference caused by a previous syntactic form.

VI. Attenuating Interference during Metaphor Interpretation

Rachel Robertson and I, in collaboration with Boaz Keyser (of the University of Chicago) have also explored the role of suppression in metaphor

interpretation. According to Glucksberg and Keysar (1990), when we interpret a metaphor such as *Lawyers are sharks,* we should enhance attributes of the metaphor's vehicle (e.g., *sharks*) that are common to the metaphor's topic (e.g., *lawyers*). So, after comprehending the metaphor *Lawyers are sharks,* we should enhance the facts that sharks are tenacious, fierce, and aggressive, among other attributes. We augmented Glucksberg and Keysar's (1990) theory by proposing that when we interpret a metaphor we also suppress the attributes that are not appropriate to (or concordant with) a metaphorical interpretation. So, for example, when we interpret the metaphor *Lawyers are sharks,* we might suppress attributes such as sharks being good swimmers, having fins, and living in the ocean. We tested both of these hypotheses by asking subjects to read a statement that might be metaphorical such as *Lawyers are sharks,* and then confirm the verity of a property statement such as *Sharks are tenacious.* In our first experiment, we used as a control condition statements that contained the same vehicle but a nonsensical topic, such as *Notebooks are sharks.*

We (Gernsbacher et al., 1995) found striking evidence that interpreting a metaphor such as *Lawyers are sharks* leads to both the enhancement of the attributes that are appropriate to the metaphorical interpretation and the suppression of attributes that are inappropriate to the metaphorical interpretation. For instance, subjects were faster to verify the statement *Sharks are tenacious* after they read the metaphor *Lawyers are sharks* than after they read the control statement *Notebooks are sharks.* This finding supports the hypothesis that interpreting a metaphor involves enhancing attributes that are appropriate to the metaphorical interpretation. In contrast, subjects were considerably slower to verify the statement *Sharks are good swimmers* after they read the metaphor *Lawyers are sharks* than after they read the control statement *Notebooks are sharks.* This finding supports the hypothesis that interpreting a metaphor involves suppressing attributes that are inappropriate to the metaphorical interpretation.

In a second experiment, we observed identical results when, instead of using a nonsensical statement as a baseline (control), we used a literal statement as a baseline. For example, we presented the literal statement *Hammerheads are sharks* as a baseline comparison for the metaphorical statement *Lawyers are sharks.* Again, we found striking evidence to support the hypothesis that interpreting a metaphor leads to both the enhancement of attributes that are appropriate to the metaphorical interpretation and the suppression of attributes that are inappropriate to the metaphorical interpretation. For example, again, subjects were faster to verify the statement *Sharks are tenacious* after they read the metaphor *Lawyers are sharks* than after they read the literal statement *Hammerheads are sharks.* And conversely, subjects were again considerably slower to verify the statements

Sharks are good swimmers after they read the metaphor *Lawyers are sharks* than after they read the literal statement *Hammerheads are sharks.* Therefore, both experiments demonstrated that interpreting a metaphor involves both enhancing the attributes that are relevant to the metaphorical interpretation and more intriguingly, suppressing the attributes that are not relevant to the metaphorical interpretation.

VII. Attenuating Interference during Inference Revision

When most of us hear or read that *George became too bored to finish the history book,* we infer that George is reading a very boring book. However, if we later hear or read that *George had already spent five years writing it,* we must revise our initially drawn inference because it was inappropriate. Brownell, Potter, Bihrle, and Gardner (1986) found that right-hemisphere-damaged patients had a whale of a time revising such inferences. They concluded that right-hemisphere-damaged patients' difficulty arose because they were unable to "let go of" the initial inferences that they drew. Perhaps revising such an inference is difficult because the revision requires suppressing the initially drawn inference. Thus, another role that suppression might play is to attenuate the interference caused by a previously drawn, but erroneous, inference.

We empirically tested this hypothesis, by investigating whether revising such inferences was difficult, not just for right-hemisphere-damaged patients but for "normal" college-aged adults. We constructed 40 two-sentence vignettes, similar to the *George became too bored to finish the history book. He had already spent five years writing it* example (other stimuli included *Jeff got a ticket after parking his car. As he headed into the movie theater, he handed the ticket to the usher; Sarah drove frantically all the way to the Emergency Room. She was already running 15 minutes late for her shift that evening; Jack painted the boat a bright red. Then he painted the ocean a deep blue and the sun a bright orange*). We measured how long subjects needed to read the inference-revising second sentence (e.g., *He had already spent five years writing it*) after they read the inference-inviting premise sentence (e.g., *George became too bored to finish the history book*). We compared how long subjects needed to read the inference-revising second sentence (e.g., *He had already spent five years writing it*) after they read the inference-inviting premise sentence or after they read a control premise sentence, which was a sentence that did not invite the inference; indeed it explicitly stated a different situation (e.g., *George became too bored to finish* writing *the history book; Jeff* bought *a ticket after parking his car; Sarah drove frantically all the way to* her job at *the Emergency Room;*

In the painting, *Jack painted the boat a bright red*). We found that subjects required substantially longer to read the second sentence (e.g., *He had already spent five years writing it*) after they read the experimental (inference-inviting) premise sentence (e.g., *George became too bored to finish the history book*) than after they read the control (inference-noninviting) premise sentence (e.g., *George became too bored to finish* writing *the history book*). We interpreted subjects' greater latency as reflecting their difficulty in suppressing a previously, but erroneously, drawn inference.

Furthermore, we found that members of a particular subject group—a group that we have previously identified to have difficulty quickly employing suppression (as I shall describe later)—were substantially slower to reject a test word that was related to the erroneously drawn inference, even after they read the inference-revising second sentence. For example, members of this group of subjects took longer to reject the test word *READ* after they read the inference-revising sentence *He had already spent five years writing it.* Members of this subject group were less skilled comprehenders, and this finding leads me to the last section of this chapter, the role of suppression in general comprehension skill.

VIII. Attenuating Interference and Comprehension Skill

A few years ago, honors student Kathy Varner, postdoctoral fellow Mark Faust, and I presented evidence in support of a construct we called "General Comprehension Skill" (Gernsbacher et al., 1990). Briefly put, we found that adults' skill in comprehending written language was highly correlated with their skill in comprehending spoken language, and both skills were highly correlated with comprehending nonverbal picture stories. We also found a critical characteristic of less skilled adult comprehenders: Less skilled adult comprehenders are less able to suppress quickly the inappropriate meanings of homonyms.

We (Gernsbacher et al., 1990) discovered this critical characteristic in the following way: We selected 64 more versus less skilled University of Oregon undergraduates on the basis of their performance on our Multi-Media Comprehension Battery (Gernsbacher & Varner, 1988). This battery tests reading, listening, and picture story comprehension. We drew the more skilled comprehenders from the upper third of a distribution of 270 subjects and the less skilled comprehenders from the bottom third. We invited these more and less skilled subjects to return to our lab (which was no easy feat, as the less skilled subjects did not have that much fun the first time they were there). When the subjects returned they read short

sentences; after reading each sentence they were shown a test word. Their task was to decide quickly whether the test word fit the meaning of the sentence that they had just read. On experimental trials, the final word of the sentence was a homonym, such as *spade,* as in *He dug with the spade.* The test word on these trials was related to a meaning of that homonym, but not the meaning implied by the sentence, for example, *ACE.* We compared how rapidly the more versus less skilled comprehenders could reject a test word that was related to the inappropriate meaning with how rapidly they could reject the same test word after reading a control sentence, for example, *He dug with the shovel.* The more time subjects took to reject *ACE* following the *spade* versus *shovel* sentence, the more interference we hypothesized they were experiencing from the inappropriate meaning. We measured this interference immediately (100 ms) after subjects finished reading the sentences and after an 850-ms delay.

Immediately after both the more and less skilled comprehenders read the homonyms, both groups experienced a reliable amount of interference, and, indeed, the two groups did not differ in the amount of interference they experienced at the immediate test point. In contrast, after the delay, the more skilled comprehenders were no longer experiencing a reliable interference, suggesting that they had successfully suppressed the inappropriate meanings of the homonyms. But for the less skilled comprehenders, they experienced the same amount of interference after the delay as they experienced immediately, suggesting that they were less able to quickly suppress the inappropriate meanings of the homonyms.

This pattern has been replicated by our colleagues around the world with whom we have shared our stimuli. For example, Leslie Twilley and Peter Dixon at the University of Alberta replicated this pattern using our measure of comprehension skill. Harvey Shulman at Ohio State replicated this pattern testing subjects who scored in the top versus bottom half of the ACT verbal test, but not the math test. Francesca Pazzaglia and her colleagues at the University of Padova replicated this pattern using Italian homonyms (and Italian subjects). Natasha Todorov at Macquarie University replicated this pattern with 7th-grade students selected according to their Nelson-Denny reading scores. In his dissertation Robert Crane at Washington State replicated this pattern testing university students, with small versus large reading spans. And we replicated this pattern testing United States Air Force recruits. Thus, this pattern replicates with Canadians, Italians, Australians, Buckeyes, and the U.S. military.

Furthermore, this pattern occurs when comparing members of other populations who hypothetically suffer from less efficient suppression with members of populations who are hypothesized to have more efficient suppression. For example, using our task and stimulus materials, McDowd and

Sundry at the University of Southern California found that healthy elderly subjects showed less efficient suppression compared with college-aged subjects. Elizabeth Schaunessy found that children diagnosed with attention deficit disorder showed less efficient suppression compared with children not diagnosed with attention deficit disorder. Mark Faust, David Balota, Janet Duchek, Stan Smith, and I found that patients with severe senile dementia of the Alzheimer's type showed extraordinarily inefficient suppression compared with patients with only moderate dementia compared with healthy age-matched controls (Faust et al., in press). Indeed, our dementia data show a dosing effect: The more severe the dementia, the more inefficient the suppression.

In all the data that I have reviewed, all the subjects showed initial interference from the inappropriate meanings, which I believe is crucial to demonstrate, but the members of the population hypothesized to suffer from less efficient suppression showed continued interference from the inappropriate meanings. Mark Faust and I (Gernsbacher & Faust, 1991b) also observed the same pattern when we examined how quickly less versus more skilled comprehenders could reject test words related to the incorrect forms of homophones, for instance, how quickly they could reject the test word *CALM* following the sentence *He had lots of patients,* versus the sentence *He had lots of students.* (Prior to collecting our data, we conducted a pilot experiment—no pun intended—to ensure that members of this population did know the correct spelling of these homophones.)

We (Gernsbacher & Faust, 1991b) also discovered that less versus more skilled comprehenders are not less able to reject the contextually inappropriate meanings of homonyms just because they do not know what is appropriate. We observed that less skilled comprehenders perform equally as well as more skilled comprehenders when the task is to accept the appropriate meaning of a homonym, for example, when their task is to correctly say "yes" that the test word *ACE* is related to the sentence *He dealt the spade.*

More recently, Rachel Robertson and I (Gernsbacher & Robertson, 1995) replicated our tried and true finding that less skilled comprehenders are worse than more skilled comprehenders when the task is to reject a test word that is related to the inappropriate meaning. For example, less skilled comprehenders are slower to reject the test word *ACE* after reading the sentence *He dug with the spade.* Presumably this is because less skilled comprehenders are less able to suppress the activation of the inappropriate meanings. We (Gernsbacher & Robertson, 1995) also replicated the finding I just mentioned, namely that less and more skilled comprehenders do not differ when the task is to accept the appropriate meaning. For example, less skilled comprehenders are just as fast as more skilled comprehenders

in accepting the test word *ACE* after reading the sentence *He dealt the spade.* Again, this suggests that less skilled comprehenders' difficulty in rejecting inappropriate meanings is *not* because they do not know what is appropriate.

However, to ensure that we didn't have a "Nancy Reagan effect" on our hands (i.e., less skilled comprehenders just can't say "no"), we (Gernsbacher & Robertson, 1995) also created a task in which the goal was to say "yes" to a meaning that was inappropriate, somewhat like what one needs to do to understand a pun. And we again found that less skilled comprehenders were worse than more skilled comprehenders. For example, less skilled comprehenders were slower to accept the test word *ACE* after reading the sentence *He dug with the spade,* perhaps because this task—accepting an inappropriate meaning—requires suppressing the appropriate meaning (recall how difficult it was to understand the pun *Two men walk into a bar and the third man ducks.* It is as though to understand the "metal bar" meaning, one needs to suppress the "tavern" meaning).

IX. Summary

To summarize, I have suggested that the cognitive mechanism of suppression attenuates interference in many language comprehension phenomena. During lexical access, the mechanism of suppression attenuates the interference caused by the activation of other lexical information, such as the inappropriate meanings of homonyms. During anaphoric reference, the mechanism of suppression attenuates the interference caused by the activation of other potential referents. In this way, the referent to which the anaphor does not refer becomes the most activated concept. Moreover, the strength of the suppression is a function of the markedness of the anaphoric device: More marked anaphors such as repeated proper names immediately lead to suppression; less marked anaphors such as pronouns take longer to enact suppression. During cataphoric reference, the mechanism of suppression attenuates the interference caused by the introduction of other concepts. In this way, a cataphorically marked concept gains a privileged status in comprehenders' mental representations. The more marked the cataphoric device is, the less interference is caused by the introduction of another concept. More marked cataphoric devices such as spoken stress protect their concepts more than less marked cataphoric devices such as the indefinite *this.* During syntactic parsing, the mechanism of suppression attenuates the interference caused by a previous syntactic form. During metaphor comprehension, the mechanism of suppression attenuates the interference caused by a literal interpretation. During infer-

encing, the mechanism of suppression attenuates the interference caused by an initial but inappropriate inference. Thus, my previous research has demonstrated the crucial role that suppression—and by that I mean a general, cognitive mechanism that attenuates interference—plays in many facets of language comprehension.

Acknowledgments

Preparation of this chapter was supported by grants from the National Institutes of Health (RO1 NS 29926) and the Army Research Institute (DASW0194-K-0004).

References

Brownell, H. H., Potter, H. H., Bihrle, A. M., & Gardner, H. (1986). Inference deficits in right brain-damaged patients. *Brain and Language, 27,* 310–321.

Carreiras, M. (1997). Manuscript in preparation.

Carreiras, M., & Gernsbacher, M. A. (1992). Comprehending conceptual anaphors in Spanish. *Language and Cognitive Processes, 7,* 281–299.

Carreiras, M., Gernsbacher, M. A., & Villa, V. (1995). The advantage of first mention in Spanish. *Psychonomic Bulletin & Review, 2,* 124–129.

Corbett, A. T., & Chang, F. R. (1983). Pronoun disambiguation: Accessing potential antecedents. *Memory & Cognition, 11,* 283–294.

Deaton, J. A., & Gernsbacher, M. A. (in press). Causal conjunctions and implicit causality cue mapping in sentence comprehension. *Journal of Memory and Language.*

Emmorey, K. A. (in press). Non-antecedent suppression in American Sign Language. *Language and Cognitive Processes.*

Faust, M. E., Balota, D. A., Duchek, J. A., Gernsbacher, M. A., & Smith, S. D. (in press). Inhibitory control during sentence processing in individuals with dementia of the Alzheimer type. *Brain and Language.*

Faust, M. E., & Gernsbacher, M. A. (1996). Cerebral mechanisms for suppression of inappropriate information during sentence comprehension. *Brain and Language, 53,* 234–259.

Foertsch, J., & Gernsbacher, M. A. (1994). In search of complete comprehension: Getting "minimalists" to work. *Discourse Processes,* pp. 271–296.

Garnham, A., Traxler, M. J., Oakhill, J., & Gernsbacher, M. A. (1996). The locus of causality effects in comprehension. *Journal of Memory and Language, 35,* 517–543.

Gernsbacher, M. A. (1985). Surface information loss in comprehension. *Cognitive Psychology, 17,* 324–363.

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

Gernsbacher, M. A. (1990). *Language comprehension as structure building.* Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1991a). Cognitive processes and mechanisms in language comprehension: The structure building framework. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 217–263). San Diego, CA: Academic Press.

Gernsbacher, M. A. (1991b). Comprehending conceptual anaphors. *Language and Cognitive Processes, 6,* 81–105.

Gernsbacher, M. A. (1993). Less skilled readers have less efficient suppression mechanisms. *Psychological Science, 4,* 294–298.

Gernsbacher, M. A. (1994). Activating knowledge of fictional characters' emotional states. In C. A. Weaver, S. Mannes, & C. R. Fletcher (Eds.), *Discourse comprehension: Essays in honor of Walter Kintsch* (pp. 141–155). Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1995). The Structure Building Framework: What it is, what it might also be, and why. In B. K. Britton & A. C. Graesser (Eds.), *Models of text understanding* (pp. 289–311). Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1996). Coherence cues mapping during comprehension. In J. Costermans & M. Fayol (Eds.), *Processing interclausal relationships in the production and comprehension of text* (pp. 3–21). Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A., & Faust, M. (1991a). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 245–262.

Gernsbacher, M. A., & Faust, M. (1991b). The role of suppression in sentence comprehension. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 97–128). Amsterdam: North-Holland.

Gernsbacher, M. A., & Faust, M. (1995). Skilled suppression. In F. N. Dempster & C. N. Brainerd (Eds.), *Interference and inhibition in cognition* (pp. 295–327). San Diego, CA: Academic Press.

Gernsbacher, M. A., & Givón, T. (Eds.). (1995). *Coherence in spontaneous text.* Philadelphia: John Benjamins.

Gernsbacher, M. A., Goldsmith, H. H., & Robertson, R. R. W. (1992). Do readers mentally represent characters' emotional states? *Cognition and Emotion, 6,* 89–111.

Gernsbacher, M. A., & Hargreaves, D. (1988). Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language, 27,* 699–717.

Gernsbacher, M. A., & Hargreaves, D. (1992). The privilege of primacy: Experimental data and cognitive explanations. In D. L. Payne (Ed.), *Pragmatics of word order flexibility* (pp. 83–116). Philadelphia: John Benjamins.

Gernsbacher, M. A., Hargreaves, D., & Beeman, M. (1989). Building and accessing clausal representations: The advantage of first mention versus the advantage of clause recency. *Journal of Memory and Language, 28,* 735–755.

Gernsbacher, M. A., & Jescheniak, J. D. (1995). Cataphoric devices in spoken discourse. *Cognitive Psychology, 29,* 24–58.

Gernsbacher, M. A., Keysar, B., & Robertson, R. R. W. (1995, November). *The role of suppression and enhancement in metaphor interpretation.* Paper presented at the 36th annual meeting of the Psychonomic Society, Los Angeles.

Gernsbacher, M. A., & Robertson, R. R. W. (1992). Knowledge activation versus sentence mapping when representing fictional characters' emotional states. *Language and Cognitive Processes, 7,* 353–371.

Gernsbacher, M. A., & Robertson, R. R. W. (1994, November). *The costs and benefits of meaning.* Paper presented at the 35th meeting of the Psychonomic Society, St. Louis, MO.

Gernsbacher, M. A., & Robertson, R. R. W. (1995). Reading skill and suppression revisited. *Psychological Science, 6,* 165–169.

Gernsbacher, M. A., & Robertson, R. R. W. (1997). *Parallel form affects sentence comprehension.* 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., & St. John, M. F. (in press). Modeling the mechanism of suppression in lexical access. In R. Klein & P. McMullen (Eds.), *Converging methods for studying reading and dyslexia.* Cambridge, MA: MIT Press.

Gernsbacher, M. A., & Varner, K. R. (1988). *The multi-media comprehension battery* (Tech. Rep. No. 88–03). Eugene: University of Orgeon, Institute of Cognitive and Decision Sciences.

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., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review, 97,* 3–18.

Haenggi, D., Gernsbacher, M. A., & Bolliger, C. M. (1993). Individual differences in situation-based inferencing during narrative text comprehension. In H. van Oostendorp & R. A. Zwaan (Eds.), *Naturalistic text comprehension: Vol. 53. Advances in discourse processing* (pp. 79–96). Norwood, NJ: Ablex.

Haenggi, D., Kintsch, W., & Gernsbacher, M. A. (1995). Spatial situation models and text comprehension. *Discourse Processes, 19,* 173–199.

Lee, J. (1992). *On-line processing of pronoun resolution in reading.* Unpublished doctoral dissertation, Korea University, Seoul.

MacDonald, M. C., & MacWhinney, B. (1990). Measuring inhibition and facilitation effects from pronouns. *Journal of Memory and Language, 29,* 469–492.

Oakhill, J., Garnham, A., Gernsbacher, M. A., & Cain, K. (1992). How natural are conceptual anaphors? *Language and Cognitive Processes, 7,* 257–280.

Prince, E. F. (1981). On inferencing of the indefinite-this NPs. In A. Joshi, B. Webber, & I. Sag (Eds.), *Elements of discourse understanding* (pp. 231–250). Cambridge, UK: Cambridge University Press.

Rayner, K., & Posnansky, C. (1978). Stages of processing in word identification. *Journal of Experimental Psychology: General, 107,* 64–80.

Rosinski, R. R., Golinkoff, R. M., & Kukish, K. S. (1975). Automatic semantic processing in a picture-word interference task. *Child Development, 46,* 247–253.

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.

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.

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.

Wald, B. (1983). Referents and topics within and across discourse units: Observations from current vernacular English. In S. Klen-Andrew (Ed.), *Discourse perspectives on syntax* (pp. 91–116). New York: Academic Press.

Wright, S., & Givón, T. (1987). The pragmatics of indefinite reference: Quantified text-based studies. *Studies in Language, 11,* 1–33.