

Good-Enough Representations in Language Comprehension

Fernanda Ferreira,¹ Karl G.D. Bailey, and Vittoria Ferraro

Department of Psychology and Cognitive Science Program, Michigan State University, East Lansing, Michigan

Abstract

People comprehend utterances rapidly and without conscious effort. Traditional theories assume that sentence processing is algorithmic and that meaning is derived compositionally. The language processor is believed to generate representations of the linguistic input that are complete, detailed, and accurate. However, recent findings challenge these assumptions. Investigations of the misinterpretation of both garden-path and passive sentences have yielded support for the idea that the meaning people obtain for a sentence is often not a reflection of its true content. Moreover, incorrect interpretations may persist even after syntactic reanalysis has taken place. Our good-enough approach to language comprehension holds that language processing is sometimes only partial and that semantic representations are often incomplete. Future work will elucidate the conditions under which sentence processing is simply good enough.

Keywords

language comprehension; satisficing; syntax; linguistic ambiguity

Over the past three decades, various theories of language comprehension have been developed to explain how people compose the meanings of sentences from indi-

vidual words. All theories advanced to date assume that the language-processing mechanism applies a set of algorithms to access words from the lexicon, organize them into a syntactic structure through rules of grammar, and derive the meaning of the whole structure based on the meaning of its parts. Furthermore, all theories assume that this process generates complete, detailed, and accurate representations of the linguistic input.

MODELS OF SENTENCE PROCESSING

Two approaches to sentence processing that have been widely contrasted are the *garden-path model* (Ferreira & Clifton, 1986; Frazier, 1978) and the *constraint-satisfaction model* (MacDonald, Pearlmutter, & Seidenberg, 1994; Trueswell, Tanenhaus, & Garnsey, 1994). According to the garden-path account, the language processor initially computes a single syntactic analysis without consideration of context or plausibility. Once an interpretation has been chosen, other information is used to evaluate its appropriateness. For example, a person who heard, "Mary saw the man with the binoculars," would tend to understand the sentence to mean that Mary used the binoculars as an instrument. If it turned out that the man had the binoculars, the initial interpretation would be revised to be compatible with that contextual knowledge.

Constraint-satisfaction theorists, in contrast, assume that all possible syntactic analyses are computed at once on the basis of all relevant sources of information. The analysis with the greatest support is chosen over its competitors. The constraint-based approach predicts that people who hear the sentence about Mary, the man, and the binoculars will activate both interpretations and then select the one that is more appropriate in the context. Thus, the two classes of models assume radically different approaches to sentence processing: According to the garden-path model, analyses are proposed serially, and syntactic information is processed entirely separately from real-world knowledge and meaning. According to constraint-based models, analyses are proposed in parallel, and the syntactic processor communicates with any relevant information source. Nevertheless, both models incorporate the assumption that interpretations of utterances are compositionally built up from words clustered into hierarchically organized constituents.

IS THE MEANING OF A SENTENCE ALWAYS THE SUM OF ITS PARTS?

This assumption of compositionality seems eminently plausible, but results in the literature on the psychology of language call it into question. For example, people have been observed to unconsciously normalize strange sentences to make them sensible (Fillenbaum, 1974). The *Moses illusion* (Erickson & Mattson, 1981) is typically viewed as demonstrating the fallibility of memory processes, but it is also relevant to issues of language interpretation and compositionality. When asked, "How many animals of each sort did Moses put on the ark?" people tend to respond

“two,” instead of objecting to the presupposition behind the question. Similarly, participants often overlook the anomaly in a sentence such as “The authorities had to decide where to bury the survivors” (Barton & Sanford, 1993).

A study conducted to examine whether sentence meaning can prime individual words (i.e., activate them so that they are more accessible to the comprehension system) also demonstrates that language processing is not always compositional, and that the semantic representations that get computed are shallow and incomplete (rather than computing the structure to the fullest degree possible, the comprehension system just does enough to contend with the overall task at hand; Duffy, Henderson, & Morris, 1989). Participants were asked to speak aloud the final word in various sentences after reading the sentences. On average, they took less time to say the word in biased sentences like (1) than in sentences such as (2), indicating that “cocktails” had been activated, or primed, earlier in the sentence. But, unexpectedly, the times were as fast for sentences like (3) as they were for sentences like (1), even though the word “bartender” has no semantic connection to “cocktails” in (3).

(1) The boy watched the bartender serve the cocktails.

(2) The boy saw that the person liked the cocktails.

(3) The boy who watched the bartender served the cocktails.

Clearly, the semantic representation that yielded priming in (1) and (3) was not detailed enough to distinguish the difference in meaning between the two sentences. The representation was “good enough” to provide an interpretation that satisfied the comprehender, but not detailed enough to distinguish the important differences in who was doing what to whom.

RECENT STUDIES OF WHETHER INTERPRETATIONS ARE GOOD ENOUGH

In two series of studies, our lab has been investigating some situations in which good-enough, or noncompositional, processing may occur.

Misinterpretations of Garden-Path Sentences

One series (Christianson, Hollingworth, Halliwell, & Ferreira, 2001) addressed the straightforward question whether people delete from memory their initial misinterpretation of a sentence after reanalysis. When people were visually presented sentence (4), they initially took “the baby” to be the object of “dressed.”

(4) While Anna dressed the baby played in the crib. (presented without commas)

As a result, readers spent a great deal of time processing the disambiguating word “played” and often reread the preceding material. Sentences such as this one are often termed *garden-path sentences*, because the first part of the sentence sends the language comprehension system in an ultimately wrong direction. The comprehender will have no difficulty with (4) if the clauses are separated by a comma or if the main clause is presented before the subordinate clause. In these cases, there is no temptation to take “the baby” to be the object of “dressed,” and therefore the reader has no difficulty integrating “played.”

It has generally been assumed that if comprehenders restructure their initial interpretation of (4) so as to make “the baby” the subject of the main clause, they will end up with an appropriate representation of the sentence’s overall mean-

ing. This assumption was tested by asking participants to respond to questions after reading (at their own pace) garden-path sentences or non-garden-path control versions of the same sentences (Christianson et al., 2001). The questions were of two sorts:

(5) Did the baby play in the crib?

(6) Did Anna dress the baby?

Question (5) assessed whether the phrase “the baby” was eventually taken to be the subject of “played.” Recall that initially it is not; the syntactic processor makes “the baby” the object of “dressed,” and so “played” ends up without a subject. Thus, successful syntactic restructuring requires that “the baby” be removed from that first clause and included in the second, making “yes” the correct answer to (5). Question (6) assessed whether comprehenders then adjusted the meaning of the sentence to correspond to that reanalysis: Under this reinterpretation, “the baby” is no longer the object of “dressed,” and so the sentence means that Anna is dressing herself. Therefore, the participants should have said “no” in response to (6).

Participants were virtually 100% correct in responding that the baby played in the crib. Performance was equally good in the garden-path and non-garden-path conditions. Yet when the sentence led the comprehenders down a syntactic garden path, they were inaccurate in answering (6). That is, people initially took “the baby” to be the object of “dressed.” Then, they restructured the sentence to make “the baby” the subject of “played,” but they persisted in thinking that the baby was being dressed. People who read the non-garden-path control version, however, almost always correctly replied that Anna did not dress the baby. In summary, the initial misinterpretation lingered and caused comprehenders to end up with a representation

in which “the baby” was both the subject of “played” and the object of “dressed.” This is clear evidence that the meaning people obtain for a sentence is often not a reflection of its true content.

Misinterpretations of Passive Sentences

The other series of experiments (Ferreira & Stacey, 2000) was designed to investigate an even more basic question: Are people ever tricked by simple, but implausible, passive sentences? Consider an active sentence like (7). People have little trouble obtaining its implausible meaning. In contrast, the passive sentence (10) is much more difficult to understand, and one’s impression is that it is hard to keep straight whether the dog is the perpetrator or the victim in the scenario.

(7) The man bit the dog.

(8) The man was bitten by the dog.

(9) The dog bit the man.

(10) The dog was bitten by the man.

In one experiment (Ferreira & Stacey, 2000), participants read sentences like (7) through (10) and were instructed to indicate whether the event described in each sentence was plausible. For the active sentences, people were almost always correct. However, they called passive sentences like (10) plausible more than 25% of the time. In another experiment, participants heard one of these four sentences and then identified either the agent or the patient of the action. Again, people were accurate with all sentences except (10). Thus, when people read or hear a passive sentence, they use their knowledge of the world to figure out who is doing what to whom. That interpretation reflects the content words of the sentence more than its compositional, syntactically derived meaning. It is as if people use a semantic heuristic rather than syntactic algo-

rithms to get the meaning of difficult passives.

OUR GOOD-ENOUGH APPROACH

The linguistic system embodies a number of powerful mechanisms designed to enable the comprehender to obtain the meaning of a sentence that was intended by the speaker. The system uses mechanisms such as syntactic analysis to achieve this aim. Syntactic structure allows the comprehender to compute algorithmically who did what to whom, because it allows thematic roles such as agent to be bound to the individual words of the sentence. The challenges in comprehension, however, are twofold. First, as the earliest work in cognitive psychology revealed, the structure built by the language processor is fragile and decays rapidly (Sachs, 1967). The representation needs almost immediate support from context or from schemas (i.e., general frameworks used to organize details on the basis of previous experience). In other words, given (10), syntactic mechanisms deliver the proper interpretation that the dog is the patient and the man is the agent; but the problem is that the delicate syntactic structure needs reinforcement. Schemas in long-term memory cannot provide that support, and so the source of corroboration must be context. Quite likely, then, sentences like this would be correctly understood in normal conversation, because the overall communicative context would support the interpretation. The important concept is that the linguistic representation itself is not robust, so that if it is not reinforced, a merely good-enough interpretation may result.

The second challenge to the linguistic system is that it must cope with potentially interfering infor-

mation. The garden-path studies show that an initial incorrect representation of a sentence lingers and interferes with obtaining the correct meaning for the sentence. In the case of implausible passive sentences, information from schemas in long-term memory causes interference. As a result, people end up believing that (10) means what their schema tells them rather than what the output of the syntactic algorithms mandates. This interfering information must be inhibited for comprehension to be successful.

FUTURE DIRECTIONS

Experiments are under way to examine the characteristics of the memory representations for garden-path sentences, and to focus on how misinformation is suppressed during successful comprehension. The studies on passives are intriguing because they demonstrate that complex syntactic structures can be misinterpreted, but what makes a structure likely to be misinterpreted? One of the experiments (Ferreira & Stacey, 2000) demonstrated that the surface frequency of the sentence form is not critical to determining difficulty. People were as accurate with sentences such as “It was the man who bit the dog” as they were with common active sentences, even though the former structure is rare. One possible explanation for why the passive structure is difficult to comprehend is that passives require semantic roles to be assigned in an atypical order: patient before agent. This hypothesis can be addressed by examining languages that permit freer word order than does English. We are currently focusing on the aboriginal Native American language Odawa, which orthogonally crosses voice and word order—that is, an active sen-

tence may have the patient either before or after the agent, as may a passive sentence. Thus, Odawa provides a unique opportunity for us to study the factors that cause linguistic representations to be particularly fragile and vulnerable to influence from schemas.

The good-enough approach also leads us in several other less traditional directions. For example, speech disfluencies that occur during conversation include pauses filled with “uh” or “um,” repeated words, repairs that modify or replace earlier material, and false starts (utterance fragments that are begun and abandoned). Disfluencies will often yield a string of words that violates grammatical principles. Nevertheless, comprehenders seem able to process such strings efficiently, and it is not clear how interpretation processes are affected by these disfluencies. Are abandoned fragments incorporated into the semantic representation of a sentence? Our work on misinterpretations of garden-path sentences suggests that the answer could well be yes. In the same way that the incorrect interpretation of a garden-path sentence lingers even though its underlying structure is ultimately corrected, an interpretation built upon an ultimately abandoned fragment (e.g., “Turn left—I mean right at the stop sign”) might persist in the comprehender’s overall representation.

We are also investigating whether syntactically ambiguous sentences such as (11) and (12) are given incomplete syntactic representations. A recent study found that people were faster at reading sentences like (11), for which the attachment of the relative clause is semantically ambiguous, than at reading semantically unambiguous versions like (12) (Traxler, Pickering, & Clifton, 1998).

(11) The son of the driver that had the mustache was pretty cool.

(12) The car of the driver that had the mustache was pretty cool.

One proposed explanation for this finding is that the syntactic representation in the ambiguous case remains underspecified. That is, perhaps the language processor does not bother to attach the relative clause “that had the mustache” to either “son” or “driver” because it does not have enough information to support one interpretation over the other.

More generally, the good-enough approach to language comprehension invites a more naturalistic perspective on how people understand utterances than has been adopted in psycholinguistics up to this point. Psycholinguists have focused on people’s ability to understand individual sentences (or short texts) in almost ideal circumstances. In laboratories, stimuli are (usually) shown visually in quiet rooms that offer no distractions. The results that have emerged from this work are central to any theory of comprehension, but examination of only those conditions will not yield a complete story. Outside the laboratory, utterances are often difficult to hear because of background noise; dialect and idiolect differences as well as competing sounds can make it difficult for the hearer to extract every word from an utterance; and speakers often produce utterances with disfluencies and outright errors that the processing system must handle somehow. We have shown in our research that, even in the ideal conditions of the laboratory, comprehension is more shallow and incomplete than psycholinguists might have suspected. In the real world, interpretations are even more likely to be “just good enough.”

Perhaps good-enough interpretations help the language system coordinate listening and speaking during conversation. Usually when

people talk to one another, turns are not separated by gaps. Therefore, comprehension and production processes must operate simultaneously. The goal of the comprehension system might be to deliver an interpretation that is just good enough to allow the production system to generate an appropriate response; after all, it is the response that is overt and that determines the success of the participants’ joint activity. An adequate theory of how language is understood, then, will ultimately have to take into account the dynamic demands of real-time conversation.

Recommended Reading

- Christianson, K., Hollingworth, A., Halliwell, J., & Ferreira, F. (2001). (See References)
- Clark, H.H. (1996). *Using language*. Cambridge, England: Cambridge University Press.
- Clifton, C., Jr. (2000). Evaluating models of human sentence processing. In M.W. Crocker, M. Pickering, & C. Clifton, Jr. (Eds.), *Architectures and mechanisms for language processing* (pp. 31–55). Cambridge, England: Cambridge University Press.
- Tanenhaus, M.K., Spivey-Knowlton, M.J., Eberhard, K.M., & Sedivy, J.E. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268, 632–634.

Note

1. Address correspondence to Fernanda Ferreira, 129 Psychology Research Building, Michigan State University, East Lansing, MI 48824-1117; e-mail: fernanda@eyelab.msu.edu.

References

- Barton, S.B., & Sanford, A.J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory & Cognition*, 21, 477–487.
- Christianson, K., Hollingworth, A., Halliwell, J., & Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology*, 42, 368–407.

- Duffy, S.A., Henderson, J.M., & Morris, R.K. (1989). Semantic facilitation of lexical access during sentence processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 791–801.
- Erickson, T.A., & Mattson, M.E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behavior*, 20, 540–552.
- Ferreira, F., & Clifton, C., Jr. (1986). The independence of syntactic processing. *Journal of Memory and Language*, 25, 348–368.
- Ferreira, F., & Stacey, J. (2000). *The misinterpretation of passive sentences*. Manuscript submitted for publication.
- Fillenbaum, S. (1974). Pragmatic normalization: Further results for some conjunctive and disjunctive sentences. *Journal of Experimental Psychology*, 102, 574–578.
- Frazier, L. (1978). *On comprehending sentences: Syntactic parsing strategies*. Unpublished doctoral dissertation, University of Connecticut, Storrs.
- MacDonald, M.C., Pearlmutter, N.J., & Seidenberg, M.S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, 101, 676–703.
- Sachs, J.S. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception & Psychophysics*, 2, 437–442.
- Traxler, M.J., Pickering, M.J., & Clifton, C., Jr. (1998). Adjunct attachment is not a form of ambiguity resolution. *Journal of Memory and Language*, 39, 558–592.
- Trueswell, J., Tanenhaus, M., & Garnsey, S. (1994). Semantic influences on parsing: Use of thematic role information in syntactic disambiguation. *Journal of Memory and Language*, 33, 285–318.