Sentence Processing

Chapter · November 2018

CITATIONS
2

READS
6,458

2 authors:

Derya Cokal
Queen Mary, University of London
23 PUBLICATIONS **66** CITATIONS
SEE PROFILE

Fernanda Ferreira
University of California, Davis
153 PUBLICATIONS **9,212** CITATIONS
SEE PROFILE

Title: Sentence Processing

* Fernanda Ferreira

University of South Carolina

Email: fernanda@sc.edu

Address: Institute for Brain and Mind, University of South Carolina, Columbia, SC 29208

Phone: 803-777-7224

Derya Çokal

University of South Carolina

Email: cokal@mailbox.sc.edu

Address: Same as above.

Phone: 803-404-0439

*Corresponding author.

Abstract

By convention, the term "sentence processing" refers to the subfield of psycholinguistics focusing on the interpretation of sentences. A range of information sources is used for successful sentence processing. Lexical and syntactic constraints are central for defining the structural alternatives, and information associated with the prosody of the sentence as well as the discourse and visual context in which the sentence occurs reinforces some interpretations and fleshes out the meaning of the sentence. Approaches to sentence processing differ on whether they assume serial versus parallel and modular versus interactive architectures; almost all assume incremental interpretation and even prediction of structure and lexical content. Current models include the garden-path model, constraint-based models, and approaches that allow the processing system to reduce and even distort the input (Good Enough Processing; Noisy Channel Models). An important trend is work designed to shed light on how sentence processing is implemented in the brain.

Key words: sentence processing models, syntax, ambiguity, modularity, online processing

The existence of a field called "sentence processing" attests to the implicit agreement among most psycholinguists that the sentence is a fundamental unit of language. In addition, by convention, the term "processing" in this context tends to refer to comprehension rather than production, and thus the topic of this chapter is people's interpretations of sentences. Our goal is to provide an overview of the findings, theories, and debates that will be discussed in more detail in the chapters in this volume making up the section on "Sentence Processing" (Chapters 53-61). The relevant issues include syntactic and semantic processing, the time-course of interpretation, and the role of other cognitive systems such as working memory in forming sentence interpretations. In this chapter we begin by examining the sources of information that are used during sentence processing. We then review the major theoretical controversies and debates in the field: the incremental nature of interpretation, serial versus parallel processing, and the extent of interaction among information sources during online processing. Then we will go over the major models of sentence processing, including syntax-based models, constraint-based models, the good-enough approach, and the very recent rational analysis approaches. We end with a few conclusions and speculations concerning future research directions.

1. Sources of Information for Sentence Processing

Since the 1980s, when psycholinguistics experienced a renaissance (Clifton, 1981) and returned to the question of how to relate formal and psychological approaches to language, the field of sentence processing has been associated with a commitment to the idea that syntactic information is critical to successful language comprehension. Not all theorists agree on the nature of those syntactic representations or the relative

importance of information sources that are nonsyntactic, but almost all assume that structure-building operations are essential for successful comprehension (Fodor, Bever, & Garrett, 1974; Frazier & Rayner, 1990). One key component is phrase-structure parsing, which refers to the process of identifying constituents and grouping them into a hierarchical structure. For example, in a sentence such as *While Mary bathed the baby played in the crib*, the parser must create a structural analysis which postulates the existence of a subordinate and a main clause; moreover, the main verb of the subordinate clause must be analyzed as intransitive and reflexive, and the subject of the main clause must be identified as *the baby*. With this analysis, the correct meaning can be derived, which is that Mary is bathing herself, and the baby is the agent of playing.

As the same example makes clear, one of the challenges to the parser is syntactic ambiguity: At various points in a sentence, a sequence of words can be given more than one grammatical analysis. In the example, the phrase *the baby* appears to be the object of bathed, but in fact turns out to be the subject of *played*. The result is a so-called "garden-path": The parser first builds an incorrect analysis, and reanalysis processes are triggered upon receipt of a constituent that cannot be incorporated into the existing structure. Because the parser obeys the rules of the grammar, including the rule mandating overt subjects, the parse will fail at *played*, and the sentence processing system must locate the alternative analysis on which *the baby* is a subject. How this happens is another point of divergence between competing sentence processing models, as will be discussed in Section 2.

An additional complication regarding the syntactic analysis of a sentence is that the grammar allows constituents to be moved from their canonical positions. One classic

example is the passive, in which the theme of an action is also the sentential subject, contrary to the general preference to align agency and subjecthood (Fillmore, 1968; Grimshaw, 1990; Jackendoff, 1990). Another type of moved constituent is wh-phrases; in English, as in many other languages, wh-phrases must be moved from their canonical position to a position at the beginning of the clause, leaving behind a trace or "gap". For example, in *Which man did the dog bite*?, the phrase *which man* receives its thematic role from *bite*. The job of the parser is to find the gap and relate it to the wh-phrase so that the sentence can receive a correct interpretation. This task is made difficult by two challenges: First, the gap is a phonetically null element in the string, and therefore the parser must identify the gap based on the application of a range of linguistic constraints. The second challenge concerns ambiguity: Because many verbs have multiple argument structures, the parser may end up postulating a gap incorrectly. The result are so-called "decoy gaps", as illustrated in *Who will the zombie eat with?*. The parser initially assumes that *who* was moved from the direct object position after *eat*, and then must reanalyze that structure when *with* is encountered.

Studies investigating the processing of filler-gap dependencies have found evidence for a filled-gap effect, which is closely related to decoy gaps. Consider the example *Which patient did the doctor expect the nurse to call?*. Most comprehenders will assume that *which patient* is the object of *expect*, but the NP *the nurse* occupies that position, which means that the parser must look further along for the correct gap (located after *call*). The existence of filled-gap effects has led researchers to postulate two parsing preferences for creating filler-gap dependencies: One is that the parser adopts an active or early filler strategy (Frazier, Clifton, & Randall, 1983; Frazier & Flores D'Arcais, 1989),

according to which a gap is postulated at the first syntactically permissible location. The second is that the parser makes use of verb argument structure information to guide the postulation of gaps. If a verb has a strong intransitive bias, then the parser is less likely to postulate a gap after it; if the verb is strongly transitive, then a postverbal gap will be more compelling.

As we have been discussing the importance of syntactic information for parsing, we have had numerous occasions to refer to lexical information as well. This is because lexical information is the fundamental bottom-up information source for sentence processing. Indeed, in lexicalist theories, syntactic information is attached to specific words so that when a word is retrieved, its associated structural possibilities become available as well (Joshi & Schabes, 1997; MacDonald, Pearlmutter, & Seidenberg, 1994). For example, retrieval of the verb *bathe* would bring up not just information associated with the syntactic category and meaning of that word, but also the word's syntactic dependents in the form of what are known as argument structures. An optionally transitive verb like *bathe* would have at least two argument structures, one specifying an agent and a patient, and the other specifying an agent and an obligatory reflexive null element. Nonlexicalist theories also assume a major role for this type of information; but in contrast with lexicalist theories, argument structures are used not to generate a parse, but rather to filter or reinforce a particular analysis, and to facilitate recovery from a garden-path.

Another type of lexical information that can be critical for parsing relates to semantic features such as number and animacy. Number information can affect how an ambiguous phrase is attached during online processing; for example, in a sentence such

as *While John and Mary kissed the baby slept*, the verb *kissed* is interpreted as intransitive because the plural subject triggers a reciprocal reading of *kissed*. A singular subject does not license this reciprocal interpretation (Ferreira & McClure, 1997; Patson & Ferreira, 2009). Similarly, animacy can help the parser avoid a garden-path, or help it recover more easily (Ferreira & Clifton, 1986; Trueswell, 1993). Specifically, if a subject is inanimate, then it is unlikely to be an agent, and that analysis in turn might lead the parser to adopt a less frequent passive or reduced relative parse (e.g., *The evidence examined by the lawyer*). These examples also show how word properties such as number and animacy interact with lexical argument structures, as those features can lead the parser to select one argument structure (e.g., a reciprocal one for a verb such as *kiss*) over another.

Next, let us consider the question of how prosodic information might influence sentence processing. The starting point for most studies published on this topic is that syntactic and prosodic structures are related, and in particular, major syntactic boundaries such as those separating clauses are usually marked by phrase-final lengthening and changes in pitch (Ferreira, 1993). Some clause-internal phrasal boundaries are also marked, although much less reliably (Allbritton, McKoon, & Ratcliff, 1996) - for example, in the sentence *John hit the thief with the bat*, the higher attachment of *with the bat*, which supports the instrument interpretation, is sometimes (but not always) associated with lengthening of *thief*. The logic of the research enterprise is to see whether prosodic "cues" can signal syntactic structure and help the parser to avoid going down a garden-path. One of the earliest studies to consider this question was conducted by Beach (1991), which demonstrated that metalinguistic judgments about sentence structure are

influenced by the availability of durational and pitch information linked to the final structures of the sentences. A couple of decades later, more sensitive online techniques including recording of event-related potentials (ERPs) and eyetracking have yielded a wealth of information about the comprehension of spoken sentences, and one of the ideas on which there is now a general consensus is that prosody indeed does influence the earliest stages of parsing (e.g., Nakamura, Arai, & Mazuka, 2012).

Another potentially influential source of information for sentence processing is context, both discourse and visual. An early analysis of the role of discourse context is known as Referential Theory (Crain & Steedman, 1985). They observed that many of the sentence forms identified as syntactically dispreferred by the two-stage model are also presuppositionally more complex. For example, the sentence *John hit the thief with the bat* allows for two interpretations: the *with*-phrase may be interpreted as an instrument or a modifier; the latter interpretation requires a more complex structure (on some theories of syntax). The "confound" here is that the more complex structure also involves modification whereas the simpler analysis does not. Moreover, a modified phrase such as *the thief with the bat* presupposes the existence of more than one thief, and thus the difficulty of the more complex structure might not be due to its syntax but rather to the lack of a context to motivate the modified phrase. Crain and Steedman predicted that sentences processed in presuppositionally appropriate contexts would be easy to process, a prediction that Ferreira and Clifton (1986) examined using eye movement monitoring in reading. Their data were consistent with the idea that context did not affect initial parsing decisions: Supportive contexts led to shorter global reading times and more accurate question-answering behavior, but early measures of processing revealed that processing

times were longer for structurally complex sentences compared to their structurally simpler counterparts.

The potential role of visual context became a topic of intense interest in the 1990s with the emergence of the Visual World Paradigm (VWP) for studying sentence processing. The idea behind the paradigm is simple: From reading studies, it was known that fixations are closely tied to attention and processing (Rayner, 1977). The VWP extends this logic to spoken language processing by pairing spoken utterances with simple displays containing mentioned and unmentioned objects. The "linking hypothesis" (Tanenhaus, Magnuson, Dahan, & Chambers, 2000) is that as a word is heard, its representation in memory becomes activated, which triggers eye movements towards the named object as well as objects semantically and even phonologically associated with it (Huettig & McQueen, 2007). The widespread adoption of the VWP occurred in part because the idea of multimodal processing was also catching on, with many cognitive scientists wanting to understand the way different cognitive systems might work together - in this case, the auditory language processing system and the visuo-attention system associated with object recognition (Henderson & Ferreira, 2004; Jackendoff, 1996). There was also a growing interest in auditory language processing generally, and in the investigation of how prosodic information might be used during comprehension, as discussed earlier. By now, hundreds of studies have been reported making use of it in one way or another (Ferreira, Foucart, & Engelhardt, 2013; Huettig, Olivers, & Hartsuiker, 2011; Huettig, Rommers, & Meyer, 2011).

The report that triggered the widespread use of the VWP is Tanenhaus, Spivey-Knowlton, Eberhard and Sedivy (1995; Spivey, Tanenhaus, Eberhard, & Sedivy, 2002).

This study adapted the ideas of Crain and Steedman (1985) concerning presuppositional support to the domain of visual contexts and spoken sentences that could be evaluated against them. To illustrate, consider the sentence *Put the apple on the towel in the box*. At the point at which the listener hears *on the towel*, two interpretations are possible: Either *on the towel* is the location to which the apple should be moved, or it is a modifier of *apple*. The phrase *into the box* forces the latter interpretation because it is unambiguously a location. Referential Theory specifies that speakers should provide modifiers only when modification is necessary to establish reference. It follows that if two apples are present in the visual world and one of them is supposed to be moved, then right from the earliest stages of processing, the phrase *on the towel* will be taken to be a modifier, because the modifier allows a unique apple to be identified. The listener faced with this visual world containing two referents should therefore immediately interpret the phrase as a modifier and avoid being garden-pathed (Farmer, Cargill, & Spivey, 2007; Novick, Thompson-Schill, & Trueswell, 2008; Spivey et al., 2002; Tanenhaus et al., 1995). Recently, however, the interpretation of these findings has been challenged. Ferreira et al. (2013) conducted three experiments manipulating properties of the utterances and the visual worlds. They concluded that listeners engage in a fairly atypical mode of processing in VWP experiments with simple visual worlds and utterances that are highly similar to each other over all experimental trials: Rather than processing utterances normally, they instead form a skeleton, underspecified representation of what they are likely to hear based on the content of the display, and then evaluate that prediction against the utterance itself. These issues concerning the use of the VWP require additional investigation.

In summary, a range of sources of information is used for successful sentence processing. Lexical and syntactic constraints are central for defining the structural alternatives considered by the language processing system, and information associated with the prosody of the sentence as well as the discourse and visual context in which the sentence occurs helps to reinforce some interpretations and flesh out the full meaning of the sentence. In the next section, we consider some of the theoretical controversies concerning the architecture of the language system and the way these sources of information are coordinated. This discussion will set the stage for our discussion of theoretical models of sentence processing.

2. Theoretical Controversies

In this section, we will review three issues that help distinguish among competing models of sentence processing: (1) incremental interpretation, (2) serial versus parallel processing, and (3) interactivity versus modularity.

Incremental interpretation refers to whether the sentence processing system builds up the meaning of a sentence word-by-word, as the input unfolds, or whether the system either falls behind or gets ahead of the input. Falling behind the input would indicate delays in interpretation; getting ahead would indicate anticipation or prediction. Essentially all current models of processing assume that interpretations are built up incrementally, and in particular, that there are no delays in incorporating new words into the ongoing representation of sentence meaning. In addition, there is some evidence that comprehenders engage in prediction (Levy, 2008; Rayner, Li, Juhasz, & Yan, 2005; Van Berkum, Brown, Zwiserlood, Kooijman, & Hagoort, 2005). The classic demonstration of prediction comes from Altmann and Kamide (1999), who used the VWP and

semantically constrained sentences such as *The boy will eat the cake*. They observed that listeners made anticipatory eye movements to a depicted cake prior to hearing the word *cake*, indicating that they predicted that continuation. In the structural domain, Staub and Clifton (2006) found that when readers processed a clause beginning with the word *either*, they predicted an upcoming *or*-clause based on the syntactic constraint that the latter must follow the former. These and other studies have been taken as evidence that the sentence processing system is not just incremental, but actually predictive, anticipating structure and even specific lexical content.

At the same time, there is some evidence that additional processing takes place at major syntactic boundaries. So-called "end of sentence wrap up" refers to the finding that reading times at the ends of clauses and sentences are longer than in other sentential positions (Aaronson & Scarborough, 1976; Just & Carpenter, 1980; Rayner, Kambe, & Duffy, 2000; Rayner, Sereno, Morris, Schmauder, & Clifton, 1989). Wrap up effects indicate that some elements of meaning are computed over a more global domain. In addition, clause boundaries might be the locations where the comprehension system evaluates the entire structure to ensure that all relevant constraints are satisfied - for example, to check that a verb has all its obligatory arguments. Evidence for underspecified representations also suggests some tendency on the part of the processing system to delay interpretations (for an excellent summary, see Frisson, 2009). Words with multiple senses (e.g., *book* as an object versus its content) seem to be processed by initially activating an underspecified meaning, and then filling out the semantics once contextually disambiguating information becomes available. Some syntactic ambiguities may also be handled in a similar manner--for example, comprehenders leave open the

interpretation of ambiguous relative clauses (*the servant of the actress who was on the balcony*), making a specific attachment decision only once it is necessary to do so (Swets, Clifton, & Ferreira, 2008). Pronouns are also often not assigned specific antecedents (McKoon, Greene, & Ratcliff, 1993).

The second theoretical issue on which theories of sentence processing differ is serial versus parallel processing, which typically refers to assumptions about whether the system considers only one interpretation at a time or multiple interpretations. For example, consider *The defendant examined by the lawyer turned out to be unreliable* (Ferreira & Clifton, 1986). The sequence *the defendant examined* could mean that the defendant examined something or that the defendant is the thing being examined (the ultimately correct analysis). The issue is whether only one of these interpretations is built and evaluated at any one time, or whether all the interpretations are simultaneously activated and assessed. On the serial view, first the system considers one analysis -- on most theories, the one which assumes that the defendant is the agent of examining, given that this analysis is syntactically simpler and more frequent -- and then reanalyzes it should a revision signal be encountered. The sentence processing system then goes into "reanalysis mode", attempting to adjust the syntactic structure that has been built so as to create a grammatical analysis (Ferreira & Henderson, 1991; Fodor & Ferreira, 1998; Fodor & Inoue, 1994). Ease of reanalysis depends on the extent to which the sentence processing system can find lexical and grammatical information that motivates an alternative structure.

The parallel view assumes that the sentence processing system activates all grammatically licensed analyses simultaneously. Considering our example, both the

incorrect and the ultimately correct interpretations of *the defendant examined* would be available in parallel, initially weighted by their frequency. The agent analysis of *defendant* is more frequent and therefore at first it will be stronger than the ultimately correct analysis. But when the word *by* is encountered, the sentence processing system must shift to the other activated interpretation. Ease of reanalysis depends on the relative activation levels of the two interpretations. If the ultimately correct interpretation is infrequent, then it will be difficult to retrieve and reanalysis might even fail. If the right interpretation has some strength based on the extent to which it conforms to a wide range of linguistic and nonlinguistic constraints, then reanalysis will be easier, and so will overall comprehension of the sentence.

A careful reader might have noticed subtle differences in the terminology used in our discussion of serial versus parallel processing. For the former, interpretations are typically described as being "built", whereas for the latter, they are often referred to as being "activated" or "retrieved". These different terms reflect fundamentally different ideas about how interpretations are stored in memory and accessed during sentence processing. The serial view tends to assume that syntactic rules are stored in memory and then used online to create a structural representation bit by bit. Reanalysis processes are a matter of editing the structure. The parallel view tends to assume that structures are stored in chunks, typically corresponding to an argument-taking word such as a verb and its arguments. Online processing involves not so much building a structure as much as activating one. These issues will be raised again when we consider models of sentence processing.

Finally, almost since the earliest days of psycholinguistics, debate has centered around the issue of whether the system considers only linguistic (and possibly even only syntactic) information when parsing a sentence, versus a system that considers all potentially relevant sources of information. Modular models assume sentence structures are assigned to words at least initially without any consideration of whether the structure will map on to a sentence interpretation that makes sense given prior knowledge or given the contents of the immediate linguistic, visual, or social context. For example, the sentence processing system would be garden-pathed not just by *the defendant examined by the lawyer*, but also by *the evidence examined by the lawyer*, even though *evidence* is inanimate and therefore cannot engage in an act of examination. In contrast, interactive models assume the immediate use of all relevant constraints. At this stage, there is widespread belief in the field that the preponderance of evidence supports interactive models, although it is possible to argue that this conclusion goes somewhat beyond the evidence (Ferreira & Nye, in progress).

3. Classes of Models of Sentence Processing

We begin with the so-called "two-stage model" or "garden-path" model, first developed by Lyn Frazier (Ferreira & Clifton, 1986; Frazier & Fodor, 1978; Rayner, Carlson, & Frazier, 1983). The model assumes that a single parse is constructed for any sentence based on the operation of Minimal Attachment, which constrains the parser to construct no potentially unnecessary syntactic nodes, and Late Closure, which causes the parser to attach new linguistic input to the current constituent. In addition, the model assumes that the only information that the parser has access to when building a syntactic structure is its database of phrase structure rules, and therefore the parser cannot consult

information associated with lexical items. For example, in the sequence *Mary knew Bill*, the noun phrase *Bill* would be assigned the role of direct object because that analysis is simpler than the alternative subject-of-complement-clause analysis, and the information that *know* takes sentence complements more frequently than direct objects could not be used to inform the initial parse.

The two-stage model has evolved over the last three decades or so to take into account changes in linguistic theory and significant findings in psycholinguistics. One important addition is the notion of "Construal" (Frazier & Clifton Jr, 1997; Frisson & Pickering, 2001), which allows some constituents to be merely associated with a specific thematic domain in a sentence rather than definitively attached into the structure. Evidence for Construal comes from the finding that readers process sentences with ambiguous relative clauses more quickly than those that have a unique attachment (e.g., *the servants of the actress who was on the balcony*), unless the sentence is followed by a question that forces the reader to provide a specific interpretation; in that case, readers take longer to read the ambiguous versions, presumably because they are trying to choose between the attachment options. Another important revision of the two-stage model is that now, prosody plays an essential role in determining how parsing proceeds from the earliest stages of processing (Millotte, Wales, & Christophe, 2007; Nakamura et al., 2012; Price, Ostendorf, Shattuck-Hufnagel, & Fong, 1991). Pitch and durational information associated with different kinds of prosodic and intonational phrasing are used to constrain the parser's syntactic analyses and assist in the construction of semantic meanings such as focus and presupposition. Nonetheless, the essential features of the two-stage model remain: It assumes that (a) information is used incrementally to build up

an interpretation, (b) different possible interpretations are built and evaluated serially, rather than in parallel, and (c) only certain kinds of information can be used during the initial stages of sentence processing - in particular, information stated in the syntactic and prosodic vocabulary of the sentence processing module.

The two-stage model was soon challenged by researchers in sentence processing who were strongly influenced by the connectionist architectures popular in the 1980s and 1990s (Rumelhart & McClelland, 1985, 1986; Seidenberg & McClelland, 1989). Thee architectures contrast with the assumptions of the two-stage model in two defining ways: First, in connectionist systems, alternative possibilities are activated and evaluated in parallel, and second, any relevant source of information can be used to modulate the activation levels and allow one possible analysis to win at the expense of the others (e.g., Macdonald et al., 1994). Applying these ideas to sentence processing, the connectionist alternative assumed the following principles. First, rather than analyses being built with the help of grammatical rules, a great deal of the burden of syntactic representation is put into the lexicon. Adapting ideas that were then timely in linguistic theory (e.g., Pesetsky, 1995), lexical representations were assumed to activate not just words and word meanings, but also syntactic frames. On this view, syntactic rules are redundant because almost all the necessary information is already stated in the lexicon. Thus, with syntactic structures being stored rather than built, it is easy to imagine an architecture in which all possible analyses are considered in parallel, weighted by their frequency of use. Lexical, contextual, and pragmatic constraints can be used to further modulate the activation levels. On this approach the sentence processing system is (a) incremental, but (b) different possible interpretations are activated in parallel. In addition, (c) any potential

source of information can be used at any stage of sentence processing, making the system interactive rather than modular.

In the last fifteen years or so, an alternative to both the two-stage and the connectionist models has emerged. There are many variants with important distinctions among them, but what they share is the idea that comprehenders sometimes end up with an interpretation that differs from the actual input received - the interpretation is either simpler (Construal), somewhat distorted (Late Assignment of Syntax Theory; Good-Enough Processing), or outright inconsistent (Noisy Channel Approaches) with the sentence's true content. Let us begin with the models that assume representations which reduce the input in some way. One implementation is to allow representations to be underspecified (Sanford & Sturt, 2002). Consider Construal: As mentioned earlier, this model assumes that syntactic structures are not always fully connected, and adjunct phrases in particular (e.g., relative clauses, modifying prepositional phrases) may instead simply get associated with a certain processing domain, "floating" until disambiguating information arrives. The parser thus remains uncommitted (Pickering et al., 2006) concerning the attachment of the relative clause and the interpretation that would follow from any particular attachment (see Frisson & Pickering, 2001; Sanford & Graesser, 2006; Sturt, Sanford, Stewart, & Dawydiak, 2004). Other studies support the idea of underspecified representations for global syntactic structures (Tyler & Warren, 1987), semantic information (Frazier & Rayner, 1990), and coercion structures (Pickering et al., 2006).

More radical variants of shallow processing models allow the comprehension system to generate an interpretation that is even more discrepant from the input.

Researchers in the field of text processing have shown that readers are sometimes remarkably insensitive to contradictions in text (Otero & Kintsch, 1992), and also often fail to update their interpretations when later information undermines a fact stated earlier (Albrecht & O'Brien, 1993). These ideas from text processing were exported to the sentence processing literature in a series of experiments showing that people do not seem to fully recover from garden-paths (Christianson, Hollingworth, Halliwell, & Ferreira, 2001). Participants read sentences such as *While the woman bathed the baby played in the crib* and then they answered a question such as *Did the woman bathe the baby?*. The surprising finding was that most people answered "yes", even though the meaning of the reflexive verb *bathe* requires that the object be interpreted as coreferential with the subject (see also Slattery, Sturt, Christianson, Yoshida, & Ferreira, 2013). It appears that comprehenders are not entirely up to the task of syntactic reanalysis, and sometimes fail to revise either all pieces of the syntactic structure or all elements of the semantic consequences of the initial, incorrect parse. And the more semantically compelling the original misinterpretation, the more likely people are to want to retain it.

Townsend and Bever's (2001) model implements an architecture similar to what has been suggested for decision-making (Gigerenzer, 2004; Kahneman, 2003), where researchers sometimes distinguish between a so-called "System 1" and "System 2" (or "Type 1" and "Type 2") reasoning. System 1 is fast, automatic, and operates via the application of simple heuristics. System 2, on the other hand, is slow and attention-demanding, and consults a wide range of beliefs - essentially anything the organism knows and has stored in memory. In Townsend and Bever's (2001) Late Assignment of Syntax Theory (LAST), sentences are essentially processed twice: First, heuristics are

accessed which yield a quick meaning, and then syntactic computations are performed on the same word string to yield a fully connected, syntactic analysis. The second process ensures that the meaning that is obtained for a sentence is consistent with its actual form. Townsend and Bever also assume that the first stage is nonmodular and the second modular; this is to account for the use of semantics in the first stage, and the use of essentially only syntactic constraints in the second.

Two models similar in spirit to LAST but which assume a modular architecture for the first stage have been proposed by Ferreira (2003) and by Garrett (2000). The Ferreira model assumes that the first stage consults just two heuristics - a version of the "NVN" strategy, in which people assume an agent-patient mapping of semantic roles to syntactic positions, and an animacy heuristic, in which animate entities are biased towards subjecthood. The 2003 Ferreira model explains comprehenders' tendency to misinterpret passive sentences, particularly when they express an implausible event with reversible semantic roles, as in *the dog was bitten by the man*. The application of heuristics in the first stage yields the dog-bit-man interpretation; a proper syntactic parse will deliver the opposite, correct interpretation, but the model assumes that it is fragile and susceptible to interference. Garrett (2000) offers a more explicitly analysis-by-synthesis model which incorporates the production system to generate what are generally thought of as top-down effects. A first pass, bottom-up process uses syntactic information to create a simple parse that in turn allows for a rudimentary interpretation; then the language production system takes over and uses that representation to generate the detailed syntactic structure that would support the initial parse and interpretation.

Finally, a family of models has recently been proposed that assume people engage in rational behavior over a noisy communication channel. The channel is noisy both because listeners sometimes mishear or misread due to processing error or environmental contamination, and because speakers sometimes make mistakes when they communicate. Thus, a rational comprehender whose goal is to recover the intention behind the utterance will normalize the input according to Bayesian priors. A body of evidence from research using Event-Related Potentials (ERPs) helped to motivate these ideas (Kim & Osterhout, 2005; Van Herten, Kolk, & Chwilla, 2005). In these experiments, it is reported that subjects who encounter a sentence such as *The fox that hunted the poachers stalked through the woods* experience a P600 rather than an N400 upon encountering the semantically anomalous word, even though an N400 might be expected given that it is presumed to reflect problems related to meaning. There is still not a great deal of consensus on what triggers P600s, but an idea that has been gaining traction is that it reflects a need to engage in some type of structural reanalysis or revision. The conclusion, then, is that when a person encounters a sentence that seems to say that the fox hunted the poachers, they "fix" it so it makes sense, resulting in a P600. Other models have taken this idea and developed it further (Gibson, Bergen, & Piantadosi, 2013; Levy, 2011; Levy, Bicknell, Slattery, & Rayner, 2009). These models are generally interactive, as the information that is accessed to establish the priors can range from biases related to structural forms all the way to beliefs concerning speaker characteristics (Van Berkum, Van den Brink, Tesink, Kos, & Hagoort, 2008). However, these noisy channel models have not yet been rigorously tested using a methodology that allows early processes to be distinguished from later ones. For example, it remains possible that comprehenders create

a simple parse in a manner compatible with modularity and then consult information outside the module to revise that interpretation, right down to actually normalizing the input. Indeed, models designed to explain the comprehension of sentences containing self-repairs and other disfluencies (*turn left uh right at the light*) assume mechanisms that allow input to be deleted so that the speaker's intended meaning can be recovered (Ferreira, Lau, & Bailey, 2004).

4. Conclusions

The field of sentence processing has changed significantly since the 1980s. Current models emphasize more detailed, context-specific information such as speaker, and there is a great deal of interest in mechanisms that allow the input to be rationally evaluated and corrected. Future work will continue to make use of behavioral techniques as well as methods from neuroscience to expand our understanding of these topics. The critical next stage is to determine how the processes assumed in models of sentence processing are actually implemented in the human brain. Our view is that the field is well-positioned for this challenge given the sophistication of extant sentence processing models.

References

Aaronson, D. & Scarborough, HS. (1976). Performance theories for sentence coding: Some quantitative evidence. *Journal of Experimental Psychology: Human Perception and Performance*, *2*(1), 56–70.

Albrecht, J. E., & O'Brien, E. J. (1993). Updating a mental model: Maintaining both local and global coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*(5), 1061-1070.

Allbritton, D. W., McKoon, G., & Ratcliff, R. (1996). Reliability of prosodic cues for resolving syntactic ambiguity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*(3), 714-735.

Altmann, G. T. M., & Kamide, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition,73*, 247-264.

Beach, C. M. (1991). The interpretation of prosodic patterns at points of syntactic structure ambiguity: Evidence for cue trading relations. *Journal of Memory and Language, 30*(6), 644-663.

Christianson, K., Hollingworth, A., Halliwell, J. F., & Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology, 42*(4), 368-407.

Clifton, C. Jr. (1981). Psycholinguistic renaissance? *Contemporary Psychology*, *26*, 919–921.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological parser. In D. Dowty, L. Karttunen, & A. Zwicky (Eds.), *Natural Language Parsing: Psychological, Computational, and*

Theoretical Perspectives (pp. 320-358). Cambridge, UK: Cambridge University Press.

Farmer, T. A., Cargill, S. A., Hindy, N. C., Dale, R., & Spivey, M. J. (2007). Tracking the continuity of language comprehension: Computer mouse trajectories suggest parallel syntactic processing. *Cognitive Science, 31*(5), 889-909.

Ferreira, F. (1993). The creation of prosody during sentence production. *Psychological Review*, *100*, 233-253.

Ferreira, F. (2003). The misinterpretation of noncanonical sentences. *Cognitive Psychology, 47*(2), 14-203.

Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language, 25*(3), 348-368.

Ferreira, F., Foucart, A., & Engelhardt, P. E. (2013). Language processing in the visual world: Effects of preview, visual complexity, and prediction. *Journal of Memory and Language, 69*(3), 165-182.

Ferreira, F., & Henderson, J.M. (1991). Recovery from misanalyses of garden-path sentences. *Journal of Memory and Language*, *30*, 725-745.

Ferreira, F., Lau, E. F., & Bailey, K. G. (2004). Disfluencies, language comprehension, and tree adjoining grammars. *Cognitive Science, 28*(5), 721-749.

Ferreira, F., & McClure, K. (1997). Parsing of garden-path sentences with reciprocal verbs. *Language and Cognitive Processes*, *12*, 273-306.

Ferreira, F., & Nye, J. (in progress). The Modularity of Sentence Processing Reconsidered. In R.G. De Almeida, & L. Gleitman (Eds.), *The Modularity of Language: A Psycholinguistic Festschrift for Jerry Fodor*. Cambridge: MIT.

Fillmore, C. J. (1968). The creation of prosody during sentence production. *Psychological Review*, *100*, 233-253.

Fodor, J. A., Bever T.G., & Garrett, M. (1974). *The psychology of Language: An introduction to Psycholinguistics and Generative Grammar*. New York: MCRAW-Hill.

Fodor, J. D., & Ferreira, F. (1998). *Reanalysis in sentence processing.* Dordrecht, The Netherlands: Kluwer.

Fodor, J.D. & Inoue, A. (1994). The diagnosis and cure of garden paths. *Journal of Psycholinguistic Research*, *23*, 407-434.

Frazier, L., Clifton, C. & Randall, J. (1983). Filling Gaps: Decision Principles and Structure in Sentence Comprehension. *Cognition*, *13*, 187-222.

Frazier, L., & Clifton, C. Jr. (1989). Successive cyclicity in the grammar and the parser. *Language and Cognitive Processes*, *4*, 93-126.

Frazier, L., & Clifton, Jr. C. (1997). Construal: Overview, motivation, and some new evidence. *Journal of Psycholinguistic Research, 26*(3), 277-295.

Frazier, L., & Flores D'Arcais, G. B. (1989). Filler driven parsing: A study of filling gap in Dutch. *Journal of Memory and Language*, *28*, 331- 444.

Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition, 6*(4), 291-325.

Frazier, L., & Rayner, K. (1990). Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language, 29*(2), 181-200.

Frisson, S. (2009). Semantic underspecification in language processing. *Language and Linguistics Compass, 3*(1), 111-127.

Frisson, S., & Pickering, M. J. (2001). Obtaining a figurative interpretation of a word: Support for underspecification. *Metaphor and Symbol, 16*(3-4), 149-171.

Garrett, M. (2000). Remarks on the architecture of language processing systems. In Y. Grodzinsky, & L. Shapiro (Eds.), *Language and the Brain: Representation and Processing* (pp. 31-69). San Diego, CA: Academic Press.

Gibson, E., Bergen, L., & Piantadosi, S. T. (2013). Rational integration of noisy evidence and prior semantic expectations in sentence interpretation. *Proceedings of the National Academy of Sciences*, *110(20)*, pp. 8051-8056.

Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. In D.J. Koehler, & N. Harvey (Eds.), *Handbook of Judgment and Decision Making* (pp. 62-88). Oxford, UK: Blackwell.

Grimshaw, J. (1990). *Argument structure*. Cambridge, MA: MIT Press.

Henderson, J. M., & Ferreira, F. (2004). Scene perception for psycholinguists. In J. M. Henderson, & F. Ferreira (Eds.), *The interface of language, vision, and action: Eye movements and the visual world* (pp. 1-58). New York: Psychology Press.

Huettig, F., & McQueen, J. M. (2007). The tug of war between phonological, semantic and shape information in language-mediated visual search. *Journal of Memory and Language, 57*(4), 460-482.

Huettig, F., Olivers, C. N., & Hartsuiker, R. J. (2011). Looking, language, and memory: Bridging research from the visual world and visual search paradigms. *Acta Psychologica, 137*(2), 138-150.

Huettig, F., Rommers, J., & Meyer, A. S. (2011). Using the visual world paradigm to study language processing: A review and critical evaluation. *Acta psychologica, 137*(2), 151-171.

Jackendoff, R. (1996). The architecture of the linguistic-spatial interface. In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and space. Language, speech, and communication* (pp. 1-30). Cambridge: MIT Press.

Jackendoff, R. S. (1990). *Semantic structures*. Cambridge, MA: MIT Press.

Joshi, A. K. & Schabes, Y. (1997). Tree-adjoining grammars. In G. Rozenberg, & A. Salomaa (Eds.), *Handbook of Formal Languages* (pp. 69–123). Berlin: Springer.

Just, M. A. & Carpenter, P.A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, *87*(4), 329-354.

Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *The American Economic Review, 93*(5), 1449-1475.

Kim, A., & Osterhout, L. (2005). The independence of combinatory semantic processing: Evidence from event-related potentials. *Journal of Memory and Language, 52*(2), 205-225.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, *106*(3), 1126-1177.

Levy, R. (2011). Probabilistic Linguistic Expectations, Uncertain Input, and Implications. *Studies of Psychology and Behavior, 9*(1), 52-63.

Levy, R., Bicknell, K., Slattery, T., & Rayner, K. (2009). Eye movement evidence that readers maintain and act on uncertainty about past linguistic input. *Proceedings of the National Academy of Sciences*, *106* (50), 21086-21090.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review, 101*(4), 676-703.

McKoon, G., Greene, S., & Ratcliff, R. (1993). Discourse models, pronoun resolution, and the implicit causality of verbs. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 1-13.

Millotte, S., Wales, R., & Christophe, A. (2007). Phrasal prosody disambiguates syntax. *Language and Cognitive Processes, 22*(6), 898-909.

Nakamura, C., Arai, M., & Mazuka, R. (2012). Immediate use of prosody and context in predicting a syntactic structure. *Cognition, 125*(2), 317-325.

Novick, J. M., Thompson-Schill, S. L., & Trueswell, J. C. (2008). Putting lexical constraints in context into the visual-world paradigm. *Cognition, 107*(3), 850-903.

Otero, J., & Kintsch, W. (1992). Failures to detect contradictions in a text: What readers believe versus what they read. *Psychological Science, 3*(4), 229-235.

Patson, N. D., & Ferreira, F. (2009). Conceptual plural information is used to guide early parsing decisions: Evidence from garden-path sentences with reciprocal verbs. *Journal of Memory and Language*, *60*, 464-486.

Pesetsky, D. (1995). *Zero Syntax. Experiencers and Cascades*. Cambridge, Mass.: MIT Press.

Price, P. J., Ostendorf, M., Shattuck-Hufnagel, S., & Fong, C. (1991). The use of prosody in syntactic disambiguation. *The Journal of the Acoustical Society of America, 90*, 2956-2970.

Rayner, K. (1977). Visual attention in reading: Eye movements reflect cognitive processes. *Memory & Cognition, 5*(4), 443-448.

Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Verbal Learning and Verbal Behavior, 22*(3), 358-374.

Rayner, K., Kambe, G, & Duffy, SA. (2000). The effect of clause wrap-up on eye movements during reading. *The Quarterly Journal of Experimental Psychology*, *53*(4), 1061–1080.

Rayner, K., Li, X., Juhasz, B. & Yan, G. (2005). The effect of word predictability on the eye movements of Chinese readers. *Psychonomic Bulletin and Review*, *12*(6), 1089-1093.

Rayner, K, Sereno, S.C., Morris, R.K., Schmauder, A.R., & Clifton, C. (1989). Eye movements and on-line language comprehension processes. *Language and Cognitive Processes*, *4*, 21–49.

Rumelhart, D.E. & McClelland, J.L. (1985). Level's indeed! A response to Broadbent. *Journal of Experimental Psychology*: *General*, *114*(2), 193-197

Rumelhart, D.E. & McClelland, J.L. (1986). PDP Models and general issues in cognitive science. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel Distributed Processing* (pp.111-146*)*. Cambridge: M.I.T. Press.

Sanford, A. J., & Graesser, A. C. (2006). Shallow processing and underspecification. *Discourse Processes, 42*(2), 99-108.

Sanford, A. J., & Sturt, P. (2002). Depth of processing in language comprehension: Not noticing the evidence. *Trends in Cognitive Sciences, 6*(9), 382-386.

Slattery, T. J., Sturt, P., Christianson, K., Yoshida, M., & Ferreira, F. (2013). Lingering misinterpretations of garden path sentences arise from competing syntactic representations. *Journal of Memory and Language, 69*(2), 104-120

Spivey, M. J., Tanenhaus, M. K., Eberhard, K. M., & Sedivy, J. C. (2002). Eye movements and spoken language comprehension: Effects of visual context on syntactic ambiguity resolution. *Cognitive Psychology, 45*(4), 447-481.

Staub, A., & Clifton, C. (2006). Syntactic prediction in language comprehension: Evidence from either… or. *Journal of Experimental Psychology*: *Learning, Memory, and Cognition*, *32*, 425-436.

Sturt, P., Sanford, A. J., Stewart, A., & Dawydiak, E. (2004). Linguistic focus and good-enough representations: An application of the change-detection paradigm. *Psychonomic Bulletin & Review, 11*(5), 882-888.

Swets, B., Desmet, T., Clifton, C., & Ferreira, F. (2008). Underspecification of syntactic ambiguities: Evidence from self-paced reading. *Memory & Cognition, 36*(1), 201-216

Tanenhaus, M. K., Magnuson, J. S., Dahan, D., & Chambers, C. (2000). Eye movements and lexical access in spoken-language comprehension: Evaluating a linking hypothesis between fixations and linguistic processing. *Journal of Psycholinguistic Research, 29*(6), 557-580.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science, 268*(5217), 1632-1634.

Townsend, D. J., & Bever, T. G. (2001). *Sentence comprehension: The integration of habits and rules* (Vol. 1950). Cambridge, MA: MIT Press.

Traxler, M. J., Pickering, M. J., & Clifton, C. (1998). Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of Memory and Language, 39*(4), 558-592.

Trueswell, J. C., Tanenhaus, M. K., & Kello, C. (1993). Verb-specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. Journal of Experimental Psychology: Learning, Memory and Cognition, *19*(3), 528-553.

Tyler, L. K., & Warren, P. (1987). Local and global structure in spoken language comprehension. *Journal of Memory and Language, 26*(6), 638-657.

Van Berkum, J. A., Brown, C.M., Zwitserlood, P., Kooijman, & Hagoort, P. (2005). Anticipating upcoming words in discourse: Evidence from erps and reading times: *Journal of Experimental Psychology*: *Learning, Memory and Cognition*, *31*(3), 443-467.

Van Berkum, J., Van den Brink, D., Tesink, C., Kos, M., & Hagoort, P. (2008). The neural integration of speaker and message. *Journal of Cognitive Neuroscience, 20*(4), 580-591.

Van Herten, M., Kolk, H. H., & Chwilla, D. J. (2005). An ERP study of P600 effects elicited by semantic anomalies. *Cognitive Brain Research, 22*(2), 241-255.