

## CHAPTER 17

# Syntactic parsing

Roger P. G. van Gompel and Martin J. Pickering

### 17.1 Introduction

A crucial part of understanding a sentence is to construct its syntactic structure. Without this, it would be very difficult for language users to determine that sentences with different word orders, such as *The man sees the woman* and *The woman sees the man* have different interpretations, or explain why sentences such as *The hunter killed the poacher with the rifle* have two possible interpretations. The processes involved in constructing syntactic structures during language comprehension are commonly referred to as *parsing* or *syntactic processing*.

Sentence processing research has shown that parsing is largely incremental, i.e. language comprehenders incorporate each word into the preceding syntactic structure as they encounter it; they do not delay syntactic structure building until, for instance, the end of the sentence or phrase (e.g. Marslen-Wilson, 1973; 1975). Evidence for incrementality comes from numerous studies that show that language comprehenders experience difficulty with temporarily ambiguous sentences well before the end of the sentence. For example, many experiments have shown that readers slow down in the region *by the lawyer* in (1) (e.g. Clifton et al., 2003; Ferreira and Clifton, 1986; Rayner et al., 1983; Trueswell et al., 1994).

1. The defendant examined by the lawyer turned out to be unreliable.

This sentence is temporarily ambiguous at *examined*, because this verb could be part of a *reduced relative structure*, in which it is a past participle, or part of a *main clause structure*, in which it is a past tense verb. The finding that people experience difficulty at *by the lawyer* is

normally interpreted as evidence that people initially favor the main clause analysis, and experience difficulty because *by the lawyer* rules out this analysis. The difficulty that people experience in such a case is often referred to as a “garden-path effect” (Bever, 1970).

There is considerable controversy about when people use different sources of information during sentence processing. Most controversially, do they immediately use all relevant sources of information, or are some sources of information delayed relative to others? Sentence processing theories can roughly be divided into *interactive* accounts, in which all relevant information can be used immediately, and *modular* accounts, in which some information can be used immediately but some information cannot.

### 17.2 Modular models

Modular models assume that the mind consists of modules that perform very specific processes (e.g. Fodor, 1983). These processes are informationally encapsulated: they use only information represented within this module. In sentence processing research, this has led to the question of whether syntactic processes are separable from other processes such as semantic and discourse processing.

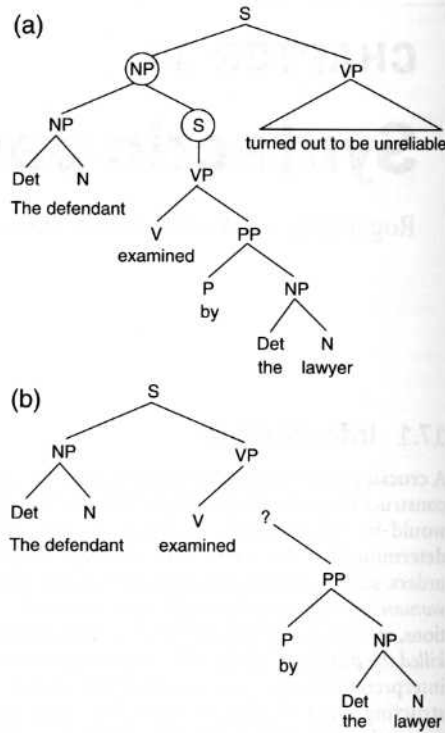
Modularity and informational encapsulation in sentence processing have usually been investigated using (temporarily) ambiguous sentences. In unambiguous sentences, syntactic information provides an extremely strong structural cue, so even according to models that are not modular, non-syntactic factors are unlikely to have much of an effect on processing. Therefore, most

sentence processing studies have investigated how different sources of information are employed during the processing of globally and temporarily ambiguous sentences, where syntactic cues allow multiple interpretations.

According to by far the most influential modular account of syntactic ambiguity resolution, the garden-path model (e.g. Frazier, 1987), the sentence processor initially employs only information about the syntactic structure of the sentence to adopt a single analysis in (temporarily) ambiguous sentences. Other, non-structural sources of information such as semantics, context, and frequency of the structures are employed during later stages of processing (e.g. Frazier, 1987; Rayner et al., 1983). When the initial analysis is inconsistent with information that becomes available later, the processor has to reanalyze, and processing difficulty ensues.

The garden-path model stipulates that the principles of minimal attachment and of late closure determine people's initial analysis of (temporarily) ambiguous sentences. These principles are language-universal, so they apply to all ambiguities in any language. According to minimal attachment, the processor incorporates an ambiguous phrase into the preceding syntactic tree structure using the fewest number of nodes. This explains the garden-path effect in sentences such as (1). Figure 17.1a shows a simplified tree structure of the reduced relative analysis (Frazier, 1979), while Figure 17.1b shows the tree structure of the main clause analysis. At *examined*, the reduced relative analysis requires more nodes (the circled NP and S nodes) than the main clause analysis is initially adopted. However, this analysis is inconsistent with the disambiguation at *by the lawyer*, so the processor cannot attach this phrase (as indicated in Fig. 17.1b) and has to reanalyze. Hence, (1) is harder to process than an unambiguous relative clause containing *that was* preceding the verb *examined* (e.g. Ferreira and Clifton, 1986; Trueswell et al., 1994). The minimal attachment principle also explains people's parsing preferences in many other syntactic ambiguities across different languages.

If two analyses of an ambiguous structure have an equal number of tree structure nodes, the late closure principle applies. It predicts that people attach an ambiguous phrase to the currently processed phrase. The late closure principle accounts for parsing preferences in many other ambiguities. For example, it predicts that in (2), the relative clause *that was tasty* prefers to attach low to the most recent noun phrase *the sauce*



**Figure 17.1** (a) Simplified tree structure of the reduced relative analysis for *The defendant examined by the lawyer turned out to be unreliable*. (b) Simplified tree structure of the main clause analysis.

rather than high to *the steak* (e.g. Traxler et al., 1998; Gilboy et al., 1995).

2. The steak with the sauce that was tasty didn't win a prize.

In many cases, late closure results in a preference for attachment to the most recent phrase in the preceding part of the sentence, and therefore it makes predictions similar to those of recency principles in other theories (Gibson, 1998; Kimball, 1973; Stevenson, 1994). Proponents of the garden-path model have conducted several studies that showed evidence for garden-path effects predicted by minimal attachment and late closure (e.g. Ferreira and Clifton, 1986; Frazier and Rayner, 1982; Rayner et al., 1983).

The garden-path account assumes a third parsing principle, known as the "active filler strategy," which accounts for the way in which people process unbounded dependencies, as found in relative clauses and questions such as (3).

### 3. Who did the housekeeper from Germany urge the guests to consider?

Following transformational grammars (e.g. Chomsky, 1981), the garden-path theory assumes that the filler *who* has been moved from its original direct object position and leaves a gap. The active filler strategy predicts that the processor fills the gap as early as possible. Because the gap following *urge* (as in *The housekeeper from Germany urged who?*) occurs earlier than following *consider* (the correct interpretation), the active filler strategy predicts a misanalysis in (3), resulting in processing difficulty. Several experiments have provided evidence consistent with this strategy (e.g. Frazier and Clifton, 1989; Frazier and Flores D'Arcais, 1989), though the data are also compatible with accounts without gaps (Traxler and Pickering, 1996); see also Phillips and Wagers (Chapter 45 this volume).

Although the garden-path model has been the dominant modular account, a number of alternative accounts also claim that the processor initially ignores certain sources of information. Most of these assume that the processor prefers analyses that involve a thematic relationship; in other words, it prefers arguments, which receive a thematic role from another word in the sentence (e.g. a verb), over adjuncts, which do not (e.g. Abney, 1989; Crocker, 1995; Pritchett, 1992). In support of this, Schütze and Gibson (1999) found that minimal attachment of the adjunct phrase *for a month* to the verb phrase containing *considered* in (4a) was harder to process than non-minimal argument of the argument phrase *for a raise* to the noun phrase *employee demands* in (4b).

- 4a. The company lawyers considered employee demands for a month but they did not act.  
 4b. The company lawyers considered employee demands for a raise but they did not act.

## 17.3 Reanalysis

An important question in modular accounts is what happens once the processor discovers that the initial analysis is inconsistent with subsequently processed disambiguating information and has to reanalyze (i.e., it has to construct an alternative analysis). We will discuss reanalysis before we move on to interactive models, because it plays such an essential role in modular models. Modular accounts are generally serial models, so they assume that the processor adopts only a single analysis at any one time. Reanalysis occurs

when the initial analysis is inconsistent with later information. Detection of the misanalysis and the subsequent reanalysis of the initial structure result in processing difficulty.

It is usually assumed that reanalysis occurs when the initially adopted analysis is inconsistent with later syntactic information, i.e. syntactic information makes the initial analysis ungrammatical. Implausibility of the initial analysis may also be able to trigger reanalysis (Pickering and Traxler, 1998). However, there is little evidence for reanalysis when the initial choice remains syntactically possible and semantically plausible, but other sources of information (e.g. recency) favor the alternative analysis (Schneider and Phillips, 2001; Sturt et al., 2001).

Models of reanalysis provide an explanation for why some types of reanalysis appear to be more difficult than others (e.g. Gorrell, 1995; Pritchett, 1992; Sturt and Crocker, 1996). These theories have often relied on intuitive judgments about reanalysis difficulty, but there are also a number of experimental studies that have tested different factors affecting reanalysis cost. For example, Sturt et al. (1999) compared reanalysis cost in object/complement clause ambiguities such as (5a) with that in object/null complement ambiguities such as (5b).

- 5a. The Australian woman saw the famous doctor had been drinking quite a lot.  
 5b. Before the woman visited the famous doctor had been drinking quite a lot.

Although the ambiguities were controlled for non-syntactic factors such as lexical frequency preferences, processing difficulty with (5b) was larger than with (5a). Pritchett's (1992) explanation of this difference is that *the famous doctor* has to move out of the thematic domain of *visited* in (5b), whereas it remains within the domain of *saw* in (5a). Alternatively, both Gorrell (1995) and Sturt and Crocker (1996) proposed that the parser has to change hierarchical relations in the tree structure in (5b) but not in (5a).

Several other studies have shown that the length of the temporarily ambiguous phrase (e.g. *the famous doctor* in (5)) affects reanalysis cost (e.g. Ferreira and Henderson, 1991; Tabor and Hutchins, 2004). Ferreira and Henderson (1991) argued that the further the head noun (e.g. *doctor*) is from the point of disambiguation, the stronger the processor commits to a thematic analysis, and the harder reanalysis is. Finally, Sturt et al. (2002) showed that during reanalysis, attachment to a recent phrase is preferred to attachment to a more distant phrase,

suggesting that reattachment to a distant phrase is costly.

However, experiments by Christianson et al. (2001) suggested that people do not always successfully abandon their initial analysis after encountering a syntactic disambiguation, contrary to the assumptions implicit in most reanalysis models. Christianson et al. showed that following temporarily ambiguous sentences such as (6), participants usually answered the question *Who spit up on the bed?* correctly, suggesting that they had correctly analyzed *the baby* as the subject of *spit up*.

6. While Anna dressed the baby that was small and cute spit up on the bed.

The more striking result was that following sentences such as (6), participants more often answered *yes* to the question *Did Anna dress the baby?* than following sentences that were disambiguated by a comma following *dressed*. Hence, Christianson et al. concluded that readers adopted the subject analysis for *the baby*, while at the same time they retained the (incorrect) analysis on which this phrase was the object of the preceding verb.

One possibility is that these results are due to strategic processes that occur when people have to answer the question. However, in a reading study, Kaschak and Glenberg (2004) showed that incorrectly adopted structures affect the reading of subsequent utterances in cases where the correct alternative is a newly learned structure. Furthermore, van Gompel et al. (2006) showed that the initially adopted but incorrect analysis primes the production of subsequent sentences. Hence, these results suggest that the initial analysis retains activation even if the disambiguation is inconsistent with it.

## 17.4 Interactive models

In contrast to modular models, interactive accounts assume that the processor immediately draws upon all possible sources of information during sentence processing, including semantics, discourse context, and information about the frequency of syntactic structures. Current interactive models are usually called *constraint-based theories* (e.g. MacDonald et al., 1994; McRae et al., 1998; Trueswell et al., 1994) and follow from earlier interactive accounts (e.g. Tyler and Marslen-Wilson, 1977). They generally assume that all syntactic alternatives are activated in parallel, with the analysis receiving most support from the various sources of information or constraints being activated most. When one

analysis has a much higher activation than its alternative(s), processing is easy, but when two analyses have an approximately equal activation, competition occurs, and this results in processing difficulty. For example, when the constraints at the beginning of the sentence highly activate one analysis, but disambiguating information later in the sentence activates an alternative analysis, the two analyses have a similar activation at the point of disambiguation. In such a case, it takes a long time before the correct analysis wins the competition and before the incorrect (but initially highly activated) analysis is inhibited. This results in processing difficulty. Note that there is no true reanalysis in this type of model, because both analyses are activated from the onset of the ambiguity, so disambiguation does not necessitate the construction of an analysis that was not initially considered.

Most constraint-based models are lexicalist: they assume that syntactic information is associated with words. For example, it is generally assumed that all verbs contain information about the frequency with which they occur in particular argument frames and that this type of information is used during syntactic ambiguity resolution. Hence, many models assume that there is a tight correspondence between sentence comprehension and production preferences: structures that are frequently produced should be easier to process than structures that are infrequent. One possible way to determine frequency constraints is to experimentally elicit production preferences by asking participants to complete sentence fragments. Constraint-based theories claim that parsing preferences should correlate with such completion preferences (e.g. Garnsey et al., 1997; McRae et al., 1998; Trueswell et al., 1993).

Constraint-based theorists have implemented computational models to explain how various sources of information interact during the processing of reduced relative clause ambiguities. One such model was proposed by McRae et al. (1998) and Spivey and Tanenhaus (1998). In their model, contextual information, semantic constraints, and information about the structures' frequency determine the activation levels of alternative syntactic analyses. At each word, the syntactic analyses compete until one reaches a threshold level of activation and the others are sufficiently inhibited. The longer it takes one analysis to reach threshold, the longer processing times are. Hence, processing is slow when two analyses receive approximately equal support from the different constraints and fast when one analysis receives much more support than its

alternative(s). McRae et al. showed that this model accurately predicted processing difficulty in reduced relatives.

Tabor and Tanenhaus (1999; Tabor et al., 1997) reported a different type of constraint-based model that learns to predict sentence fragment continuations using sentences generated by a grammar which employs information about the frequency with which structures occur. The sentence fragments are represented in a multidimensional space, with fragments which have similar continuations forming clusters that function as attractors. Tabor and colleagues assume that processing times for ambiguous sentences can be modeled as the time it takes for a representation of a sentence fragment to drift to one of the attractors in the space: the more similar a sentence fragment representation is to those in a single cluster and the denser that cluster is, the faster it reaches the attractor. The model predicts that the faster the attractor is reached, the faster processing should be. Because the clusters reflect both syntactic and semantic similarities between sentence fragments, both sources of information should be used immediately during syntactic ambiguity resolution. Tabor and Tanenhaus (1999) showed that in this way reading times for reduced relatives could be accurately modelled.

## 17.5 Semantic effects

Semantic information often provides strongly constraining information for syntactic analysis, so an important question has been whether this information is used immediately to guide sentence processing. According to constraint-based models, semantic information should have an immediate effect on sentence processing, whereas according to modular models, the use of this information should be delayed.

A number of studies have investigated the reduced relative/main clause ambiguity in (7) (see above).

- 7a. The defendant examined by the lawyer turned out to be unreliable.  
 7b. The evidence examined by the lawyer turned out to be unreliable.

When people encounter *examined* in (7a), both the main clause and the reduced relative analysis are plausible. All studies show that reading times for *by the lawyer* are longer than for the same region in unambiguous sentences containing *that was preceding examined*, suggesting that people initially adopt the main clause analysis, and have to revise this when they reach the

disambiguation. The crucial question is whether similar difficulty occurs in (7b), where semantic or animacy information rules out the main clause analysis (evidence cannot examine anything). Some eye-movement reading studies found no immediate effects of semantics, so (7b) did not differ from (7a) (Ferreira and Clifton, 1986; Rayner et al., 1983), whereas another found no sign of difficulty with (7b) at all, in comparison to an unambiguous control (Trueswell et al., 1994). However, a more recent study by Clifton et al. (2003), which used more materials and additional eye-movement measures, did observe difficulty with (7b). Similarly, experiments on different types of ambiguity also show that semantic information fails to override syntactic preferences (e.g. Hoeks et al., 2006; Schriefers et al., 1995). Hence, it appears that, if anything, semantics has only a weak effect on sentence processing.

However, this is not necessarily inconsistent with constraint-based models, so long as semantics provides a fairly weak constraint. Other methods may therefore be more helpful in discriminating between the models. Using a speed/accuracy trade-off method (in which participants are forced to respond quickly whether or not they are confident of the appropriate response), McElree and Griffith (1995) showed that semantics had a slower effect on grammaticality judgements than syntax, even though both provided equally strongly constraining information. Furthermore, in an ERP study, Hagoort (2003) showed that semantic effects, reflected by an N400 ERP response, were larger when the sentence contained a syntactic violation. In contrast, syntactic effects, reflected by a P600 response, were unaffected by semantic violations. Together, these results suggest that syntax affects semantic interpretation and therefore that syntax functionally precedes semantics. In conclusion, most evidence seems to suggest that semantics does not constrain initial syntactic analysis. This is most consistent with modular models.

## 17.6 Frequency effects

Constraint-based theories assume that people make immediate use of information about structural frequency. However, there are different possibilities here. Mitchell et al. (1995) distinguished between fine-grained, lexical frequency information, which takes into account how often specific words (especially verbs) occur in particular structures, and coarse-grained information, which simply considers the frequency of the

structure itself. They argued that a coarse-grained frequency account might explain why different languages have different relative clause attachment preferences in sentences such as (8).

- 8a. The journalist interviewed the daughter of the colonel who had the accident.  
 8b. El periodista entrevistó a la hija del coronel que tuvo el accidente.

In the English sentence (8a) and its Spanish translation (8b), the relative clause *who had the accident* may be attached to either *the daughter* (high in the tree structure) or to *the colonel* (low). Cuetos and Mitchell (1988) found a low attachment preference in English, but a high attachment preference in Spanish. More recent studies have shown either no preference or a weak low attachment preference in English (e.g. Carreiras and Clifton, 1993; 1999; Traxler et al., 1998), whereas many other languages, such as Spanish, Dutch, and French, show a preference for high attachment (e.g. Brysbaert and Mitchell, 1996; Carreiras and Clifton, 1993; 1999; Zagar et al., 1997). Mitchell et al. (1995) argued that attachment preferences may be different between languages because in some languages, high relative clause attachment is most frequent, whereas in a language like English, low attachment is most frequent.

Cross-linguistic differences in relative clause attachment present a problem for the garden-path theory, because late closure predicts a universal preference for low attachment. Hence, Frazier and Clifton (1996) proposed that the garden-path theory only holds for *primary phrases* (roughly, arguments), and that non-structural information can have an immediate effect on the processing of other phrases. Essentially, this implies that the processor is no longer modular for non-primary phrases such as relative clauses.

However, the coarse-grained frequency account has difficulty explaining why there is a high attachment relative clause preference in Dutch (Brysbaert and Mitchell, 1996) even though low attachment is more frequent (Mitchell and Brysbaert, 1998). Furthermore, relative clause attachment preferences are strongly affected by particular words. For example, languages show a strong low attachment preference when the preposition in the noun phrase is *with* (e.g. *the colonel with the daughter*), but either no preference or a much weaker preference with *of* (Gilboy et al., 1995; Traxler et al., 1998); and attachment preferences are also affected by animacy and concreteness (e.g. Desmet et al., 2006). A frequency-based model therefore has to assume that the

processor takes more fine-grained frequency information into account.

Many constraint-based theories assume that the processor employs both coarse-grained and fine-grained lexical frequency information (e.g. McRae et al., 1998; Tabor and Tanenhaus, 1999), but tests of the theories have largely focused on the latter. Trueswell et al. (1993) tested sentences such as (9), where *the solution* is temporarily ambiguous between an object analysis (the student forgot the solution) and the correct complement-clause analysis.

- 9a. The student forgot the solution was in the book.  
 9b. The student hoped the solution was in the book.

The verb *forgot* occurs more frequently with an object, whereas *hoped* occurs most frequently with a complement clause. Trueswell et al. observed that (9a) took longer to read than sentences disambiguated by *that* following the critical verb, whereas there was no difference for sentences such as (9b). This suggests that people use lexical frequency information during syntactic ambiguity resolution (see also Garnsey et al., 1997; Mitchell and Holmes, 1985; Trueswell, 1996). These results are difficult to reconcile with structurally based models such as the garden-path model unless one assumes that frequency information can be used very rapidly to revise initial structural decisions (e.g. Frazier, 1987; 1995). However, a number of studies suggest that lexical frequency information is not used to guide initial processing (e.g. Kennison, 2001; Mitchell, 1987; Pickering et al., 2000). For example, Pickering et al. showed that readers experienced difficulty in (10) shortly after *her exercises* (an implausible object) even though the verb *realised* is biased towards the complement clause analysis.

10. The young athlete realised her exercises one day might make her a world-class sprinter.

One way of explaining these conflicting results is to assume that the verb bias facilitates the complement-clause analysis, but does not completely rule out the object analysis. Therefore, there is some difficulty when the object analysis is implausible.

## 17.7 Discourse effects

Crain and Steedman (1985) argued that many parsing preferences occur because the sentences are presented in isolation. In the absence of a context, people initially prefer to attach the

prepositional phrase *with the dynamite/new lock* in (11) to the verb phrase containing *blew open* rather than to the noun phrase *the safe* (Rayner et al., 1983).

- 11a. The burglar blew open the safe with the dynamite and made off with the loot.  
 11b. The burglar blew open the safe with the new lock and made off with the loot.

But when the same sentences are presented with specific discourse contexts, the preferences may change. Altmann and Steedman (1988) had participants first read a context sentence that referred to either one or two safes. If only one safe had been mentioned, then the complex noun phrase *the safe with the new lock* is unnecessarily specific, so the prepositional phrase *with the new lock* took a long time to read. But if two safes had been mentioned, then the simple noun phrase *the safe* fails to pick out a particular safe, and so the phrase *with the dynamite* took a long time to read. Altmann and Steedman claimed that people initially adopt whichever analysis is compatible with the discourse context, attaching the prepositional phrase to the verb phrase when one safe has been mentioned, but to the noun phrase when two safes have been mentioned. In the absence of any context, verb-phrase attachment is preferred because the processor has to assume one unmentioned safe, which is easier than assuming more than one unmentioned safe. The findings from Altmann and Steedman's (1988) study are consistent with several other studies investigating referential context effects (e.g. Altmann et al., 1992; Van Berkum et al., 1999). However, it appears that one important factor is the strength of the bias when the syntactically ambiguous sentence is presented in the absence of a context. Several studies suggest that referential contexts may affect the processing of weakly biased structures, but not of more strongly biased structures (Altmann et al., 1998; Britt, 1994; Britt et al., 1992; Spivey and Tanenhaus, 1998). In particular, Britt (1994) showed that referential contexts neutralized the preference to attach a prepositional phrase to the verb phrase if the prepositional phrase was an optional argument of the verb (12a, b), but not if it was an obligatory argument (12c, d).

- 12a. He dropped the book on the chair before leaving.  
 12b. He dropped the book on the battle onto the chair.  
 12c. He put the book on the chair before leaving.  
 12d. He put the book on the battle onto the chair.

According to constraint-based theories, discourse information is just one of the many factors that affect sentence processing, so one might assume that it is overridden by other factors in ambiguities that are strongly biased towards one analysis. In such ambiguities, discourse effects may be relatively weak and may therefore show up in later measures of processing. In weakly biased ambiguities, by contrast, discourse may overpower other factors, so discourse effects should be clearer.

Spivey et al. (2002) argued that information from the linguistic context may be forgotten or may not be salient, so it may exert a relatively weak effect on sentence processing. They investigated whether information provided by a visual context that is present during the auditory presentation of a sentence affects processing. They asked people to follow auditory instructions such as (13).

- 13a. Put the apple on the towel in the box.  
 13b. Put the apple that's on the towel in the box.

Spivey et al. used the *visual world eye-movement method* (Tanenhaus et al., 1995). People's eye movements were monitored while they were presented with either a one-referent scene containing a single apple on a towel, or a two-referent scene containing two apples, one of which was on a towel. Both scenes also contained an empty towel without an apple, and a box. When viewing the one-referent scene, people looked more often at the empty towel when hearing (13a) than when hearing (13b). Hence, they appeared to initially misinterpret the prepositional phrase *on the towel* in the temporarily ambiguous sentence (13a) as modifying the verb *put* and took it as the destination for the apple, rather than as the modifier of *the apple*, which is the correct interpretation. But most importantly, in the two-referent scene, no such difference was observed, suggesting that participants immediately interpreted *on the towel* as a modifier of *the apple*. Hence, Spivey et al. argued that visual referential context immediately affected syntactic ambiguity resolution. Chambers et al. (2004) showed that the use of visual context during syntactic ambiguity resolution is affected by action-relevant properties of objects, termed *affordances*. If the only one of the referents in the two-referent scene could be picked up (e.g. the scene contained a liquid and a solid egg), people misinterpreted the temporarily ambiguous prepositional phrase as modifying the verb, so they essentially processed the sentence as in a one-referent scene.

Results from Knoeferle et al. (2005) indicate that people also use visual information that

provides cues about the event roles in the sentence. They tested word order ambiguities in German and showed that depicted actions influenced whether people analyzed an ambiguous noun phrase as an agent (the subject) or a patient (the object) of the verb. However, Snedeker and Trueswell (2004) found less strong effects of visual context with ambiguities that were only slightly different from (13). They showed that visual context influenced syntactic ambiguity resolution, but that it did not completely eliminate the preference for the verb modifier analysis. Interestingly, Trueswell et al. (1999) and Snedeker and Trueswell (2004) found no evidence that children use the visual-world context at all during syntactic ambiguity resolution.

Recent research has also investigated other discourse effects on parsing. Hoeks et al. (2002) proposed that people adopt the simplest possible topic structure. They showed that when a sentence occurs in a neutral context, people preferred the analysis that introduces the fewest number of sentence topics. But when the discourse had already introduced these topics, this preference was neutralized.

Altmann et al. (1998; see also Liversedge et al., 1998) showed that contexts created by indirect questions also affected syntactic ambiguity resolution. In the absence of a question, reading times for *next week* in (14a) were longer than for *last week* in (14b), indicating that readers prefer to attach the temporal phrase to the second verb phrase (*she proposed to the committee*) rather than to the first (*She'll implement the plan*).

- 14a. She'll implement the plan she proposed to the committee next week, they hope.  
 14b. She'll implement the plan she proposed to the committee last week, they hope.

But when these sentences were preceded by an indirect question such as (15), which creates the expectation that the first clause is modified, reading times were longer for (14b) than (14a).

15. The committee members wonder when Fiona will implement the plan she proposed.

Hence, it appears that indirect questions can set up a context that affects sentence processing.

An interesting issue is whether the processing of word order ambiguities in languages with flexible word order is affected by discourse factors. Functional linguistic theories assume that different word orders reflect different information structures of the sentence (e.g. Givón, 1984; 1990). Non-canonical word orders should be straightforward when justified by prior context but difficult otherwise. In accord with this, Kaiser and

Trueswell (2004) showed that sentences with non-canonical object–verb–subject order in Finnish took longer to read than canonical subject–verb–object sentences (see also Hyönä and Hujanen, 1997), but this effect was reduced when the context introduced the object. A visual-world eye-movement experiment suggested that the residual difficulty associated with object–verb–subject sentences was due to people anticipating new information: before hearing the postverbal noun, participants looked more often at new information following object–verb–subject than following subject–verb–object sentences.

In sum, it appears that discourse information often has a very early influence on sentence processing, though it does not always completely override parsing preferences that exist in the absence of a context. If the use of discourse information is delayed, as claimed by modular models, then the delay must be very short—too short to be detected by current psycholinguistic methods.

## 17.8 Testing other properties of the models

In addition to the distinction between modularity and interaction, the dominant models in the sentence processing literature differ in other ways. Most modular models are serial models, i.e. they assume that the processor adopts only a single analysis at a time, whereas most constraint-based models assume that syntactic analyses are activated in parallel in cases of ambiguity. Unfortunately, it has been notoriously difficult to test whether the processor is serial or parallel (e.g. Gibson and Pearlmuter, 2000; Lewis, 2000), because both serial accounts and ranked parallel accounts (in which one analysis is initially favored over others) predict comparable garden-path effects.

A more fruitful way of discriminating between the models is to investigate whether processing difficulty is due to reanalysis, as claimed by the garden-path model and other two-stage accounts, or competition, as claimed by most constraint-based theories. Van Gompel et al. (2005) compared the processing of globally ambiguous sentences such as (16a), where either the bodyguard or the governor may be retiring, with semantically disambiguated sentences (16b/c) and unambiguous sentences (16d).

- 16a. I read that the bodyguard of the governor retiring after the troubles is very rich.  
 16b. I read that the governor of the province retiring after the troubles is very rich.

- 16c. I read that the province of the governor retiring after the troubles is very rich.  
 16d. I read quite recently that the governor retiring after the troubles is very rich.

According to constraint-based competition models, strong competition should occur in (16a), because the constraints equally support both analyses: they are equally plausible and roughly equally preferred. In contrast, competition should be much weaker in (16b/c), because plausibility should immediately affect syntactic ambiguity resolution. However, (16a) was actually *easier* to process than (16b) and (16c) and in fact did not differ from (16d). These results also present difficulty for the garden-path model, which cannot explain why low attachment sentences such as (16c) are harder than globally ambiguous sentences such as (16a). Van Gompel et al. (2005) accounted for these results in terms of the "unrestricted race model" (Van Gompel et al., 2001). This model claims that when there is syntactic ambiguity, the possible analyses are engaged in a race, and that the analysis which is constructed fastest is adopted. The more strongly syntactic and non-syntactic information prior to the point of ambiguity (at *retiring* in (16)) support an analysis, the faster it is constructed, and therefore the more likely it is to be adopted. For balanced ambiguities such as (16), the processor initially adopts each analysis about half the time, because both analyses are about equally preferred. It therefore has to reanalyze about half the time in (16b) and (16c), because plausibility information at *retiring* is inconsistent with the initial analysis. However, it never has to reanalyze when both analyses are plausible, as in (16a), or when the sentence is unambiguous, as in (16d) (see also Traxler et al., 1998; Van Gompel et al., 2001).

Recently, Green and Mitchell (2006) have argued that Van Gompel et al.'s (2005) results can be explained by the competition-integration model proposed by McRae et al. (1998) and Spivey and Tanenhaus (1998), even though this model assumes competition. To make this model work, Green and Mitchell postulated that competition between high and low attachment in (16) occurs from the very first word in the sentence. By the time the relative clause is encountered, only a single analysis is highly activated. Therefore, there should be no competition in the globally ambiguous sentences. In the disambiguated sentences, by contrast, the highly activated analysis may be implausible, and this should result in difficulty. However, given that an infinite number of structures for the rest of the sentence is possible at the beginning of the

sentence, the construction of all these structures should result in a massive working memory load. Because there is no evidence for extreme difficulty at the beginning of the sentence, this assumption seems implausible.

To conclude, there is no reason to assume that competition occurs during syntactic ambiguity resolution. Instead, the results are more consistent with the unrestricted race model, which claims that a single analysis is adopted in a probabilistic fashion, and that difficulty occurs when the initial analysis is implausible. This model also fits well with research showing that non-syntactic information has an early influence on syntactic ambiguity resolution, as it claims that both syntactic and non-syntactic information affect the chance with which an analysis is adopted.

## 17.9 Working memory capacity

Although much research on sentence processing has investigated modularity and interaction, several other issues have also been prominent over the years, and several new strands of research have recently started to emerge. One issue that has been the focus of much interest is the role of working memory capacity in sentence processing. Much debate has centered around the question of whether the working memory resources employed during syntactic processing are different from the working memory resources used for other, more conscious verbal tasks. Just and Carpenter (1992) proposed a *shared resources account* of working memory, in which all linguistic processes draw upon the same limited pool of working memory resources. When people's working memory capacity is exceeded, because either storage or processing demands are very high, this should result in either a processing slow down or a failure to maintain linguistic information in memory. They claimed that individual differences in people's verbal working memory lead to individual differences in sentence processing. These differences can be assessed with a *reading span test* (e.g. Daneman and Carpenter, 1980), which determines how many unrelated words people can remember while reading sentences.

In contrast, Caplan and Waters (1999) proposed the *dedicated resources account*, which assumes that the working memory resources dedicated to obligatory and automatic linguistic processes such as sentence processing are different from those used for more strategic and controlled linguistic processes such as those used in the reading span test.

