resources that are necessary for understanding the basic meaning (i.e., who's doing what to whom) of sentences instantiating the constructions described in the previous section. Because very little work has been done to date on developing a processing model for RRG, my proposals will be pitched at a fairly general level and must be considered tentative.

3.2.1 Parsing and Interpretation

"Parsing" and "interpretation" are technical terms that are frequently used in the psycholinguistic literature. From the point of view of RRG, these terms can be defined as follows: parsing is the process of creating the constituent structures and assigning the syntactic relations of sentences, and interpretation is the process of establishing correspondences between NPs, macroroles, the arguments of predicates, and the concepts expressed by specific nouns. These two kinds of processes are essential for syntactic comprehension, and I will discuss the general properties of each one in turn.

3.2.1.1 Parsing

There are two ways to approach parsing in RRG. One way is to view it as a process of incremental syntactic tree formation driven by simple input-output mapping operations. This is the conception of parsing that is most widely assumed in the psycholinguistic literature (e.g., Dowty et al. 1985; Frazier 1987; Kempen & Vosse 1989; Caplan 1992; Clifton et al. 1994). Basically, these operations take lexical or syntactic categories as input and create elements of constituent structure as output. For example, if a sentence begins with the expression *The dog . . . ,* the series of operations shown below on the left would lead to the immediate assembly of the constituent structure shown on the right ("-->" means "given the unit on the left as input, activate the unit on the right as output"):

a. N --> REF

b. REF --> NUC_N

c. NUC_N --> $CORE_N$

d. $CORE_N$ --> NP

e. NP --> ARG

ARG

NP

$CORE_N$

NUC_N

REF

N

Additional operations would enable the parser to go beyond the tall tree extending from N to ARG and predict that other nodes and branches should appear in the constituent structureùin particular, that an ARG must be dominated by a CORE, and that a CORE must have a CLAUSE and a SENTENCE above it and a NUC and a PRED below it. These anticipated elements of the constituent structure would then get confirmed when the verb of the sentence is encountered, since the lexical category of verb triggers the firing of another series of operations which lead to the construction of these same elements. Further operations must be devoted to assigning syntactic relationsùe.g., pivot, direct core argumentùto appropriate NPs. Thus, according to this view of parsing, the syntactic comprehension system contains a large but finite set of basic mapping operations that are collectively sufficient for creating the constituent structures and assigning the syntactic relations of all possible sentences.

An alternative way to handle parsing in RRG is to treat it as a process of activating and combining syntactic templates that contain precompiled information. This approach is not as popular as the previous one, but an increasing number of researchers are exploring its potentialùresearchers coming not only from a background in sentence proces-sing (e.g., Trueswell & Tanenhaus 1994; Trueswell et al. 1995; MacDonald et al. 1994; Pearlmutter & MacDonald 1995), but also from a background in grammatical analysis (e.g., Jurafsky, in press; Langacker 1987, 1991; Van Valin & LaPolla, in press). The basic idea is that in addition to having a mental dictionary or lexicon that stores mor-

phemes, words, and fixed multiword expressions, people have a syntactic inventory or "syntacticon" that stores complex syntactic units consisting of already assembled constituent structure and, in some cases, already assigned syntactic relations. For instance, there is a family of templates for NPs, including a template for the tall thin tree shown on the previous page, a template for the genitive construction "NP of NP" (e.g., *the box of oranges, the father of the bride*), a template for the possessive construction "NP's NP" (e.g., *the cat's tail, the play's final act*), and so on. There is also a family of templates for cores, including a template for a core with a nucleus and a single argument position (see the intransitive constructions in Figure 19), a template for a core with a nucleus and two argument positions, one a pivot and the other a direct core argument (see the transi-tive active construction in Figure 11), a template for a core with an attached periphery (see the backgrounding passive construction in Figure 12), and so on. Furthermore, there are various templates for complex sentences, such as a template for a clause con-taining two cores, one matrix and the other dependent (see Figures 14-18), a template for a sentence containing two clauses (e.g., for clausal coordination like *Steve went running and then he took a shower*), etc. According to this approach, templates in the syntactic inventory have a resting threshold of activation that is determined by their frequency of occurrence in the language. During the course of on-line sentence processing, multiple templates are activated in parallel to different degrees, and the ones that are most con-sistent with the input are preserved, whereas the ones that do not fit the input are sup-pressedùa functional organization which is in accord with numerous constraint satis-faction models of pattern recognition (Bechtel & Abrahamsen 1991; Churchland & Sejnowski 1992). The complete constituent structure of a sentence is then assembled
by joining together templates at different levels of hierarchical structure, like snapping together Lego pieces.

The two different ways of viewing parsing in RRG are equally coherent from a theoretical standpoint, and I am not aware of any empirical data that strongly favors one over the other (although there is, of course, ongoing debate over the relative merits and shortcomings of each general approachùe.g., see Frazier 1995). In what follows, I will assume the second approach, since it is the view adopted by RRG. With respect to the processing requirements for the specific English constructions described in section 3.1.2, I propose a rough distinction between, on the one hand, parsing operations for creating simple constituent structures, which I define as those containing a single core, and, on the other hand, parsing operations for creating complex constituent structures, which I define as those containing more than one core. This leads to the classification of constructions shown in Table 1 below:

	Construction Type															
Parsing Operation	A	P	SS	SO	OS	OO	SC	OC	SSc	SSn	OSc	OSn	UCa	UCp	AI	UI
Assemble simple constituent structure	x	x													x	x
Assemble complex constituent structure			x	x	x	x	x	x	x	x	x	x	x	x		

Table 1: Syntactic STM for Constructions (Abbreviations: A=active, P=passive, SS=subject-subject relative, SO=subject-object relative, OS=object-subject relative, OO=object-object relative, SC=subject cleft, OC=object cleft, SSc=canonical subject-to-subject raising, SSn=noncanonical subject-to-subject raising, OSc=canonical object-to-subject raising, OSn=noncanonical object-to-subject raising, UCa= active undergoer control, UCp=passive undergoer control, AI=actor intransitive, UI=undergoer intransitive)

3.2.1.2 Interpretation

I turn now to the second general kind of process in syntactic comprehensionùnamely, interpretation. As stated earlier, from the perspective of RRG, interpretation is essentially a matter of linking; more specifically, it involves establishing correspondences betweeen NPs, macroroles, the arguments of predicates, and the concepts expressed by particular nouns. I will discuss each of these types of correspondence in turn.

Several factors influence the process of linking NPs to macroroles, including constituent structure, morphology, and verb-specific properties. The canonical linking pattern for English is manifested in the transitive active construction shown in Figure 11 (p. 67). Here the preverbal pivot NP is linked to the actor macrorole and the postverbal direct core NP is linked to the undergoer macrorole. Many researchers have argued that because this pattern is highly frequent, the syntactic comprehension system treats it as a kind of default (e.g., Bever 1970; Bates & MacWhinney 1989; Caplan 1992). The passive construction shown in Figure 12 (p. 68) reverses this canonical linking pattern, since it requires that the preverbal pivot NP be linked to the undergoer macrorole and the object of *by* be linked to the actor macrorole. This deviation from the standard syntactic-semantic mapping relation is signaled explicitly by three different "cues": the auxiliary, the perfect participial form of the verb, and the preposition *by*. Hence, in order to interpret passive sentences correctly, the syntactic comprehension system must be able to detect these morphosyntactic cues. Because passive sentences have an atypical linking pattern, one would expect them to be more difficult to understand than their active counterparts, and this has been confirmed in several psycholinguistic experiments (Slobin 1966; Forster & Olbrei 1973; Osterhaut & Swinney 1993).[1]

---

[1] It is worth noting, however, that Bever et al. (1989) found significant individual differences in the processing of active and passive sentences as a function of familial handedness. While right-handed individuals who have all right-handers in their families comprehend active sentences much more quickly than corresponding passives, right-handed individuals who have some left-handers in their families comprehend passive sentences slightly faster than corresponding actives. This is part of a more general tendency for familial right-handers to rely

The importance of verb-specific properties in interpretation is exemplified by the subject-to-subject raising constructions shown in Figures 14 and 15 (pp. 74-5) and by the undergoer control construction shown in Figure 18 (p. 78). In the two raising constructions, the predicate **seem** is marked with the feature [0MR], which has the effect of blocking the normal interpretive process of linking the pivot NP in the constituent structure to a macrorole associated with the predicate's LS. And in the undergoer control construction, the process of linking an NP in the matrix core to a macrorole associated with an argument in the embedded verb's LS is guided by the semantic properties of the matrix verb, in accordance with the RRG "theory of control." Thus, in order to correctly interpret subject-to-subject raising sentences and undergoer control sentences, the syntactic comprehension system must be sensitive to special properties of the semantic representations of verbs.

Another important point about establishing correspondences between NPs and macroroles is that although the standard linking process involves mapping an NP in a core to a macrorole associated with the LS of the predicate in the same core, many of the constructions described in section 3.1.2 require a more complex kind of linking process. One such process, which I refer to as cross-core linking, involves mapping an NP in a matrix core to a macrorole associated with the LS of a predicate in a dependent core. This kind of linking is necessary for the noncanonical subject-to-subject raising construction (Figure 15, p. 75) and the noncanonical object-to-subject raising construction (Figure 16, p. 76). In fact, it is worth emphasizing that in the case of these two constructions, the pivot NP is linked *only* to a macrorole associated with the LS of the predicate in the dependent core. Another complex form of linking, which I refer to

---

on morphosyntax more than familial left-handersùsince the morphosyntax of passives is more challenging than that of actives, familial right-handers get slowed down more than familial left-handers.

as cross-clausal linking, involves mapping an NP in a matrix clause to a macrorole associated with the LS of a predicate in a peripheral clause. This is required for all four relative clause constructions and for both cleft constructions, exemplified below for ease of reference (the gaps are strictly for expository purposes; there are no actual empty categories in the constituent structures):

a. subject-subject relative: *The man$_i$ [that ____$_i$ saw Sally] knows me.*
b. subject-object relative: *The man$_i$ [that Sally saw ___$_i$] knows me.*
c. object-subject relative: *I know the man$_i$ [that ____$_i$ saw Sally].*
d. object-object relative: *I know the man$_i$ [that Sally saw ____$_i$].*
e. subject cleft: *It was the man$_i$ [that ____$_i$ saw Sally].*
f. object cleft: *It was the man$_i$ [that Sally saw ____$_i$].*

In the object-subject and object-object relative clause constructions as well as in the two cleft constructions, the head of the complex NP is first linked to a macrorole associated with the LS of the predicate in the matrix clause, and is then linked to a macrorole associated with the LS of the predicate in the peripheral clause. By contrast, in the subject-subject and subject-object relative clause constructions, the head of the complex NP is first linked to a macrorole associated with the LS of the predicate in the peripheral clause, and is then linked to a macrorole associated with the LS of the predicate in the matrix clause.

Although the foregoing consideration of how correspondences are established between NPs and macroroles is far from complete, it provides a useful framework for classifying the constructions described in section 3.1.2 according to the operations that they do and do not share. Such a classification is presented in Table 2.

Before moving on to discuss how correspondences are established between macroroles and the arguments of predicates, I would like to briefly consider one further issue. Recent work in linguistics has shown that grammatical constructions such as passive, dative, causative, locative, etc., are typically associated with rather specific semantic properties (Wierzbicka 1988; Pinker 1989; Jackendoff 1990; Levin 1993). For instance, the prepositional dative construction [NP V NP to NP] is associated with the meaning "X causes Y to go to Z" (e.g., *Sally threw the frisbee to Harry*, *Sally handed the box to Harry*), whereas the double object dative construction [NP V NP NP] is associated with the meaning "X causes Z to have Y" (e.g., *Sally threw Harry the frisbee*, *Sally handed Harry the box*). Given that such construction-specific meanings can exist, it is natural to wonder if some, even many, of the syntactic templates in the "syntacticon" include long-term memory associations between particular NPs and particular macroroles. Thus, it may be the case that the template for the transitive active construction is stored in memory with already established links between the pivot NP and the actor macrorole on the one hand, and the direct core NP and the undergoer macrorole on the other. Similarly, the template for the backgrounding passive construction may be stored in memory with already established links between the pivot NP and the undergoer macro-role on the one hand, and the oblique NP and the actor macrorole on the other; indeed, to get even more concrete, this template may also have a long-term association between the preposition category in the constituent structure and the lexical node for the specific preposition *by*. An approach like this is currently being pursued by several different researchers working independently, and it will be interesting to see where it will lead (e.g., Langacker 1987, 1991; Fillmore et al. 1988; Goldberg 1995; Van Valin & LaPolla, in press). In what follows, I will assume that such an approach is on the right track.

I turn now to the second type of correspondence that must be established when interpreting sentencesùnamely, correspondences between macroroles and argument positions in the LSs of predicates. This issue is essentially about how the semantic relations of predicates are processed. Specific semantic relations (i.e., notions like agent and patient, possessor and possessed, perceiver and perceptual target) are presumably not computed on-line, since they are directly determined by the content and configur-ation of the predicate's LS. For instance, because the predicate **see** (x,y) expresses an activity of visual perception, its first argument is necessarily a perceiver and its second argument is necessarily a perceptual target. With regard to higher-order macroroles, they could either be computed on-line according to the actor-undergoer hierarchy, or they could be stored as components of the long-term memory representations of predi-cates. Although I do not know of any empirical data that bears on this issue, I will adopt the working hypothesis that the latter possibility is true. When this assumption is added to the assumption made earlier that syntactic templates often include long-term associa-tions between NPs and macroroles, it becomes clear that a large part of the process of interpretation involves forming "bridges" between NPs and arguments via activation of the same macroroles. This is illustrated in Figure 20, where the red lines symbolize the correspondences:

SENTENCE

CLAUSE

CORE

ARG NUC ARG

PRED

NP V NP

Actor Undergoer

see' (x,y)

Figure 20: Correspondences between NPs, macroroles, and arguments

In order to fully interpret a sentence, i.e., determine "who's doing what to whom," one last type of correspondence must be establishedùspecifically, between the argu-ments of predicates and the concepts expressed by the nouns in the sentence. This kind of correspondence is what enables a stable representation of the meaning of the sentence to be maintained in short-term semantic memory. Consider, for instance, the processing of the sentence *Harry saw Sally,* which is illustrated in Figure 21. The first step is to activate the abstract lexical units for the words (these units are triggered by the phonological forms of the words, which are not shown in the figure; for evidence supporting the reality of abstract lexical units, see Damasio et al. 1996). Each of these lexical units then activates, in parallel, its associated syntactic and semantic representations. Thus,

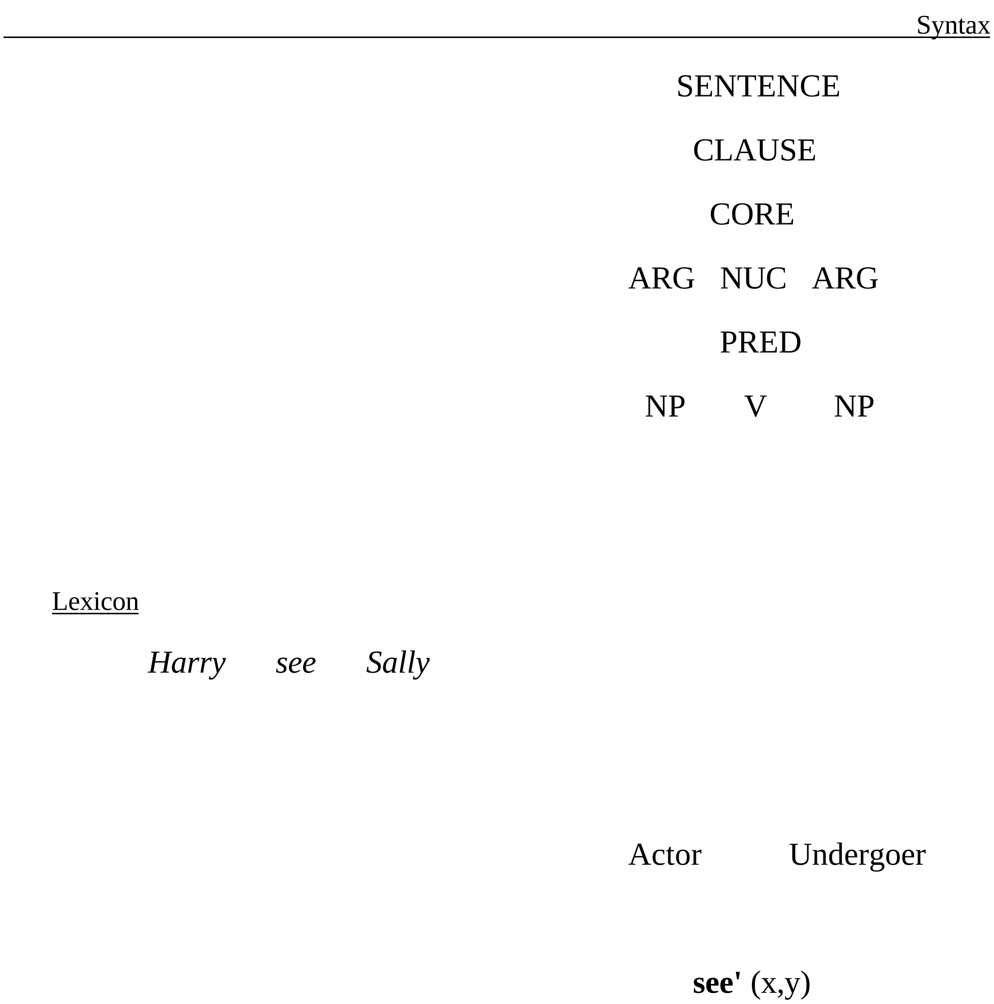

Harry **Sally**

Semantics

Figure 21: Correspondences between argument positions and noun concepts

Harry and *Sally* activate NPs in the syntactic component and the concepts **Harry** and **Sally** in the semantic component, and *see* activates a V in the syntactic component and the predicate **see** in the semantic component. Within the syntactic component, multiple templates compete for dominance until the one that is most consistent with the input is selected, that being the transitive active template. The pivot NP of this template causes the actor macrorole to be activated, and the direct core NP causes the undergoer macro-role to be activated; these correspondences are established because of long-term memory associations between the respective syntactic and semantic units. Meanwhile, within the semantic component the first argument of **see** automatically activates the actor macro-role, and the second argument automatically activates the undergoer macrorole; again, these correspondences are established by virtue of long-term memory associations.

The final step of the interpretation process is to establish the appropriate correspondences between the two arguments of the predicate and the two noun concepts. A close inspection of Figure 21 reveals that the linking between arguments and noun concepts is actually already availableùjust follow the black lines: a chain of correspondences exists between the concept **Harry**, the lexical unit *Harry*, the pivot NP, the actor macro-role, and the first argument of **see**; and another chain of correspondences exists between the concept **Sally**, the lexical unit *Sally*, the direct core NP, the undergoer macrorole, and the second argument of **see**. These two long chains of correspondences provide grammatically mediated linkings between the arguments of the predicate and the concepts expressed by the nouns. But once these indirect correspondences have been established, it is possible to form direct correspondences between the arguments and noun concepts, so that the meaning of the sentence can be maintained in semantic short-term memory after the lexical units and syntactic template have been deactivated. Such direct correspondences are marked with red lines in the figure. I will discuss a possible mechanism for establishing correspondences later in this chapter (see º3.3.2.3, esp. pp. 128-30).

3.2.2 Processing Resources

In addition to requiring operations for parsing and interpretation, syntactic comprehension also requires several different kinds of processing resources that enable the system to function efficiently, especially when dealing with unusally challenging types of constructions. I will focus on the following resources: syntactic short-term memory (henceforth, syntactic STM), and attentional control. Both of these resources have been the subject of recent research on sentence processing, although the former has been studied far more intensely than the latter. I will discuss each one in turn.

3.2.2.1 Syntactic STM

Syntactic STM consists of a limited-capacity buffer that retains constituent structures in an activated or semi-activated state until they can be fully interpreted, after which point they are deactivated so that further syntactic information can enter the buffer (Caplan 1992; Carpenter et al. 1994; Gibson, in press). In short, syntactic STM is a resource for "the remembrance of things parsed" (Pinker 1994: 201). This resource is necessary for processing a variety of constructions. Perhaps its most straightforward function is simply to hold "dangling" elements of constituent structure until they can be completed and mapped into the semantic representation of the sentence. For example, the constituent structure for the initial NP of a sentence, such as *The big red apple* . . . , must be retained in syntactic STM until the predicate of the sentence is encountered and its LS is accessed; then correspondences can be established between the NP, the appro-priate macrorole, and the appropriate argument in the LS of the predicate.

Syntactic STM also plays an important role in the processing of constructions that involve local syntactic ambiguity. For instance, whenever the complementizer *that* is encountered after a noun (e.g., *The man that* . . .), it signals that a relative clause is coming up. This leads to the immediate assembly of a constituent structure for an NP with a periphery containing a dependent clause. However, the complementizer does not provide any information whatsoever about the internal structure of the upcoming clauseùthat is, it doesn't indicate whether the clause is a subject-subject relative, a subject-object relative, or some other type of relative.[2] Recent research suggests that in cases of ambiguity like this, the syntactic comprehension system adopts the strategy of creating several possible constituent structures and maintaining all of them until disam-biguating input is encountered (Hickok 1993).

Another context in which syntactic STM is important is constructions in which the pivot NP is separated from the matrix predicate by intervening material, such as a sequence of prepositional phrasesùe.g., *The park in the central part of the city next to the zoo is a good place to run*). In cases like this, syntactic STM is needed to "bridge the distance," so to speak, from the pivot NP to its predicate. Two of the constructions described in section 3.1.2 require this kind of processingùnamely, the two center-embedded relative clause constructions:

a. subject-subject relative: *The man [that saw Sally] knows me.*
b. subject-object relative: *The man [that Sally saw] knows me.*

---

[2] By contrast, other languages have relative pronouns that do provide such information and thus ease the burden of syntactic comprehension for the listener. While English does preserve the distinction between *who* and *whom*, this is fading out of usage.

As indicated by the brackets, in both of these constructions the relative clause intervenes between the pivot NP and the matrix predicate, and hence syntactic STM is needed to keep the NP in an activated state.

Yet another context in which syntactic STM is crucial is constructions that require filler-gap integration. In such constructions, an NP does not appear in its normal position adjacent to its predicate but rather in a "higher" syntactic position. As a result, the NP cannot be interpreted immediately and hence must be retained in memory until the appropriate predicate is encountered (or, as some researchers say, until the syntactic "gap" where the NP would normally appear is encountered), at which point the NP can finally be interpreted. Several of the constructions described in section 3.1.2 require this kind of processing. First of all, consider the four relative clause constructions and the two cleft constructions, instances of which are shown below with the filler-gap relations marked explicitly:

a. subject-subject relative: *The man*$_i$ *[that ____*$_i$ *saw Sally] knows me.*
b. subject-object relative: *The man*$_i$ *[that Sally saw ____*$_i$*] knows me.*
c. object-subject relative: *I know the man*$_i$ *[that ____*$_i$ *saw Sally].*
d. object-object relative: *I know the man*$_i$ *[that Sally saw ____*$_i$*].*
e. subject cleft: *It was the man*$_i$ *[that ____*$_i$ *saw Sally].*
f. object cleft: *It was the man*$_i$ *[that Sally saw ____*$_i$*].*

In the subject-subject relative (a), object-subject relative (c), and subject cleft (e) constructions, syntactic STM is not required for filler-gap integration because the predicate in the embedded clause is encountered immediately after the complementizer (as noted above, however, syntactic STM is still needed for the subject-subject relative construction in order to hold the pivot NP until the matrix predicate is identified). By contrast, in the subject-object relative (b), object-object relative (d), and object cleft (f) constructions, syntactic STM is needed for filler-gap integration, since the NP must be retained until the predicate in the embedded clause is encountered.

Consider now the two noncanonical raising-to-subject constructions exemplified below:

g. subject-to-subject raising: *Sally*$_i$ *seems to Harry [____*$_i$ *to be tall].*
h. object-to-subject raising: *Sally*$_i$ *is easy [for Harry to see ____*$_i$*].*

Both of these constructions require syntactic STM for purposes of filler-gap integration. In the subject-to-subject raising construction (g), the pivot NP cannot be interpreted until the predicate in the dependent core is identified, and while this NP is being held in memory the oblique NP must be associated with an argument in the LS of the matrix predicate. Similarly, in the object-to-subject raising construction (h), the pivot NP must be retained in memory until the predicate in the dependent core is encountered, at which point this NP as well as the other NP must be interpreted simultaneously. With regard to the two canonical raising-to-subject constructions, neither one involves filler-gap integration, and therefore neither one depends on syntactic STM for this function.

Finally, consider the undergoer control construction, exemplified by the sentence *Harry persuaded Sally*$_i$ *[____*$_i$ *to be nice].* In this construction, an NP in the matrix core must be linked to a macrorole associated with an argument in the LS of the verb in the dependent core. However, because the controller NP is followed immediately by the embedded verb, syntactic STM should not be necessary. On the other hand, when the matrix core is passivizedùe.g., *Sally*$_i$ *was persuaded by Harry [____*$_i$ *to be nice]*ùthe controller NP is separated from the embedded verb by intervening material, and for this reason the NP must be held in syntactic STM until the verb is encountered.

As with the various operations for parsing and interpretation, it is useful to summarize the preceding discussion by classifying the constructions described in section 3.1.2

according to whether or not they require syntactic STM. Such a classification appears in Table 3 below.

A number of studies have focused on the time-course of filler-gap integration (Fodor 1989, 1995; Garnsey et al. 1989; Nicol & Swinney 1989; Boland et al. 1990; Tanenhaus et al. 1990; Kluender & Kutas 1993; Osterhaut & Swinney 1993; Nicol 1994). These studies have employed sophisticated methodologies involving event-related potentials (ERPs) and cross-modal lexical priming (CMLP). Overall, the studies reveal two important properties of on-line filler-gap integration. First, although the constituent structure of the "filler" NP is maintained in a fully activated state throughout the time that it is held in syntactic STM, the semantic representation associated with this NPùthat is, the

	Construction Type															
	A	P	SS	SO	OS	OO	SC	OC	SSc	SSn	OSc	OSn	UCa	UCp	AI	UI
Syntactic STM			x	x		x		x		x		x		x		

Table 3: Syntactic STM for Constructions (Abbreviations: A=active, P=passive, SS=subject-subject relative, SO=subject-object relative, OS=object-subject relative, OO=object-object relative, SC=subject cleft, OC=object cleft, SSc=canonical subject-to-subject raising, SSn=noncanonical subject-to-subject raising, OSc=canonical object-to-subject raising, OSn=noncanonical object-to-subject raising, UCa= active undergoer control, UCp=passive undergoer control, AI=actor intransitive, UI=undergoer intransitive)

concept expressed by the head nounùis maintained in only a partially activated state; however, when the appropriate predicate is identified, the semantic representation of the NP is reactivated for purposes of being associated with an argument in the predicate's LS. Second, this reactivation process generally occurs approximately 400 msec after the predicate is encountered.

Before moving on to discuss attentional control, I should point out that syntactic STM is distinct from another kind of linguistic memory resource that is often discussed in the sentence processing literatureùnamely, verbal STM. As I mentioned in Chapter 2 (see pp. 17, 22), this latter memory resource has two components, one articulatory and the other auditory. It is used primarily for rehearsing single words, multiword sequences, and sentences, and it is typically measured by span tasks that require the subject to remember a list of semantically unrelated items for a given period of time (Baddeley 1986, 1992; Baddeley & Hitch 1994). In much of the early research on sentence processing, it was assumed that verbal STM is essential for on-line syntactic comprehension (Saffran & Marin 1975; Caramazza et al. 1981; Vallar & Baddeley 1984; Caramazza & Berndt 1985). However, more recent research suggests that this is not so. In particular, several studies have shown that brain-damaged patients who have severe impairments of verbal STM are nonetheless able to understand complex constructions that require long-distance filler-gap integration as well as other sorts of long-distance syntactic dependencies (Caplan & Waters 1990; Martin 1990; Waters et al. 1991; Martin & Romani 1994). Still, it may be the case that verbal STM contributes to syntactic comprehension in special situationsùe.g., it may provide a backup phonological representation of a sentence that can be consulted when on-line processing bogs down because of syntactic ambiguities or other challenging operations (Romani 1994), and certainly it is used when one rehearses a sentence in order to make sure that the interpretation derived from "first pass processing" is accurate.

3.2.2.2 Attentional Control

I turn now to attentional control, which is the second major processing resource for syntactic comprehension. Although the role that attention plays in syntactic comprehension has not been the subject of much investigation, there are good theoretical reasons for believing that its role is important, and there are a few studies that have provided empirical support for this view.

As in Chapter 2 (see º2.2.3, pp. 24-25), I will adopt the view that attentional control serves two closely related functions. The first is to amplify the processing efficiency of the syntactic comprehension system, usually at the expense of other mental domains. For example, this aspect of attention may be important when you are listening to someone speaking and there is a great deal of background noise, such as at a party or while standing on a busy streetcorner. The second function is to monitor the activities of the syntactic comprehension system for signs of trouble, so to speak, and when such a sign is detected, to intervene by influencing the selection of structures (e.g., syntactic templates and their associated linking patterns) in a top-down manner. For instance, this aspect of attention is what facilitates recovery from parsing breakdown after pursuing the wrong analysis of garden-path sentencesùe.g., *I thought that the Vietnam war would end for at least an appreciable chunk of time this kind of reflex anticommunist hysteria* (Pinker 1994: 213). In addition, the second function of attentional control may contribute to the processing of contructions that involve noncanonical linking patterns between NPs and macroroles, especially those constructions that are both complex and have few or no overt morphosyntactic cues for noncanonical linking.

It is precisely this last type of situation that I am most concerned with. For example, consider from a purely theoretical standpoint how the sequence of words *the man that Sally saw* might be processed when encountered in a sentential context. The words *the man* trigger the activation of an NP, and the complementizer *that* triggers the activation of a syntactic template for a relative clause. In addition, as I mentioned in the discussion of syntactic STM (see pp. 93-4), the templates for at least two possible kinds of relative clause are also activated as a way of anticipating what is likely to come next. One of these templates is for the subject-relative construction, and the other is for the object-relative construction. The subject-relative template leads to the tentative association of the head NP *the man* with the actor macrorole, and the object-relative template leads to the tentative association of the head NP *the man* with the undergoer macrorole. Moreover, since subject relatives are used more frequently in English than object relatives (Fox & Thompson 1990), the subject-relative template is activated more strongly than the object-relative template; this constitutes a "best bet" prediction about what is going to be encountered downstream. However, when the next few words are encounteredù . . . *the woman saw*ùthe prediction is violated. Since this sequence of words is more consistent with the object-relative template than with the subject-relative template, it causes the activation level of the former to increase and the activation level of the latter to decrease. Notice, though, that the only explicit cue indicating which template is the appropriate one is the order of words in the sequence. It may be the case that this bottom-up input is not sufficient by itself to enable the object-relative template and its associated linking pattern to fully overcome the subject-relative template and its associated linking pattern.

This is presumably the point at which attentional control comes into play. While it is not clear exactly how attentional control operates in computational terms, some very general speculations can be made. Imagine that a monitoring mechanism detects an "impasse" within the syntactic comprehension system and acts on this information by recruiting a decision-making mechanism that is dedicated to resolving such problems. This decision-making mechanism may then operate in either or both of two ways. First, it may intervene in a direct manner by adjusting the activation levels of the relevant templates in the right directions; in particular, it may enhance the activation level of the object-relative template and reduce the activation level of the subject-relative template. Second, it may intervene in a more indirect manner by retrieving the original sequence of auditory word forms from verbal STM and running it through the syntactic comprehension system again, but this time with extra attentional amplification so that the critical word order cue will be sufficient to determine the correct syntactic template and linking pattern. This second type of intervention is undoubtedly more time-consuming, effortful, and consciously mediated than the first and is referred to in the psycholingui-stic literature as "second-pass processing" (e.g., Caplan & Waters 1990; see also Cohen et al. 1990 on how the distinction between automatic, involuntary, and unconscious processes on the one hand, and controlled, voluntary, and conscious processes on the other, is better seen as a continuum than as a rigid dichotomy).

The computational details of how attentional control actually functions are hidden from view just like the secrets behind a magician's tricks. However, the general idea that such a processing resource is frequently needed for syntactic comprehension should not be controversial. People often have the subjective feeling that comprehending sentences that are complex and involve noncanonical linking patterns is more difficult and requires more concentration, sometimes even rehearsal, than comprehending sentences that are comparatively simple and involve canonical linking patterns. The theoretical notion of attentional control is meant to provide a scientific basisùalbeit a very rough one at presentùfor explaining this intuition as well as other data gathered from experimental research.

Within the past few years, several psycholinguistic and neurolinguistic studies have provided empirical support for the view that attentional control plays an important role in syntactic comprehension. Most of these studies have been conducted by a single research team consisting of Carpenter, Just, King, and Miyake (King & Just 1991; Just & Carpenter 1992, 1993; Carpenter et al. 1994, 1995; Miyake et al. 1994, 1995; King & Kutas 1995). Some of this team's most impressive findings come from investigations of the processing of subject-subject and subject-object relative clauses such as the ones shown below:

subject-subject relative: *The reporter that attacked the senator admitted the error*
subject-object relative: *The reporter that the senator attacked admitted the error*

In a replication of previous experiments by Holmes and O'Regan (1981) and Ford (1983), King and Just (1991) demonstrated that in a self-paced word-by-word reading task, the visual fixation times for the two consecutive verbs in subject-object relatives are significantly longer than the fixation times for the corresponding noun and verb positions in subject-subject relatives. While this difference in fixation times is probably due in part to the increased syntactic STM demands for subject-object relatives, it most likely also reflects the increased need for attentional control to regulate template selec-tion and linking in these sentences. Additional evidence consistent with this view was obtained by Just and Carpenter (1993), who showed that in a self-paced word-by-word reading task, pupil dilation increases significantly more for the two consecutive verbs in subject-object relatives than for the corresponding noun and verb positions in subject-subject relatives. Since the degree of pupil dilation is a reliable index of the intensity
of processing (Beatty 1982), it is reasonable to interpret this finding as supporting the hypothesis that template selection and linking are guided by attentional control more in subject-object relatives than in subject-subject relatives. Finally, King and Kutas (1995) observed in an electrophysiological study that the two consecutive verbs in subject-object relatives elicit a distinctive brainwave pattern at the left central frontal and left lateral frontal recording sites, whereas the corresponding noun and verb positions in subject-subject relatives do not. This accords well with the other findings, since, as I will argue in section 3.3.2.5 (p. 136), both the anterior cingulate cortex and the ventro-lateral prefrontal cortex contribute to attentional control for syntactic comprehension.
In addition, King and Kutas found another processing difference between subject-object and subject-subject relatives, one that did not show up in either of the other two studies. In particular, they observed that the determiner immediately following the complementizer in subject-object relatives (e.g., *The reporter that* <u>*the*</u> *senator attacked* . . .) elicited an N400 response at the left Wernicke's and occipital recording sitesùa response which typically indexes the violation of an expectation or the inability to integrate an item into

its preceding context. This suggests that, as I hypothesized earlier, the occurrence of the complementizer causes the subject-relative template and linking pattern to be activated more strongly than the object-relative template and linking pattern, so that when the determiner is encountered, it is a surprise, so to speak, for the syntactic comprehension system.[3] In order to get past this roadblock thrown into the path of sentence proces-sing, it makes sense to assume that top-down attentional control in required to suppress the subject-relative template and linking pattern and promote the object-relative template and linking pattern. However, it may take some time for this intervention to take place: the "impasse" signal must be detected by a monitoring mechanism; the monitoring mechanism must then recruit a special-purpose decision-making mechanism; and finally, the decision-making mechanism must specify a course of action. Hence, the observable effects of attentional intervention do not show up until the predicate of the relative clause is encountered.

It is worth noting that in all three of the studies just described, performance varied across the subjects. Specifically, while the general processing differences between the two relative clause constructions were valid for all of the subjects, they were more pronounced for some of the subjects than for others. This may be due to underlying individual differences in syntactic STM capacity, attentional capacity, or both.

Although the studies conducted by Carpenter and her colleagues focused on the differential involvement of attention in the processing of just two constructionsùsubject-object and subject-subject relative clausesùit is possible to draw inferences from these studies about the degree to which this resource contributes to the processing of the other types of constructions described in section 3.1.2. Consider first the other two relative clause constructions and the two cleft constructions, which are shown below:

---

[3] King and Kutas speculate that the determiner may not have caused longer fixation times or greater pupil dilations in the other studies because the subjects were using a performance strategy of trading accuracy for speed.

a. object-subject relative: *I know the man that saw Sally.*
b. object-object relative: *I know the man that Sally saw.*
c. subject cleft: *It was the man that saw Sally.*
d. object cleft: *It was the man that Sally saw.*

Since the object-object relative (b) and the object cleft (d) both contain a noncanonical complex NP just like in the subject-object relative, the processing of these constructions may require attentional control to suppress the inappropriate syntactic template and linking pattern and facilitate the appropriate ones. By contrast, since the object-subject relative (a) and the subject cleft (c) both contain a canonical complex NP just like in the subject-subject relative, attentional control should not be needed to regulate template selection and linking.

Now consider the raising-to-subject constructions exemplified below:

e. subject-to-subject raising:
   i. canonical: *It seems to Harry that Sally is tall.*
   ii. noncanonical: *Sally seems to Harry to be tall.*
f. object-to-subject raising:
   i. canonical: *It's easy for Harry to see Sally.*
   ii. noncanonical: *Sally is easy for Harry to see.*

The linking patterns in the two canonical constructions (e-i, f-i) are fairly straightforward, so it is not likely that attentional control is needed for on-line processing. On the other hand, the linking patterns in the two noncanonical constructions (e-ii, f-ii) are atypical, and this atypicality is only signaled by a single explicit cue in each case: in the subject-to-subject raising construction, the only explicit cue is the preposition *to,* which indicates that the following NP is the experiencer of *seem;* and in the object-to-subject raising construction, the only explicit cue is the complementizer *for,* which indicates that the following NP is the actor of the predicate in the dependent core. Because there are so few explicit cues for the atypical linking patterns, it is reasonable to assume that during the processing of these constructions, attentional control may be needed to sup-press certain heuristic templates and linking patterns and facilitate the correct ones. For the subject-to-subject raising construction, the heuristic strategy is to treat the NP that is syntactically closest to the embedded predicate as being semantically associated with it; and for the object-to-subject raising construction, the heuristic strategy is to treat the first and second NPs as the actor and undergoer, respectively, of the predicate in the dependent core.

Next, consider the transitive active construction and the two passive constructions:

g. transitive active: *Harry awakened Sally.*
h. passive:
   i. foregrounding: *Harry was awakened.*
   ii. backgrounding: *Harry was awakened by Sally.*

With regard to the transitive active construction (g), it is highly unlikely that attentional control is necessary for on-line processing, since the constituent structure is very simple and the linking pattern is perfectly canonicalùin fact, it's the default. By contrast, the foregrounding and backgrounding passive constructions (h) both involve noncanonical linking patterns, with the pivot NP being mapped to the undergoer macrorole and, in the backgrounding passive, the oblique NP being mapped to the actor macrorole. Hence, one might suppose that attentional control would be needed in order to inhibit the incorrect "active" template and linking pattern and promote the correct "passive" template and linking pattern. I suspect, however, that the situation is not as straightforward as this, since the two constructions not only have very simple constituent structures but also have multiple explicit morphosyntactic cues that signal the noncanonical linking pattern: the backgrounding passive has three such cuesùthe auxiliary, the perfect participial verb form, and the preposition *by*ùand the foregrounding passive has twoùthe auxil-iary, and the perfect participial verb form. Hence, I do not think that attentional control is generally required for processing these sentences. It is worth noting, however, that if attention were needed, it would be needed more for the foregrounding passive than for the backgrounding passive, since the former construction has fewer explicit cues.

I turn now to the active and passive undergoer control constructions:

i. undergoer control:
   i. *Harry persuaded Sally to be nice.*
   ii. *Sally was persuaded by Harry to be nice.*

In order to comprehend undergoer control sentencesùeither active or passiveùan NP in the matrix core must be linked to a macrorole associated with an argument in the LS of the verb in the dependent core. However, which NP must be linked in this fashion is not signaled by any explicit marker whatsoever; instead, it is determined solely by implicit semantic properties of the matrix verb. For this reason, one might think that special attention would be needed for processing undergoer control sentences. I do not think this is the case, however, since it is likely that during the course of on-line sentence processing, the grammatically relevant semantic properties of verbs are strongly activated in an automatic fashion so that attention is not needed to detect or amplify certain features, such as the control features of control verbs (Shapiro et al. 1989; Boland et al. 1990; Garrett 1990). One might still think that attention is required for processing passive undergoer control sentences, since the default strategy of selecting the direct core NP as "controller" may have to be overriden. However, it should not be necessary to suppress one linking strategy and promote an alternative one, since there are multiple explicit cues signaling the noncanonical status of the matrix core (the auxiliary verb, the perfect participial suffix, and the preposition *by*).

Last of all are the actor and undergoer intransitive constructions:

j. intransitive:
   i. actor intransitive: *Harry applauded.*

ii. undergoer intransitive: *Harry drowned.*

Because pivot NPs are typically interpreted as actors, the actor intransitive construction (j-i) has a canonical linking pattern whereas the undergoer intransitive construction (j-ii) has a noncanonical linking pattern. In addition, the noncanonical linking pattern of the latter construction is not explicitly signaled; rather, it is determined by the implicit semantic properties of the verb. Hence, one might suppose that attentional control would be useful for establishing the correct template and linking pattern of undergoer intransitive sentences. I do not think that this inference is valid, however, for the following reasons: first, the constituent structure of the undergoer intransitive construction is very simple; and second, as I argued above in the case of undergoer control sentences, it is likely that the grammatically relevant semantic properties of verbs are strongly activated in an automatic fashion when they are encountered in the course sentence processing. Thus, for an undergoer intransitive sentence like (j-ii), the LS of the achievement predicate **drown** is probably accessed quickly, and the fact that the single argument of this LS is associated with the undergoer macrorole means that the pivot NP can only be an undergoer. Attentional control should therefore not be needed to suppress the alternative interpretation of this NP as actor. It is worth noting, however, that during the processing of undergoer intransitives like *Harry drowned*, there may be a brief period of ambiguity, since the strongest cue that the sentence is in fact an undergoer-intransitive, as opposed to a transitive sentence like *Harry drowned Sally*, is the absence of a direct core NP, and this cue cannot be registered until after the verb has been encountered. During this period of ambiguity, both the intransitive and transitive templates are probably activated, and both the achievement LS and the accomplishment LS of the verb are probably activated. I suspect, though, that the temporary ambiguity is quickly and automatically resolved once the intransitive status of the sentence is established. It may even be the case that intonational cues allow the ambiguity to be resolved before the absence of a direct core NP is registered. In English, focal stress typically falls on the final word (Ladefoged 1993), so that the intransitive sentence *Harry drowned* has focal stress on *drowned*, whereas the transitive sentence *Harry drowned Sally* has focal stress on *Sally*. During on-line processing, then, detection of focal stress on *drowned* may rapidly "tip the balance" in favor of the undergoer-intran-sitive analysis.

The foregoing discussion of the contribution of attentional control to syntactic comprehension is quite general and does not address a number of important questions, perhaps the most challenging of which is how this processing resource functions in precise computational terms. Nonetheless, I hope to have shown that there are good theoretical and empirical reasons for believing that attention plays an important role in syntactic comprehension. My overview of its contribution to each of the constructions described in section 3.1.2 is summarized in Table 4:

	Construction Type															
	A	P	SS	SO	OS	OO	SC	OC	SSc	SSn	OSc	OSn	UCa	UCp	AI	UI
Attentional Control				x		x		x		x		x				

Table 4: Attentional Control for Constructions (Abbreviations: A=active, P=passive, SS=subject-subject relative, SO=subject-object relative, OS=object-subject relative, OO=object-object relative, SC=subject cleft, OC=object cleft, SSc=canonical subject-to-subject raising, SSn=noncanonical subject-to-subject raising, OSc=canonical object-to-subject raising, OSn=noncanonical object-to-subject raising, UCa= active undergoer control, UCp=passive undergoer control, AI=actor intransitive, UI=undergoer intran-sitive)

By way of concluding this section, it is useful to represent together all of the proces-sing operations and resources that are necessary for comprehending the constructions described in section 3.1.2. Such a synthesis is provided in Table 5 below. Because this table contains a great deal of detailed information, it is worthwhile to present a more simplified table in which the various constructions are categorized according to just four critical processing factors: (1) complex parsing, (2) noncanonical linking, (3) syntactic STM for filler-gap integration, and (4) attentional control. This information is provided in Table 6.

	Construction Type															
Processing Factor	A	P	SS	SO	OS	OO	SC	OC	SSc	SSn	OSc	OSn	UCa	UCp	AI	UI
Complex Parsing			x	x	x	x	x	x	x	x	x	x	x	x		
Noncanonical Linking		x		x		x		x		x		x		x		x
Syntactic STM			x	x		x		x		x		x		x		
Attentional Control				x		x		x		x		x				

Table 6: Four Critical Processing Factors for Constructions (Abbreviations: A=active, P=passive, SS=subject-subject relative, SO=subject-object relative, OS=object-subject relative, OO=object-object

relative, SC=subject cleft, OC=object cleft, SSc=canonical subject-to-subject raising, SSn=noncanonical subject-to-subject raising, OSc=canonical object-to-subject raising, OSn=noncanonical object-to-subject raising, UCa=active undergoer control, UCp=passive undergoer control, AI=actor intransitive, UI=under-goer intransitive)

3.3 Neurobiology

In the previous two sections, I characterized the syntactic comprehension system at the levels of structure and processing. In this section, I shift to the final level of analysis, where the aim is to describe how the syntactic comprehension system is physically realized in the brain. I will adopt a methodological strategy called "hier-archical decomposition," which amounts to first establishing the general neural sub-strates of the system as a whole, then attempting to identify the brain areas that support each major subsystem, and ultimately moving further down the scale of functional-