Lexical Structure and Parsing Complexity

Suzanne Stevenson
Department of Computer Science,
and Center for Cognitive Science
Rutgers University
suzanne@ruccs.rutgers.edu

Paola Merlo
Département de Linguistique
and Faculté de Psychologie
Université de Genève
merlo@fapse.unige.ch

To appear in *Language and Cognitive Processes*, 1997.

We gratefully acknowledge the financial support of the Rutgers Center for Cognitive Science (RuCCS), the administrative support provided by Trish Anderson of RuCCS, and the time volunteered by our informants at RuCCS and at the University of Geneva. We thank Charles Clifton, Sven Dickinson, Uli Frauenfelder, Jane Grimshaw, Maryellen MacDonald, Luigi Rizzi, Mike Tanenhaus, John Trueswell, and three anonymous reviewers for helpful comments and discussion on earlier drafts of this manuscript.

Abstract

In recent work on sentence processing, lexical frequencies have been proposed as a primary mechanism for syntactic and lexical disambiguation. In this paper, we instead focus on the consequences that the structural configuration of lexical knowledge has for the time-course of parsing. We concentrate on reduced relative clauses and propose a new lexical-structural analysis for those verbs that are difficult in this construction, manner of motion verbs.

The interaction between the proposed lexical structure and the competitive attachment parser (Stevenson, 1994b) explains the persistent difficulty of this construction with a manner of motion verb, even in disambiguating contexts (for example, a reduced relative in object position) or with non-ambiguous past participle verb forms. Weighted influences on the activation competition are possible with other verbs, and the model can therefore also explain data that demonstrate contextual effects on reduced relatives with simple transitives.

Consequences for the frequency-based models, and all models which base the explanation of the difficulty on the ambiguity of the past participle form, are discussed.

1 Introduction

The last several decades of development in syntactic theory have seen a shift from large systems of complex rules to systems of simple interacting constraints (e.g., Chomsky, 1981; Grimshaw, 1996). As rule systems have become simpler, lexical representations correspondingly have become richer (cf. Bresnan, 1982), and it has become accepted terminology to speak of syntactic structure as being "projected from the lexicon." Recently, the line between the lexical and syntactic components of the grammar has become even more blurred. Hale and Keyser (1993) have put forward a proposal in which lexical entries consist of entire phrases, with potentially embedded complement structure. Their theory of lexical structure accounts for lexical regularities using the same constraints on structural configurations as those that apply in the syntax. Thus as the phrase structure rules governing syntactic structure have become de-emphasized, lexical information has become more and more articulated, until lexical entries have themselves become, in a very real sense, syntactic.

The same evolution of the relation between syntax and the lexicon has occurred within models of human sentence processing. Ford, Bresnan, and Kaplan (1982) provided an early and influential investigation of the effect of individual lexical information on human parsing. Currently, the crucial idea of syntax being projected from the lexicon has developed to a point at which, at least in some central theories, the lexicon is the primary repository of syntactically relevant information. Current constraint-based approaches to modeling human sentence processing suggest that perhaps all of the relevant processing distinctions can be traced to distinctions in lexical information (e.g., MacDonald, Pearlmutter, and Seidenberg, 1994; Spivey-Knowlton and Sedivy, 1995; Spivey-Knowlton, Trueswell, and Tanenhaus, 1993; Trueswell, 1996). Furthermore, these approaches have recently focused on frequency as the key aspect of lexical information in determining preferred interpretations. For example, MacDonald et al. (1994) specifically propose that possible syntactic phrases are stored with an individual lexical item, and these phrases are activated in proportion to the frequency of the lexical item or its conjunction of features. Cooccurrence frequencies may come into play among features of a single lexical item, as well as among features across combinations of lexical structures.

We are in general agreement with syntactic and processing approaches that emphasize rich lexical representations as the source of structural analyses. It is also clear that lexical frequencies contribute to the ease or difficulty of processing particular sentences. However, we feel that the recent attention to frequency effects has been too narrowly focused, and has largely disregarded the critical role of structure in determining interpretation preferences. We note that the very richness of lexical representations makes it unlikely that frequency is the only aspect of lexical structure that is relevant to processing, and we propose that the structural configuration of the lexical knowledge itself has consequences for processing behavior. Our claim is that the structural properties of an analysis are crucial, including the structural distinctions arising from individual lexical information. In particular, we examine the lexical structure of certain verb subclasses, and the effect that the various structures have on the processing of reduced relative clause constructions.

Sentences (1) and (2) each exhibit a main verb/reduced relative (MV/RR) ambiguity at the italicized verb, which can be either a main verb or a past participle.

(1) The horse *raced* past the barn fell.

(2) The butter *melted* in the microwave was lumpy.

The striking aspect of these examples is their sharp contrast in acceptability: although they exhibit the same structural ambiguity, the first is a classic garden path, while the second is readily understood. Pure structure-based accounts rely on the greater complexity of the reduced relative construction to explain its unacceptability in the first sentence (e.g., Frazier and Fodor, 1978; Frazier, 1978; Ferreira and Clifton, 1986), but are unable to account for the ease of the second sentence. Similarly, grammar-based accounts proposed by Pritchett (1992) and Gibson (1991) have claimed that the optional intransitivity of *raced* is the factor causing difficulty in sentence (1), but this also incorrectly predicts that *melted* too should cause a garden path in sentence (2). In contrast to these structural explanations, frequency-based accounts have proposed instead that the ease or difficulty of this construction depends in crucial ways on the frequencies of the ambiguous verb in its various relevant forms—main verb/past participle, active/passive, intransitive/transitive (MacDonald et al., 1994; Trueswell, 1996). In relying on continuously valued frequency information, the constraint-based accounts can potentially address the range of difficulty that has been observed in the processing of this ambiguity (e.g., MacDonald, 1994; Trueswell, Tanenhaus, and Garnsey, 1994).

However, there are two problems with accounts that rely primarily on frequency differentials as the explanation of the degree of difficulty of reduced relative clauses. First, the extremely difficult reduced relatives, such as in sentence (1), remain unacceptable even in unambiguous contexts (e.g., *The stablehand groomed the horse raced past the barn*). Since constraint-based models emphasize relative activation of structures, they do not explain why the reduced relative interpretation would be completely unavailable to the sentence processor, when it is the only possible alternative. Second, while frequency is proposed as a crucial factor in resolving ambiguity, the source of frequency differentials has heretofore been completely unaddressed. Our work is partially motivated by the importance of discovering possible underlying causes of frequency differences, rather than accepting them as primary influences in and of themselves.

Our claim is that the structure of reduced relative clauses has a critical effect on processing behavior. In contrast to earlier structure-based theories, we contend that the relevant structural distinctions arise from differences in individual lexical structure, not solely from general syntactic information. Our proposal is that there is a structural distinction at the syntactic level between some well-known lexical semantic verb classes. This difference in structure correlates with whether a verb yields easy or difficult reduced relative constructions. We note that the relevant differences between verbs is finer-grained than the optional intransitivity proposed by Pritchett (1992) and Gibson (1991) mentioned above. Crucially, we observe that the prototypically difficult RR sentences are those that contain a manner of motion verb (for example, *raced* or *floated*), which are a subclass of *unergative* verbs. We extend the lexical semantic work of Hale and Keyser (1993) to account for the linguistic theoretic properties of this verb class. We argue that the transitive use of an unergative verb has a particularly complex syntactic structure that arises from its lexical properties.

We examine the consequences of this structural analysis within the framework of the competitive attachment parser (Stevenson 1994a, 1994b), a computational model of the human sentence processing mechanism (HSPM). The parser is a parallel network model in which the activation of structure is restricted by the properties of the competitive attachment mechanism. These specific restrictions lead to a severe garden path effect in the processing of reduced relatives with unergative verbs. Other verbs yield reduced relatives that are interpretable within the model, and are therefore open to the weighted influences (such as lexical frequencies) that the parser encodes with activation. Thus we will show that the interaction of the proposed linguistic structure with the critical computational assumptions of the processing model yield a precise account of the range of observed difficulty in MV/RR ambiguities.

2 Lexical and Syntactic Representations

We begin our account with the observation that the prototypically difficult reduced relatives, *The horse raced past the barn fell* and *The boat floated down the river sank* (Bever, 1970), both contain a manner of motion verb. We first describe the properties of verbs of this lexical semantic class (the unergatives) and a related class (the unaccusatives), and discuss their relative acceptability in the reduced relative (RR) construction. We then develop a linguistic analysis of the two verb classes, and show how linguistic factors motivate the assumption of a more complex syntactic structure for RRs with unergatives.

2.1 Verb Classes and Reduced Relative Clauses

On inspection of a set of manner of motion verbs, related to those in the classic garden path RRs, we find that the following verbs all lead to a severe garden path in the RR construction: *advance*, *sail*, *march*, *rotate*, *walk* and *glide*, illustrated in the RRs in (3).

(3) a. The students advanced to the next grade had to study very hard.
b. The clipper sailed to Portugal carried a crew of eight.
c. The troops marched across the fields all day resented the general.
d. The model planet rotated on the metal axis fell off the stand.
e. The dog walked in the park was having a good time.
f. The ship glided past the harbor guards was laden with treasure.

Surprisingly, manner of motion verbs are difficult in a RR even when they have an *unambiguous* past participle form; e.g., *run*, *withdraw*, and *fly*, exemplified in the sentences in (4).

(4) a. The greyhound run around the track all day was tired.
b. The children withdrawn from the religion class could study music instead.
c. The albatross flown to Australia was very tired.

In order to confirm our intuitions, we asked a number of our non-linguist colleagues for informal acceptability judgments on these verbs. All of the sentences with an ambiguous main verb/past participle form, as well as the sentence with *run*, were judged completely unacceptable. The sentences with the unambiguous past participle forms (other than *run*),

Test for Unergativity	Example
The "X's way" construction	The horse raced its way to victory.
The false object construction	The athlete ran herself ragged.
The use of a durative phrase	The troops marched for three hours.

Table 1: Standard tests for unergativity with some manner of motion and sound emission verbs.

although eliciting less uniform responses, were mostly judged very degraded to completely unacceptable. Note that all of the examples were intended to be neutral or biased toward a RR interpretation. For instance, we anticipated that *The albatross flown to Australia was very tired* would be perceived as odd, if presented out of context. Informants were asked their judgments within a discussion of a children's movie in which they had just been told that two mice flew an albatross to Australia. In other words, in addition to the fact that the past participle *flown* is unambiguous, the discourse situation was biased toward a RR reading.

We also found that *ring*, another verb with an unambiguous past participle, was difficult in the RR construction, as in (5).

(5) The bell rung only at Easter rusted from disuse.

Sound emission verbs such as *ring* are related to manner of motion verbs in having the relevant lexical and syntactic properties discussed here, and our results apply equally to them. To simplify the presentation, we will focus on manner of motion verbs, but the interested reader should bear in mind that our claims extend to the sound emission verbs as well.

In English, manner of motion verbs form a subclass of verbs that share critical syntactic properties (Levin, 1993; Levin and Rappaport Hovav, 1995). In the intransitive form, these verbs may be classified as unergative verbs, as we see from the application of some standard tests for unergativity in Table 1 (Levin and Rappaport Hovav, 1995). A subset of manner of motion verbs appear in an intransitive/transitive alternation, as shown in sentences (6) and (7); it is this subclass of verbs that interests us, since the RR construction requires the transitive use of a verb.

(6) a. The horse raced past the barn.
    b. The rider raced the horse past the barn.

(7) a. The lions jumped through the hoop.
    b. The trainer jumped the lions through the hoop.

We observe that the transitive form of unergative verbs is the causative counterpart of the intransitive form (see Brousseau and Ritter (1991) and Hale and Keyser (1993), but also cf. Levin and Rappaport Hovav (1995)). That is, the subject of the intransitive form becomes the object of the transitive form, and the new subject is the causer of the same action as expressed in the intransitive. A first conjecture then might be that the causative nature of the transitive form is the critical component of unergative verbs that leads to processing difficulty in the RR construction.

In order to explore this hypothesis, we turned our investigation to unaccusative verbs (Levin, 1993), such as *melt* or *grow*, since they can also occur in an intransitive/transitive alternation related by causativization, as shown in sentences (8) and (9).

(8) a. The butter melted in the pan.
    b. The cook melted the butter in the pan.

(9) a. Rice grows in Northern Italy.
    b. Farmers grow rice in Northern Italy.

In contrast to the unergative verbs, we found that the following unaccusative verbs, with either an ambiguous or unambiguous main verb/past participle form, were easily interpreted in a RR construction: ambiguous—*melt*, *mutate*, *pour*, *rotate*, as in (10); unambiguous—*begin*, *break*, *freeze*, *grow*, *fly*, as in (11).

(10) a. The witch melted in the Wizard of Oz was played by a famous actress.
    b. The genes mutated in the experiment were used in a vaccine.
    c. The oil poured across the road made driving treacherous.
    d. The picture rotated 90 degrees was easy to print.

(11) a. The lecture begun with a Robin Williams joke was very entertaining.
    b. The window broken in the soccer game was soon repaired.
    c. The chicken frozen to zero degrees lasted for 3 months.
    d. The rice grown in Northern Italy contains more starch than Indian rice.
    e. The kite flown over the water soon went out of sight.

Again, we informally surveyed our non-linguist colleagues and found that unaccusative RRs were overwhelmingly judged completely fine, with a few responses of them being slightly degraded.

Note that two of the verbs, *rotate* and *fly*, were also included in our list of unergative verbs. Following Levin and Rappaport Hovav (1995), we assume that the lexical semantics of a verb and its arguments *combine* to determine the classification as unergative or unaccusative. For example, Levin and Rappaport Hovav note that, in order for a manner of motion verb to be unergative, the agent of the action—namely, the agent of the action expressed by the verb in either the intransitive or transitive form—must have internally caused motion, such as with animals or engines. In *The mice flew an albatross to Australia*, *fly* is unergative because an albatross has internally caused motion. On the other hand, in *The girls flew a kite on the beach*, *fly* is unaccusative, as a kite does not have internally caused motion. The same is true for *rotate*, which distinguishes between *to rotate a model planet* (unergative), and *to rotate a picture* (unaccusative), where only the model planet can have internally caused motion by spinning on its axis. The classifications and acceptability judgments of these two verbs thus depend on the specific sentences used. Both ambiguous verbs, *rotate* and *fly*, were difficult in a RR in their unergative use, and easy in their unaccusative use. We can therefore account for differences in difficulty of the same verb in RR constructions when modifying different types of NPs in terms of the lexical semantic classification of the verb with that type of argument.

The clear patterning of difficulty with the verbs in question reveals two important facts about the unergative and unaccusative verb classes in the RR construction. First, the unergative RRs are all mostly or completely unacceptable, while the unaccusative RRs are all completely acceptable or only slightly degraded. Second, the ease or difficulty of the reduced relative construction appears to be relatively insensitive to the ambiguity of the past participle. Even unergative verbs whose past tense/past participle forms are not ambiguous cause difficulty in a reduced relative clause (for example, *run* and *withdraw*). Given these observations, our basic conclusion is that RRs containing a manner of motion verb never appear to be easy.[1] Moreover, we must abandon our conjecture that causativization alone is responsible for the difficulty of unergative RRs. We are left with a puzzle in which two similar verb classes—those having an intransitive/transitive alternation related by causativization—behave very differently in the RR construction, largely along verb class lines. Next we delve below the surface similarity of these verb classes to determine the critical differences between the two.

2.2 The Lexical Structural Analysis

We propose that the crucial distinction lies in the difference in assignment of thematic, or theta, roles between unergative and unaccusative verbs. Theta roles are labels that identify the lexical semantic content of the argument of a verb. These labels depend on the type of verb, and are therefore said to be assigned from the verb to the arguments. Consider examples (6) and (8), repeated here as (12) and (13).

(12) a. The horse raced past the barn.
b. The rider raced the horse past the barn.

(13) a. The butter melted in the pan.
b. The cook melted the butter in the pan.

In an intransitive unergative, such as (12a), the subject NP *the horse* is an Agent, which performs the core action. In a transitive unergative, for example (12b), the subject NP *the*

[1] It was pointed out to us that the ease of *The car driven on the long trip was in bad shape* presents an apparent counterexample to this statement, since *drive* expresses a manner of motion (Rick Lewis, personal communication). However, *drive* does not have the same causative alternation as we have seen with *race* or *jump*. In *Sara drove the car to school*, note that Sara is causing the car to move, but she is not causing the car to *drive*. Note also that the intransitive alternant of *Sara drove the car to school* is *Sara drove to school*, not *The car drove to school* (cf. *The trainer jumped the lions through the hoop/The lions jumped through the hoop*). These observations lead us to conclude that the transitive use of *drive* should not be classified as a causative unergative. Thus, the verb does not present the relevant unergative alternation that we analyze here, but rather belongs to the class of verbs that can occur with an implicit object, such as *eat* or *read*. Although we did not investigate these verbs in detail, we have no reason to expect them to cause processing difficulty in the RR construction, since they do not have the same lexical or syntactic structure as unergative verbs under our proposal.

rider is the Causal Agent—the agent of the causing action—but the object NP *the horse* is the Agent or Actor of the actual racing. These relationships are perhaps even more evident in the sentence *The trainer jumped the lions through the hoop*. By contrast, the subject of an intransitive unaccusative is clearly a Theme or Patient, the entity undergoing the action, as in sentence (13a). In the transitive form, (13b), the NP subject *the cook* is again the Causal Agent, and the NP object is, as in the intransitive form, a Theme.

While the peculiarly agentive nature of the object of manner of motion verbs has been remarked before (Cruse, 1972; Hale and Keyser, 1987), several accounts describe the object of the verb as a Theme (Levin and Rappaport Hovav, 1995; Harley, 1995). We think that the latter approach is seriously flawed, as it voids the notion of theta role of any semantic sense, since the object has a clear agentive role in the action. The agentive nature of the object of transitive manner of motion verbs is confirmed by the fact that the causer argument can only be a true agent, never an instrument or a natural force (Cruse, 1972; Levin and Rappaport Hovav, 1995); cf. the unergatives in (14) (from Levin and Rappaport Hovav, 1995) with the unaccusatives in (15).

(14) a. *The downpour marched the soldiers to the tents.
b. *The firecracker jumped the horse over the fence.

(15) a. The ice storm froze the pipes.
b. The radiation mutated the genes.

This pattern suggests that the object of an unergative cannot be subordinated to a subject that receives a lesser role in the thematic hierarchy (Jackendoff, 1972). In fact, as (16) shows, the object is necessarily agentive for the causative unergative to be grammatical (Cruse, 1972).

(16) a. John flew the falcon.
b. *John flew the sparks.

Building on this difference in thematic role assignments of the two verb classes, we develop an account based on Hale and Keyser's (1993) approach to lexical syntax, which, in turn, draws on Baker's (1988) incorporation theory and on the proposal of a uniform mapping between theta roles and syntactic structure (the Uniformity of Theta Assignment Hypothesis, Baker (1988, 46)) at the syntactic level. Hale and Keyser (1993) assume that thematic roles correspond to a given structural configuration within a lexical semantic structure. Thus the difference in thematic role assignment between the two classes of verbs under discussion is evidence of a difference in underlying lexical structure. Hale and Keyser argue that the lexical entry of a verb captures its lexical semantics in a syntactic structure *within* the lexicon. The syntactic nature of the lexical entry allows the application of syntactic processes at the level of lexical representation. Verb forms may be created through the combination of lexical items with abstract verbal morphemes, such as DO, BECOME, and CAUSE (Dowty, 1979). The combination operation is one of syntactic head movement within the lexical structure; see Figure 1. The verb incorporates into the morpheme (Baker, 1988), leaving behind a trace, which is subject to the head movement constraint. Consequently, the verb with the incorporated morpheme(s) enters the syntax as a single unit.

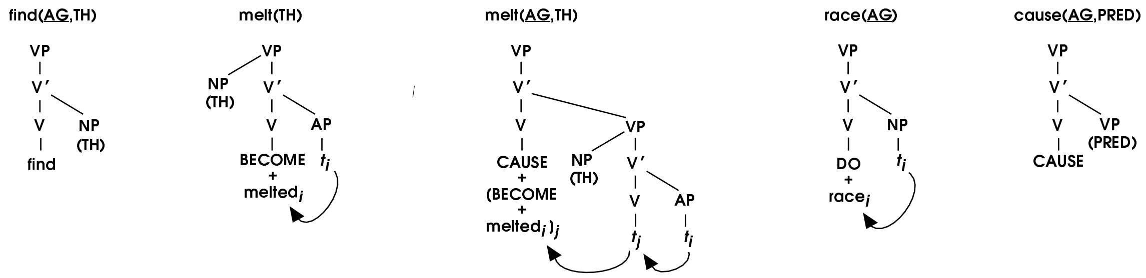

Figure 1: Lexical entries, showing thematic grids and verb incorporation, for various verb types. Thematic role abbreviations: AG=Agent, TH=Theme, PRED=Predicate. External arguments are underlined.

For example, continuing Hale and Keyser's analysis, an intransitive unaccusative such as *melt*, which is a change of state verb, consists of an adjectival head *melted* that incorporates into the abstract verbal head BECOME. On the other hand, an intransitive unergative such as *race*, which is an activity verb, is generated by the lexical incorporation of a noun *race* into the abstract verbal head DO .[2] Theta roles are assigned in these structures as follows. The specifier position (i.e., the sister to the X′ level) of the BECOME morpheme is designated in the lexical structure as the Theme of the incorporated verb, by predication with the adjectival phrase in complement position (see Figure 1). However, the specifier of the abstract verb DO is left empty because, according to Hale and Keyser, a specifier cannot be predicated of an NP in complement position (such as the complement of DO), since the NP is not a predicate. The Agent of the *race* activity must therefore be inserted at the syntactic level, at which point predication with the incorporated verb is possible. The Agent is a true external argument in the sense that it is external to the lexical structure. Thus the Agent theta role is not mapped from the lexicon, but is instead projected by a syntactic rule by virtue of the Extended Projection Principle (Chomsky, 1981).

The causative counterpart for each of these types of verbs requires further incorporation with the abstract verbal morpheme, CAUSE. In the case of the unaccusative verb, *melt*, the incorporation can occur directly in the lexical structure, because the argument structure of the verb is complete—that is, all of its theta roles have been assigned to positions within the lexical structure. Specifically, Hale and Keyser propose that because the Theme argument of *melt* (that is, BECOME + *melted*) is internal to the VP headed by BECOME in the lexical entry, the VP can be further embedded as the lexical complement of another verb. Figure 1 shows the lexical derivation of the causative form of *melt*. Note that, again, the agent of the action, in this case the Causal Agent, cannot be inserted in the lexical structure. An argument can only be inserted in the lexicon when it is predicated of an AP or PP. The

[2]Note that this analysis entails that ambiguous verbs such as *rotate* and *fly*, discussed earlier, have two lexical entries—one for the unergative use and one for the unaccusative use. But this is true for any account that recognizes the thematic distinction between an object that has internally caused motion (an Agent) and an object that does not (a Theme). Ambiguous verbs such as these must then, at a minimum, have multiple thematic assignment possibilities listed in the lexicon; e.g., (Agent) and (Theme) for the intransitive usages.

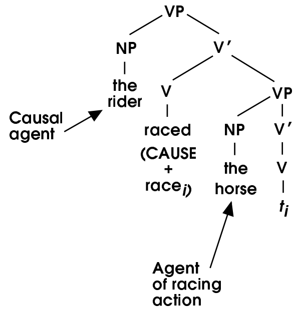

Figure 2: An unergative verb is made transitive through an operation of causativization (incorporation of the verb with the abstract causative morpheme, CAUSE), which applies in the syntactic structure, not in the lexicon.

Agent of CAUSE must be inserted in the syntax.[3]

Within this framework, an unergative verb like *race* cannot be causativized in the lexicon, because the predicate (DO plus *race*) is not saturated—it has an Agent theta role that can only be inserted in the syntax. Because the Agent argument is missing from the lexical structure, the predicate cannot be further embedded as a complement to an additional verbal head within the lexicon. Hale and Keyser (1993) do not address the problem that, nevertheless, a causative form for a subset of these verbs does occur, as evidenced by the transitive alternations above in sentences (6) and (7). We extend their account by proposing that an operation of syntactic causativization may apply in the case of an unergative verb. For the transitive use of *race*, for example, both the abstract morpheme CAUSE and the incorporated activity verb *race* (DO + *race*) must be projected into the syntax; see Figure 2. There, the Agent of the racing action can be inserted into the specifier position of the verb *race*. The abstract verb CAUSE takes the VP headed by *race* as its complement, and an Agent of causation is inserted into its specifier position.

The proposal of a different form of causativization for unaccusative and unergative verbs is not without precedent. Romance languages show a clear asymmetry of these two types of verbs with respect to causativisation, as the unaccusative alternation is quite productive both in French and Italian, while unergative verbs cannot be transitive under any circumstances. As observed in Levin and Rappaport Hovav (1995), morphological processes also confirm the different nature of the causative alternation in these two classes of verbs, as in Hebrew, which allows causatives of both unaccusatives and unergative verbs. Unaccusatives show morphological marking of the intransitive member of the pair (like Italian and French), while unergative causatives show that the transitive form is morphologically derived from

[3]Belletti and Rizzi (1988) observe that the Agent theta role seems to have a special status in the mapping of theta roles to syntax. They notice that psych verbs can be argued to have all the mappings of the pair (Experiencer, Theme) onto syntactic configurations, i.e. the Theme can be the object or the subject, and the Experiencer can be the object, the subject or the indirect object. But there is no verb that has a similar alternation between (Agent, Theme). If there is an Agent it seems to have to be the theta role that is mapped onto the external argument position, i.e. the subject. This supports Hale and Keyser's idea that theta roles have a functional relation to structure and that they are "invisible" as such in syntax, and it derives that Agents are always the external argument.

the intransitive. Thus, Levin and Rappaport Hovav also conclude that unergative verbs are formed by a different causativization process than the one that applies to unaccusatives.

At this point, it is worthwhile reporting why we did not just assume that the relevant difference between unaccusative and unergative verbs is that unaccusative verbs have an adjectival passive (Bresnan, 1982; Levin and Rappaport, 1986). The claim is that an adjectival passive entails a simpler RR structure, without an empty operator and a trace in object position, and therefore could account for the relative ease of unaccusative RRs. On applying the standard tests for adjectival passives, this distinction did not seem to correspond to verb classes. For instance, we found that while unaccusatives are easier in the adjectival *un-* affixation than manner of motion verbs (*unmelted*, *unpoured*, **unrun*, **unraced*), some sound emission verbs are acceptable as complements of a raising predicate (e.g., *The bell remains unrung*)—another test of the adjectival nature of the passive form. (Recall that sound emission verbs are the other subclass of unergative verbs that fall within our analysis.) Moreover, both kinds of verbs are hard in *wh*-fronting constructions (e.g. **How grown was the rice? *How flown was the kite? *How marched were the troops? *How flown is the albatross?*), thus showing that they are verbal passives. Neither of the two classes appear to be easily modified by adjectival adverbs (e.g. **The very grown rice*, **The very run greyhound*), while unaccusatives, if truly adjectival, should be acceptable. Finally, we observe that unaccusative past participles have a verbal nature in other languages, as is shown by clitic-affixation in languages like Italian that show this construction, exemplified in (17).

(17) a. Il burro sciolto nella pentola venne usato per ammorbidire la torta.
'The butter melted in the pan was used to moisten the cake'
b. Scioltolo nella pentola, il burro venne usato per ammorbidire la torta.
'melted-it in the pan,the butter was used to moisten the cake'
c. Scioltone un pezzo nella pentola, il burro venne usato per ammorbidire la torta.
'melted-of-it a piece in the pan, the butter was used to moisten the cake'

Therefore we conclude that the verbal/adjectival distinction is unsatisfactory as it fails to select the right subset of verbs with respect to processing difficulty.

Although we believe that semantic factors are important to the classification of verbs as unergative or unaccusative (cf. the discussion above regarding the ambiguous verbs *fly* and *rotate*), we think that the processing difficulty differential is better explained at the syntactic level, and that such an account is no more complex than a purely semantic one. Hale and Keyser's proposal accounts for distributional properties of lexical items in English (for example, why causativization of unergatives is so limited), which would remain unexplained under a purely semantic account. The explanation is available at no extra theoretical cost, as a mapping between argument structure and syntactic structure is independently needed for the definition of the semantic-syntax interface. In other words, there has to be a set of linking rules between roles and grammatical functions which determines for a given lexical entry what argument becomes the subject, and which the object, for instance. The mapping adopted here (modifying Baker, 1988, and Hale and Keyser, 1993) is the simplest one, where the Theme is always an object and the Agent always a subject at the relevant level of representation, D-structure. Lastly, a syntactic account is supported by a vast body of studies showing that the lexical semantics of unaccusative and unergative verbs entails different syntactic behavior, for instance in the distribution of partitive and reflexive clitics in Romance languages (Belletti and Rizzi, 1988; Burzio, 1981), in the alternation of auxiliary verbs (Burzio, 1981; Burzio, 1986), and in the form of resultative constructions (Levin and Rappaport Hovav, 1995).

2.3 Syntactic Consequences of the Analysis

For verbs that have undergone incorporation within the lexicon, we propose that while the structure of a lexical entry is visible to lexical operations, these articulated lexical structures are not visible to the syntactic level of processing. Figure 3 shows the relation between the underlying (articulated) lexical structures, and the resulting syntactic structures for various verbs. As can be seen in the figure, a verb that results from lexical incorporation—a sequence of head movements within a nested lexical entry—appears in the syntax in a conflated form, within a simple VP. All arguments from the various levels of lexical structure appear as direct complements to the projected verb, except for the external argument (if there is one), which is inserted in the syntax into the specifier of the verb phrase (the [spec, VP] position), and moved into the [spec, IP] position in order to satisfy independent constraints of the grammar (Koopman and Sportiche, 1991).

We make the assumption that all internal arguments appear in the syntax as complements of the verb, regardless of the underlying lexical form, because of the dual nature of certain thematic roles in the lexical and syntactic structures. For example, as shown in Figure 3, intransitive unaccusative verbs such as *melt* have their sole Theme argument in the specifier position of the lexical structure headed by the incorporated verb, capturing the predication relation between the argument and the adjective. However, there is also clear cross-linguistic evidence that this same argument shares syntactic properties with objects in complement position (Burzio, 1981; Burzio, 1986). If the lexical entries themselves are simplified, we lose the lexical structural relationships that give content to the thematic roles. But if the entire lexical structure is projected, then the crucial syntactic generalizations are lost. Furthermore, the thematic requirements encoded by the lexical structure are irrelevant to the syntactic level of processing. By partitioning the information into the relevant level of structure (lexical structure vs. syntactic structure), with a well-defined mapping between the two, we capture the full set of regularities that apply at each level of representation (cf. the approach taken in lexical-functional grammar, Kaplan and Bresnan, 1982).

Recall that our central proposal is that transitive unergatives cannot be formed through lexical incorporation, and must therefore undergo causativization in the syntax. A transitive unergative, such as *race*, requires that two verbal structures be projected into the syntax, the lexically incorporated *race* (DO plus the NP *race*) and the abstract verbal morpheme CAUSE; see Figure 3. Since both are projected into the syntax, the nested verbal structure cannot be conflated in the same way as described above for lexically incorporated verbs. As can be seen in Figure 4, the syntactic operation of causativization entails that the syntactic representation of the transitive form of unergatives is more complex than the representation for transitive unaccusatives—in terms of number of nodes and number of binding relations, and in having the embedded complement structure.

Table 2 summarizes the predicted relative processing difficulty for the different verb classes based on the difference in syntactic complexity under this linguistic analysis. The ta-

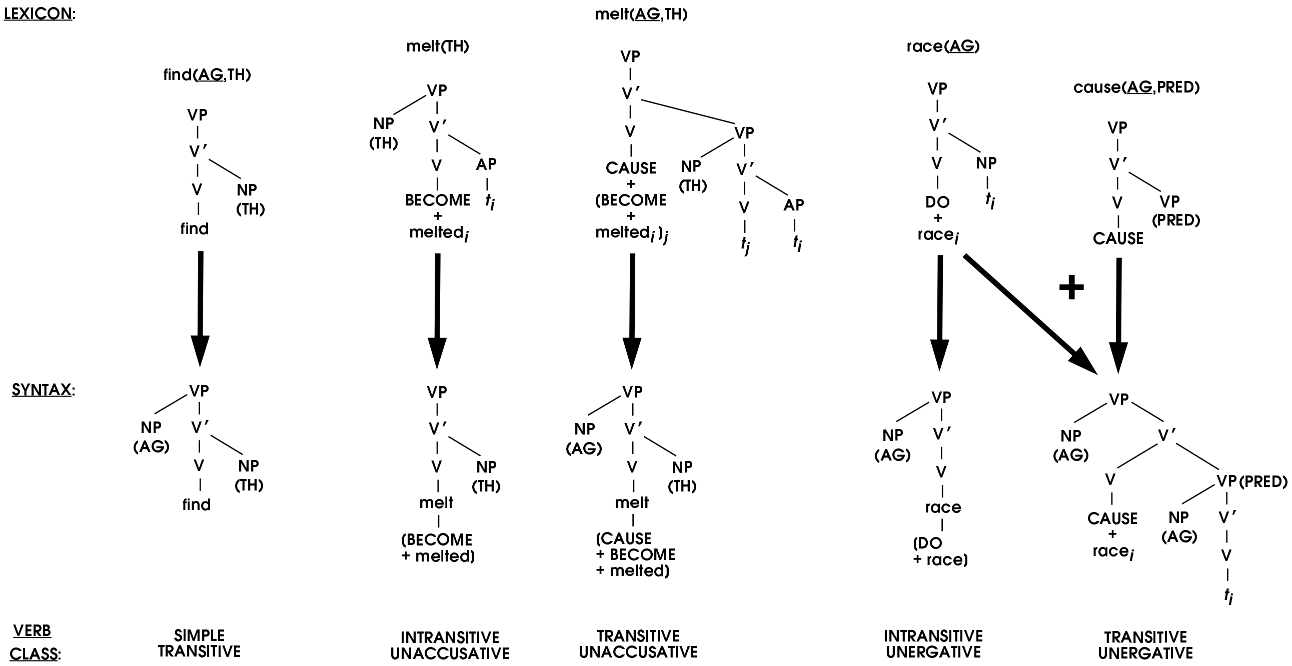

Figure 3: Lexical entries and the corresponding syntactic structures for various verb types.

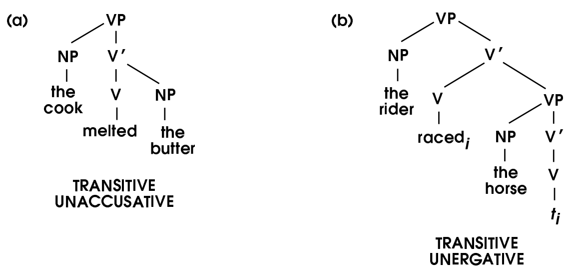

Figure 4: The syntactic structure of the VP for a transitive unergative (b) is more complex than that of a transitive unaccusative (a).

ble distinguishes manner of motion verbs as unergative or unaccusative according to whether they occur with an object which has internally caused motion or not, as discussed earlier. Note that, according to our analysis, the simple passive of an unergative is also expected to be more difficult than others, because of the increased complexity of structure. In fact, some of these verbs are actually unacceptable in the passive form, e.g. *glide* (Jane Grimshaw, personal communication); we discuss this issue further in Section 4.

We are claiming, then, that there is a qualitative difference in the processing of unaccusatives and unergatives due to the greater structural complexity of unergatives. This does not mean, however, that other factors such as lexical frequencies, semantic fit, context, etc., play no role in processing such verbs. However, those factors will be limited in helping an unergative RR, which is inherently more complex (or worse, as we discuss in the processing sections below) than the unaccusative RRs. Unaccusative RRs are more open to

		Reduced Relative		Passive
		Ambiguous	Unambiguous	
Unaccusative	change of state	easier	easier	easier
	MOM+external motion object	easier	easier	easier
Unergative	sound emission	harder	harder	harder
	MOM+internal motion obj	harder	harder	harder

Table 2: Predictions of relative difficulty for different verb classes with different types of objects. MOM = manner of motion verb; "internal motion object" is an object with internally caused motion, and "external motion object" is an object without internally caused motion; see text for discussion.

non-structural influences because those influences are not overwhelmed by structural complexity effects, as we will see shortly. Unaccusative RRs therefore will exhibit a range of difficulty depending on the specific lexical items and specific context.

Given these gradient effects, a natural question then is whether unaccusatives can be as difficult in a RR as unergatives, which would weaken our claim of a qualitative difference between the two. For instance, the following sentences of questionable acceptability were brought to our attention by an anonymous reviewer:

(18) a. The candy caramelized in an hour burned.
b. The wax solidified into abstract shapes melted.
c. The paper yellowed in the sun shrank.

The RRs in (18) each contain a change of state verb, which we have classified as unaccusative, yet these sentences are more difficult than might be expected under our account. However, a structure-based description of their difficulty might also be available. It can be observed that the verbs in these examples belong to that subclass of change of state verbs that express an internally caused eventuality. Compare, for instance, the different acceptability of a sentence where the change of state is indirectly brought about by some natural force, but is not caused by a volitional agent, as in (19a), to a sentence where the subject can be volitional, as in (19b). Notice, moreover, that although the subject of (19b) is volitional, an interpretation of the sentence in which the yellowing of the papers was done without intention is very strongly preferred.

(19) a. The sun yellowed the papers.
b. ??The chain-smoker yellowed the papers.

The same observation appears to be true of the verb *solidify*, since the change of state is probably lexicalized as internal, making the expression of a volitional causal agent awkward, as in sentence (20a). The verb *harden*, on the other hand, expresses an externally caused eventuality, making (20b) perfectly acceptable.

(20) a. ??The sculptor solidified the wax.
b. The sculptor hardened the wax.

The difference between these two subclasses of verbs of change of state—internally caused and externally caused eventualities—has led Levin and Rappaport Hovav (1995) to claim that the correct characterization of the lexical semantics of those verbs that map onto the unergative argument structure is "internally caused verbs," while externally caused verbs correspond, syntactically, to unaccusatives (see also our discussion of *fly* and *rotate* above). From this perspective, it is not surprising to find that the verbs in (18) are harder than other change of state verbs. Whether they have the same structure as unergative manner of motion verbs will be addressed in future research.

To briefly summarize, the linguistic analysis presented in this section posits lexical and syntactic differences between unergative and unaccusative verbs. Next we describe a processing model that is sensitive to these distinctions.

3 Overview of the Processing Model

The competitive attachment model is a parallel spreading activation architecture, consisting of a network of simple processing nodes whose configuration directly represents partial parse trees. The model dynamically builds syntactic structures by incrementally activating processing nodes in response to a sentence, and determining the appropriate network connections among them. Parse tree attachments between individual syntactic phrases represented within the network are determined through distributed processing involving both symbolic and numeric information. The propagation and checking of symbolic features determines the grammaticality of potential phrase structure attachments, based on Government-Binding theory (Chomsky, 1981; Rizzi, 1990). Numeric activation captures the relative strengths of the valid alternatives, and accumulates within a stable configuration representing a partial parse tree for the input seen thus far. The competition for activation among the possible syntactic analyses is the central mechanism of the architecture, since it is responsible for resolving both lexical ambiguity and attachment ambiguity. The core aspects of the symbolic and numeric processing components of the model, including the competitive attachment mechanism, are implemented in a computer system; detailed results of experiments with the implemented parser have been previously reported (Stevenson 1993a, 1994b).

In what follows, we present the properties of the competitive attachment parser that are relevant to the analysis of reduced relative constructions. First we describe the competitive processing within the model, which determines preferences among possible structural representations within the network. Then we turn to another essential aspect of the model—the limitations on the structures that can be represented. The parser is constrained by several computationally motivated restrictions arising from the requirements of the dynamic network creation and competitive activation processes of the model. These restrictions precisely define a narrow range of possible structural hypotheses. In particular, the constraints on the model generally limit the occurrence of processing nodes that represent empty structure. Because empty structure (empty operators and traces) plays an important role in the analysis of reduced relative clauses, these constraints are of central concern; we will see in Section 4 that they are crucial to explaining the garden path status of unergative RRs. We conclude the overview of the parser with an explanation of the processing of verbal projections and syntactic causatives under the linguistic analysis proposed above.

Figure 5: A sample $\overline{X}$ template and its instantiation as a verb phrase.

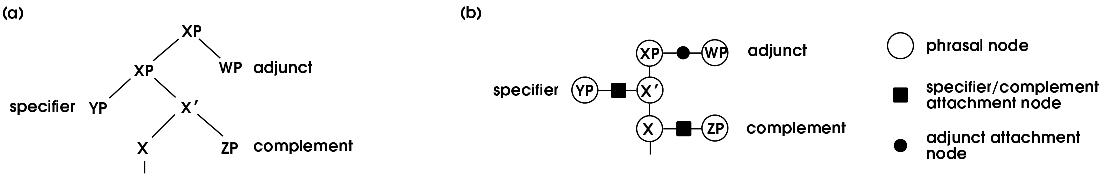

Figure 6: (a) The basic structure of an $\overline{X}$ phrase and (b) its representation in the competitive attachment network.

3.1 The Competitive Attachment Operation

The network of alternative parse tree structures is dynamically created in the model by activating two types of processing nodes. Phrasal nodes are activated from a pool of $\overline{X}$ templates, and their symbolic features and activation level are initialized based on an input token's lexical entry. Figure 5 shows a sample $\overline{X}$ template and its instantiation. A second type of processing node represents the potential attachments between phrasal nodes. Attachment nodes are established between two phrasal nodes that are potential *sisters* in the parse tree; Figure 6 shows the resulting representation of basic $\overline{X}$ relations in the network. The symbolic features and activation level of an attachment node depend on the information it receives from the two phrasal nodes it connects to (the potential parse tree sisters). An attachment node that is highly activated (above an empirically determined threshold) represents an actual sisterhood relation within the current syntactic analysis. If an attachment node is not activated above threshold, its activation level goes to zero, and the attachment does not participate in the parse. Attachment nodes that are activated above threshold can take on a range of activation levels, but for simplicity attachment nodes are shown in all figures as black when fully activated, grey when partially activated, and white when inactive.

When an input token is read, the parser activates the appropriate phrasal nodes. For an ambiguous input token (that is, one with more than one lexical entry), multiple syntactic phrases are activated, each in proportion to the frequency of the corresponding lexical entry. Since the current implementation of the model has no learning component, all frequency values must estimated (e.g., from corpus data or norming studies). After activating the appropriate phrasal nodes, the parser allocates attachment nodes to represent the potential attachments between the new phrasal nodes and previously structured phrasal nodes. As shown in Figure 7, attachment nodes are established between a new phrasal node and any potential sister node along the right edge of an existing parse tree structure. The network then enters its competitive attachment phase that determines which of these potential

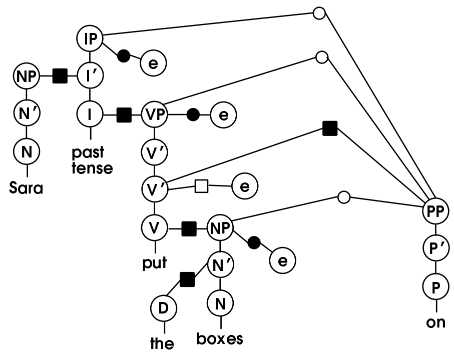

Figure 7: The attachment competition for the current phrase (a PP) determining its attachment along the right edge of the existing parse tree.

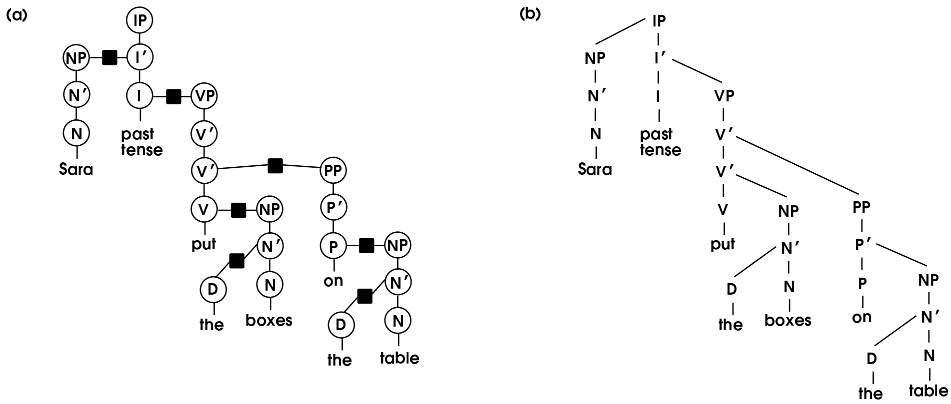

Figure 8: (a) An example network structure for a complete sentence and (b) the corresponding parse tree representation.

attachments will become highly activated. Each phrasal and attachment node iteratively updates and outputs its symbolic features and numeric activation, until the network settles into a stable state in which all attachment nodes either have a high activation level or have an activation level of zero. The highly activated attachment nodes represent the winning attachment decisions established during the current competitive attachment phase. Figure 8 shows the final network structure for a complete sentence, and the parse tree that it represents.

Two types of competitions occur during the network processing described above. An *attachment competition* is a competition at a single phrasal node for the best attachment node to activate. For example, in Figure 7, the PP node must decide whether to activate the attachment to the NP node (as an adjunct of the NP) or the attachment to the V′ node (as the second complement of the verb), or an even higher adjunct attachment to the VP or IP. Since multiple attachments for a phrasal node may be grammatical, a competitive output function (Reggia, 1987) weighs the evidence for an attachment, and allocates activation from the phrasal node to its attachment nodes in proportion to the evidence for each potential attachment. The numeric evidence for an attachment includes frequency information, such as subcategorization preferences, as well as the current activation level of the attachment node. Phrasal nodes thus "hedge their bets" by favoring attachments that are already more highly activated.

This type of interaction is important, since attachment competitions at a single phrasal node do not play out in isolation, but are influenced by the competitive activation decisions being made by other phrasal nodes. For example, referring once more to Figure 7, the competition that takes place among the potential attachments to the PP node is crucially affected by the competitions at the NP, V′, VP and IP nodes, which in turn each have two attachment nodes to choose between. (The nodes labelled "e" in the figure are *empty nodes*, which are described in the following subsection.) Competitive effects thus propagate through the network over time. When all of the attachment competitions stabilize, the set of active attachment nodes for each syntactic phrase represents the current (sisterhood) attachments for that phrase, and inactive attachment nodes connected to that phrase are deleted. The behavior of the implemented system is evaluated in terms of the attachments that are activated, the number of iterations of network processing (i.e., propagation of symbolic and numeric information) that are required, and the precise activation levels of active attachment nodes.

A *lexical competition* uses precisely the same competitive process to determine which of multiple syntactic phrases to activate in response to an ambiguous input token. For example, we mentioned that some verbs, like *rotate* and *fly*, are ambiguous between an unergative and unaccusative classification; such verbs have multiple lexical entries and activate a set of phrasal nodes corresponding to each. Of greater relevance to the results we report below, a verb that has a main verb/past participle ambiguity will activate two VPs, one with the features appropriate to the main verb usage, and one with the features of the past participle usage. Each input token uses a competitive function to output activation to its possible phrasal realizations and choose between them, analogous to the attachment competition process. Since multiple ambiguous input tokens can lead to a large fan-out of competing structures, the competitive activation function must focus activation quickly onto a single alternative. A phrase that makes successful attachments within the current parse tree structure increases in activation, and is thus more likely to be favored with additional activation from the input token and to win the competition. This lexical competition has not been incorporated into the computer system. However, based on the prior extensive testing of the implemented attachment process, we can accurately predict the actions of the parser on the examples here, since the behavior depends on the same critical properties of the competition mechanism. In the simulations described below, the predicted relative activation levels of the structures are reported.

Note that all attachment decisions are made through the parallel competition which involves the attachment nodes connecting a new phrase to the right edge of an existing partial parse tree, as shown in Figure 7. When more than one analysis is possible, the outcome of the competition is influenced by factors such as recency, frequency, and context (Stevenson 1994a, 1994b). However, not all logically possible attachment configurations are possible within the network, since any reanalysis of structure is limited to the consistent reattachment of phrases along the right edge of a tree—previously structured phrases cannot be freely restructured to accommodate new input (cf. a similar restriction in Sturt and

Crocker, 1996). This right edge restriction follows from computational limitations of the competitive attachment function (Stevenson 1993a, 1994b). Only a precisely limited form of reanalysis can be accommodated within the normal attachment operation of the parser, and reanalyses which fall outside of the possible structural reconfigurations have been shown to correspond to a range of well-known garden path effects (Stevenson 1993a, 1994a, 1995). The model therefore exhibits a range of difficulty within the class of constructs parsable within its normal attachment operation, but may also exhibit qualitatively different behavior due to constraints on the possible network configurations that can lead to sharp garden paths.

3.2 Restrictions on Empty Structure

The model has additional constraints on possible network configurations that arise due to computational limitations that affect the hypothesizing of empty phrase structure. These restrictions are of particular relevance here since the representation of a reduced relative clause requires the use of empty syntactic elements. In this section, we simply motivate and describe the constraints, and leave discussion of their ramifications to Section 4, where we show how their interaction with the linguistic analysis of unergative verbs explains the difficulty of unergative RRs.

3.2.1 Empty Nodes

In order for a particular network configuration to represent a valid partial parse tree, each phrasal node must activate a certain number of attachment nodes to satisfy its grammatical properties. However, the competitive function, which determines the amount of activation to output to attachment nodes cannot be successfully parameterized to activate a variable number of attachments. Specifically, the competitive output function *can* be tuned to effectively activate exactly one attachment—as needed for an obligatorily transitive verb. But it cannot be adjusted to allow the activation of zero or one attachment node—as needed for an optionally transitive verb, or even an obligatorily transitive verb whose complement has not been reached in the input. Since the competitive function can bring about activation of exactly one attachment node, but not zero or one, those phrasal nodes that license an attachment must have an attachment node to activate *at all times*. To accommodate this—and to maintain uniformity of numeric processing across all phrasal nodes—every phrasal node automatically posits an attachment to an *empty* node. The attachment to this "dummy" node provides an attachment node to activate even when no overt phrase is available to attach to.

Figure 9 shows the empty nodes that are activated with an $\overline{\text{X}}$ template; each attachment site connects to an attachment node that in turn connects to an empty node. (Empty nodes are labeled "e" in the figures, and are omitted in remaining figures where they are irrelevant.) Because empty nodes occur at all possible attachment sites, they can be used to represent null elements required in a parse—i.e., the empty categories or traces of Government-Binding theory. Traces are phonetically null elements that appear in the position of a syntactically displaced element—e.g., in the underlying position of the NP that moves from object to subject position in the passive construction. For example, in the reduced relative construction, empty nodes are needed to represent both the trace of passive movement, and the empty

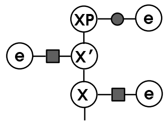

Figure 9: Empty nodes are activated in conjunction with every phrase, at each attachment site.

operator in the specifier of the relative clause. The role of an empty node as an empty category is determined within the network by the attachment and binding relations in which the node participates.

Because they initially have no overt evidence, empty nodes at first have only minimal activation, and the attachment node they connect to is therefore only weakly activated. As the parse of the input progresses, empty nodes may participate in further processing in a number of different ways. Firstly, as noted above, empty nodes are sometimes required in a syntactic analysis to represent empty categories. The coindexation of an empty node with an overt phrase licenses the node as an empty category and instantiates its symbolic features, thereby increasing its activation level. (See Stevenson (1993b, 1994a) for the effects of the active coindexation mechanism on filler-gap preferences and reanalysis.) Secondly, if an empty node occurs in a position that requires an attachment and it is *not* coindexed, it will remain only weakly activated. If an overt phrase competes with the empty node for attachment to the same site, the overt phrase—which is always more highly activated—will win the competition over the weakly activated empty node. Thirdly, a default attachment to an empty node is often replaced in the attachment competition by an attachment to an overt phrase projected from the input. For example, in Figure 7 above, the illustrated attachment competition replaces the default attachment of an empty node to the V′ node with the attachment of the new PP node to the V′ node. Finally, an empty node may occur in a structural location that does not require attachment of a phrase (e.g., as the complement to an intransitive verb). In this case, its grammatical constraints will be vacuously satisfied, and it will therefore immediately become fully activated.

3.2.2 Empty Phrases

Because the parse tree is dynamically allocated within a distributed processing environment, the hypothesizing of structure without incontrovertible evidence in the input must be prohibited in general. Although this limitation is necessary to curb the potential explosion of nested empty structure (which could quickly exhaust the available processing nodes in the parser), it must be relaxed in some highly restricted cases. For example, empty nodes are required to achieve consistent behavior in the competitive output function, but an empty node is only a single, "dummy" phrasal node—the minimum required to provide an attachment site for each phrasal node to activate an attachment to. Allocating articulated empty structure (i.e., a full phrase) is also permitted in some circumstances in order to keep a connected parse tree at all times. Specifically, an *empty phrase*—X-X′-XP—is activated in a situation where an overt phrase requires it in order to make an attachment to the existing parse tree. For example, given a node that can occur as the specifier of a following phrase, the node can activate an attachment as sister to the X′ of an empty phrase to its right. (This is clearly specific to English; in a head final language, a node may activate an empty projection to which it can attach as the complement.) Since an empty phrase is composed solely of empty phrasal nodes, it will initially be only weakly activated. Hence, an input phrase will only successfully activate an attachment as the specifier of an empty phrase if there is no other attachment node (to non-empty structure) that it can activate.

It should be emphasized that both forms of empty structure—empty nodes and empty phrases—are independently motivated, by the requirements of the competitive output function in the case of empty nodes, and by the requirement to maintain a connected structure in the case of empty phrases. Furthermore, they both take the minimal form necessary to support the aspect of processing that necessitates them. Their properties are therefore independently justified, and not just established to obtain some particular parsing behavior.

3.3 Projecting Structure in the Parser

The computational principles of the competitive attachment model, in conjunction with the linguistic representations presented in Section 2, lead to the following processing approach to projecting syntactic structure:

1. Structure is projected from a lexical item, where projection entails the activation of potential extended projections.

2. A verb activates extended projections based on its final (possibly incorporated) position in the lexical structure.

3. The internal structure of the VP of the extended projection encodes features from the predicate argument structure of the verb.

Following Grimshaw (1991), we assume that projections in the syntax consist of a single *lexical* phrase optionally dominated by some number of compatible *functional* phrases. Although our results are entirely compatible with an extended projection analysis of nouns (i.e., as either NPs or DPs), here we will only be concerned with the extended projection of verbs. For example, a verb may be projected into the syntax as a single verb phrase (VP), an inflection phrase (IP) with the VP as its complement, or a complementizer phrase (CP) with the IP and VP embedded; see Figure 10. The VP of each potential extended projection has the specifier and complement requirements determined by the lexical features of the (possibly incorporated) verb. The three potential extended projections will compete for activation from the single verb that initiated their activation, and must also compete for attachments to the existing and developing partial parse tree structures. The winning extended projection is thereby determined by the syntactic requirements of the construction in which the verb participates. Here we assume that the subject of the clause must appear in the [spec, IP] position, and the tense features or auxiliary in the head of IP, so at least the IP level is required. For simplicity, then, we will not display the bare VP possibility in the figures. The CP projection is also required when an element of the analysis must occur in either the head of the CP (e.g., a complementizer such as *that*) or in the specifier of the

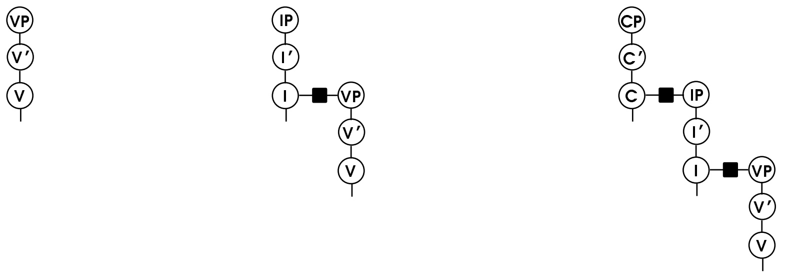

Figure 10: The parser activates a VP, an IP-VP, and a CP-IP-VP, in response to a verb. The three possible extended projections compete for activation, and only one will be fully activated in a stable network.

CP (e.g., a WH operator, such as an overt WH question word, or an empty operator as in the reduced relative).

Initially, the attachment sites of each of the $\overline{\text{X}}$ phrases of an extended verbal projection (VP, IP, or CP) will have attachments to empty nodes, as described above. The lexical information associated with the verb will put constraints on the interpretation of the empty nodes at the specifier and complement attachment sites. If an attachment is not required at the site—e.g., as in the case of the complement position of an intransitive verb—then the category of the attachment site is set to NONE. This specification indicates that no phrase need be attached at that site; an empty node attached there will therefore be fully activated, representing a grammatical structure with no attachment at that site in the parse tree. If an attachment *is* required at the site—e.g., as in the case of the complement position of an obligatorily transitive verb—then the category of the attachment site is set to an actual grammatical category. This specification indicates that the attachment site requires an attachment to a phrasal node of that category. There are two grammatical possibilities in this case. Either the empty node at that site must participate in a valid binding relationship in order to license it as a trace, or the attachment to the empty node must be replaced by an attachment to an overt phrase. The latter situation was illustrated in Figure 7, in which the overt PP replaces the empty node as the second complement of the verb.

Figure 11 shows the VP of the extended projections that are initially activated in response to several different types of verb. A simple transitive verb, such as *find*, activates extended verbal projections in which the VP directly reflects the lexical structure of the verb. The empty node attached to the V has its category feature set to NP, for the internal argument. The empty node in [spec, VP], attached to the V′, has its category feature set to NP, for the external argument. (Recall that the external argument is inserted into the specifier position of the VP, and then moves to the specifier position of the IP.) An incorporated verb, such as *melt*, activates extended verbal projections based on its position following incorporation within the lexical structure. The form of *melt* created from incorporation of BECOME and *melted* is realized as an extended projection of V, in which the VP has a single internal argument and no external argument. The NP that occurs in the [spec, VP] of the abstract element BECOME (the Theme) is realized in the syntax as a direct object of the incorporated verb. Thus this form of *melt* will have an empty node in the complement attachment position whose category is set to NP by the V node. Since it has no external argument, the empty

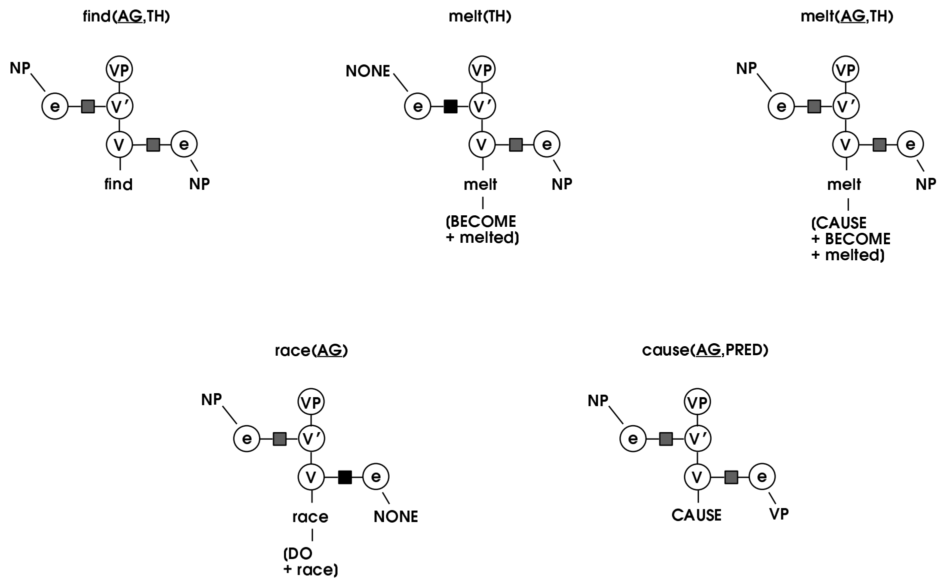

Figure 11: The VP structures initially activated by the parser in response to various types of verbs.

node attached to the V′ has its category set to NONE. The form of *melt* created from incorporation of CAUSE, BECOME and *melted* is also realized as an extended projection of V with a single internal argument. Again the internal NP argument is realized in the syntax as a direct object of the incorporated verb. The difference between this form of *melt* and the intransitive form is that this form also has an external argument to assign (the Causal Agent). Here the empty node attached to the V′ has its category set to NP.

Unergative verbs are distinguished from unaccusative verbs in that the lexical rule of causativization cannot apply, as described in Section 2. The only lexical entry available for a verb such as *race* is the intransitive form (see Figure 3 in Section 2). In the extended projections of an unergative verb, the core VP is therefore one in which the specifier position must be filled by an NP subject, but the complement position is required to be empty. Thus, as shown in Figure 11, the category of the empty node attached to the V node is NONE. In order for an unergative verb to be used in the transitive alternation, it must be combined *in the syntax* with the abstract causative morpheme, CAUSE. For example, Figure 12 shows the network structure for the VP *race the horse*. Only in the syntax can the [spec, VP] of *race* be filled, since it is an external argument. In contrast to a causative unaccusative, the syntactic representation here must include a verbal projection for both CAUSE and *race*, and the two phrases must be explicitly attached (using an attachment node) in the network.[4]

[4]A precise account of which unergatives can undergo causativization remains elusive, both under our proposal and others' (e.g., Hale and Keyser, 1993; Levin and Rappaport Hovav, 1995). One possibility is that causativization of unergatives is a productive process, as proposed by Levin and Rappaport Hovav (1995), but that the added syntactic complexity we propose discourages its use. For now we simply assume that CAUSE is linked in the lexicon to the unergative verbs with which it can so combine. However, although there may be an explicit link indicating those unergatives which can act as complements to CAUSE, this is distinct from the process of *lexical* causativization of the unaccusatives. In the unergative case, the two

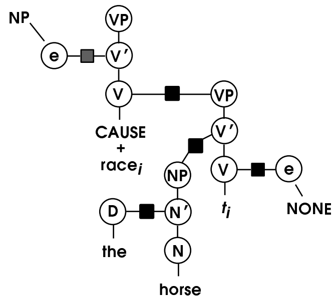

Figure 12: The network structure for the VP of the syntactic causativization of an unergative verb, in which the verb is incorporated with the abstract causative morpheme, CAUSE.

The processing of an unergative in a causative construction in the competitive attachment parser is described in the following section.

3.4 Processing a Syntactic Causative

Consider the sentence *The parents withdrew the child from the class.* When the unergative verb *withdrew* is processed, two interpretations are possible: either the lexical form of the verb is used directly, or the verb is syntactically incorporated with the abstract causative morpheme CAUSE. The parser activates the possible extended projections IP-VP and CP-IP-VP for each of the two alternatives, all four of which compete for activation. The CP portion of each CP-IP-VP projection is an empty structure with no overt evidence, and so neither of them can become fully activated (cf. Grimshaw (1996) and Baković (1995), concerning the projection of minimal structure). At this point then, the highly activated competitors are the two IP-VP forms, one corresponding to the intransitive form of the verb and one to the causative form, as shown in Figure 13. In each case, the core VP has an empty node attached as sister to the V. In the intransitive alternative, the argument requirements of *withdraw* dictate that the empty node has category NONE—that is, the verb licenses no attachment, and so the empty node becomes fully activated. Since all the phrasal and attachment nodes are fully activated, the complete structure is fully activated. By contrast, in the causative alternative, the argument requirements of CAUSE dictate that the empty node has category VP—that is, the verb takes an obligatory VP complement, and so this empty node needs to be bound in order to have full activation. Because the empty node and its attachment node are not fully activated, the complete structure corresponding to the causative alternative has less activation than the intransitive structure, and is therefore less preferred at this point in the parse.

When the NP for *the child* is activated, it cannot attach directly as the complement of either alternative of the verb. However, the NP also triggers activation of an empty phrase, to which it can attach as the specifier. This empty phrase *can* make a successful attachment, as the verbal complement of the incorporated causative verb. At this point, *withdrew* can

verbal structures cannot be incorporated because of the lack of an argument in the [spec, VP] position of the lexical structure of the unergative (e.g., as shown for *race* above in Figure 3).

Figure 13: The two best competitors activated in response to the input *the parents withdrew.*

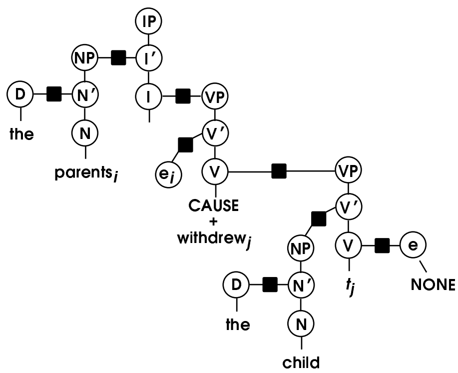

Figure 14: The network encoding the syntactic structure for *The parents withdrew the child*, containing the transitive alternation of the unergative verb *withdrew.*

also bind the empty head of the hypothesized verbal projection, forming a valid syntactic structure, as shown in Figure 14.

4 The Processing of Reduced Relative Constructions

This section presents the central results of the paper: the consequences of our linguistic analysis and processing model for the interpretation of reduced relative clauses. We show that inherent restrictions on the parsing mechanism interact with the proposed syntactic analysis to precisely account for the difficulty of unergative reduced relatives in contrast to other verb types. We begin with the processing of verbs that have a main verb/past participle ambiguity. We assume that such a verb has three lexical entries that are homophonous: an active main verb, an active past participle, and a passive past participle. We also assume that a reduced relative clause is a CP-IP-VP projection with an empty operator in [spec, CP]. The empty operator in the competitive attachment parser is realized as an empty node in [spec, CP] that binds an empty node in [spec, IP], which in turn binds an empty node complement of the VP. (Note that other linguistic analyses are possible (e.g., Siloni, 1994), but whether the RR is assumed to be a CP or DP is irrelevant for the results here.)

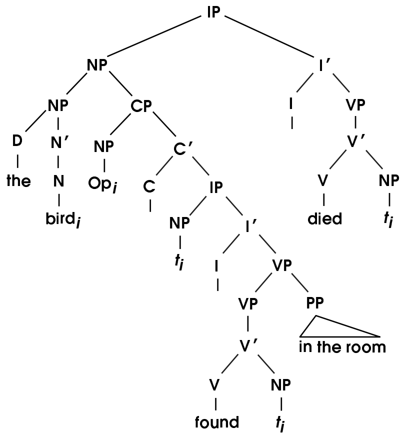

Figure 15: The syntactic structure for a reduced relative construction, in *The bird found in the room died.*

4.1 Obligatorily Transitive Verbs

First consider a reduced relative construction with an obligatorily transitive verb, as in *The bird found in the room died* (Pritchett, 1988), whose syntactic structure is shown in Figure 15. At the point that *found* is processed, the main verb alternative will yield the two competing extended projections shown in Figure 16. Because the empty node in [spec, CP] cannot be interpreted as an empty operator, the CP has no evidence (neither a head nor a specifier), and the CP-IP-VP projection becomes less active. Thus, the IP-VP is the most activated projection for the main verb reading. The past participle reading will similarly activate two extended projections, as shown in Figure 17. We will only consider the attachment between the past participle projection and the NP *the bird* in which the verbal projection is a modifier of the NP, since the NP cannot grammatically attach to the specifier position in any of the extended projections. The IP-VP of the past participle form will have an empty node in [spec, IP] that can bind the empty node that is complement to the V node, but the empty node in [spec, IP] is left unbound, and the projection cannot support the predication relation with the noun modified by the reduced relative. The CP-IP-VP has the same binding relation between the empty nodes attached to the [spec, IP] and complement of V positions. But here the empty node in [spec, CP] is able to bind the node in [spec, IP] and participate in the predication relation with the noun *bird*. This establishes the empty node in [spec, CP] as an empty operator, leading to a grammatical reduced relative construction. Because the empty nodes in the CP-IP-VP projection are all interpreted, this extended projection has more activation than the IP-VP alternative.

This leaves two competing alternative interpretations: (1) the IP-VP projection corresponding to the main verb lexical entry of *found*, which has the NP *the bird* attached as its subject, and has an empty node as the complement of the VP; and (2) the CP-IP-VP projection corresponding to the past participle lexical entry of *found*, which modifies the NP *the bird*. Neither alternative emerges as a clear winner, since the first is waiting on a required complement, and the second is waiting on a (presumably required) predicate. When the PP *in the room* is processed, it can attach to either structure, and does not settle the competi-

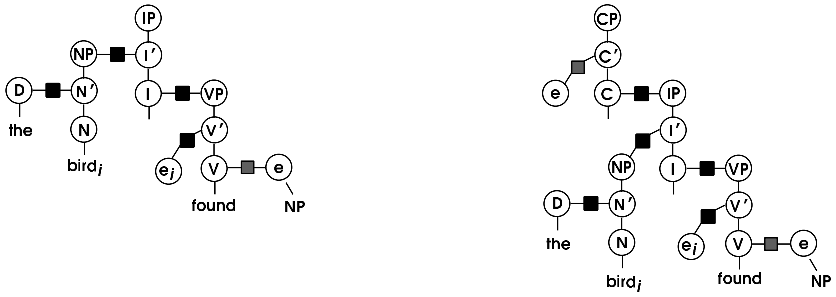

Figure 16: The extended projections within the network corresponding to the main verb reading of *found*.

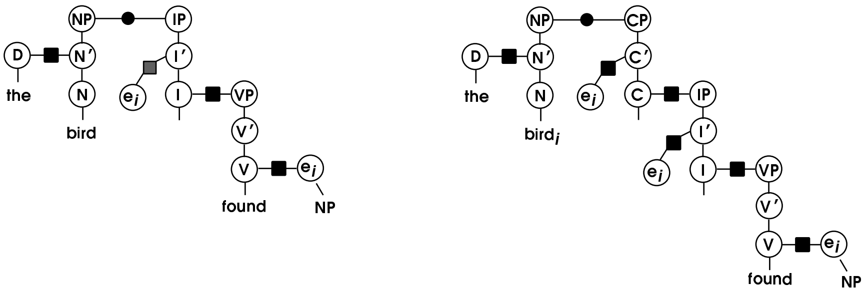

Figure 17: The extended projections within the network corresponding to the reduced relative reading of *found*, and their attachment as modifier of *the bird*.

tion, although the activation of the main verb reading is decreasing for lack of a binder for the empty object. The closeness of the competition allows for the effects of context and theta role compatibility to influence preferences at this stage (MacDonald, 1994; Trueswell et al., 1994). Given the compatibility of the NP *the bird* with either the external or internal theta role of *found*, and the lack of contextual bias, both structures will continue to remain active. When the final verb *died* is processed, the only grammatical alternative is for the NP *the bird*, which is modified by the reduced relative clause, to combine with the IP-VP extended projection of *died*, yielding the correct structure for the sentence, shown in Figure 18.

4.2 Unergative Verbs

The ease of sentences like *The bird found in the room died*, which contains an obligatorily transitive verb in the reduced relative construction, has received much attention because of the stark constrast with the prototypical garden path sentence, *The horse raced past the barn fell.* Figure 19(a) shows the syntactic structure of the latter sentence, under our analysis of the transitive alternative of *race* as an explicit syntactic causative. The proposed structure is clearly more complex than that of a reduced relative with a simple transitive verb such as *found*, repeated here in Figure 19(b). The passive participle *raced* requires a VP complement with two traces. One is a trace of the NP-movement of the passive construction; it is in the [spec, VP] position and is bound by the empty operator, through the trace in [spec, IP].

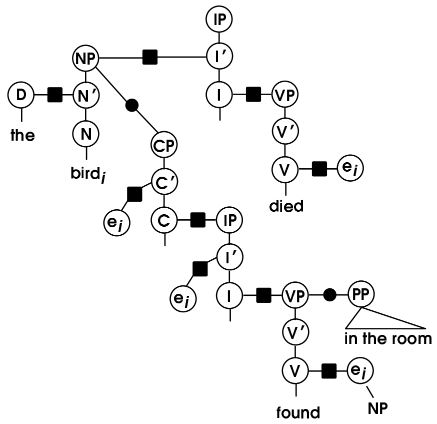

Figure 18: The final network configuration for *The bird found in the room died.*

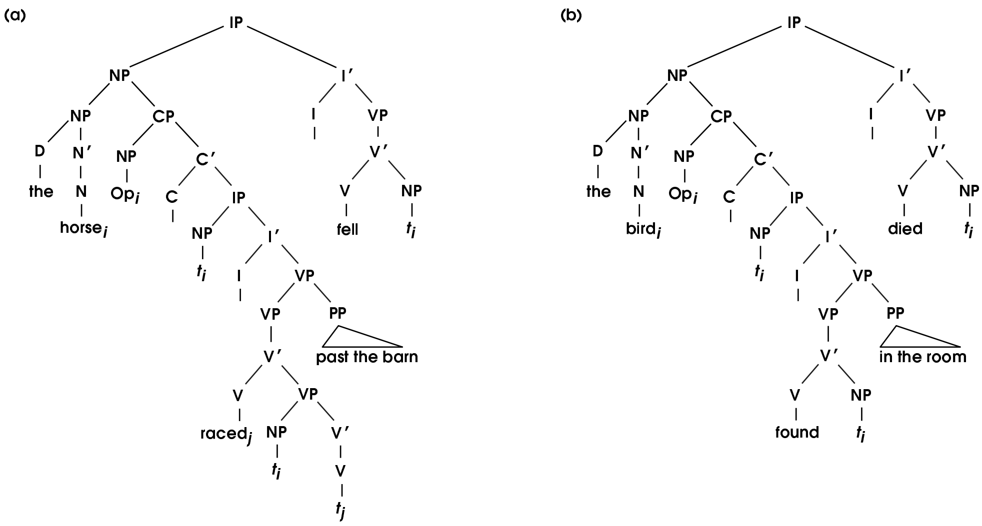

Figure 19: (a) The syntactic structure for *The horse raced past the barn fell*; (b) the syntactic structure for *The bird found in the room died*, repeated here for comparison.

The other one is the trace of verb movement resulting from the syntactic incorporation of the verb with the CAUSE morpheme; it occurs in the embedded V and is bound by *raced*. Thus the structure not only contains an empty VP complement, but one which has internal structure with two traces.

We can now examine how this structural analysis affects the competitive attachment model in its processing of an MV/RR ambiguity with an unergative verb. Like the unambiguous unergative verb *withdrew*, the verb *raced* has two main verb readings, the intransitive and the causative. For reasons analogous to those discussed above in relation to *withdrew*, the two best competitors of the possible extended projections of the main verb reading of *raced* are the two IP-VP forms corresponding to the intransitive reading and the causative reading; these are shown in Figure 20. The empty node that is sister to V in the intransitive form, in contrast to the main verb reading of *found*, is specified to be empty, so all gram-

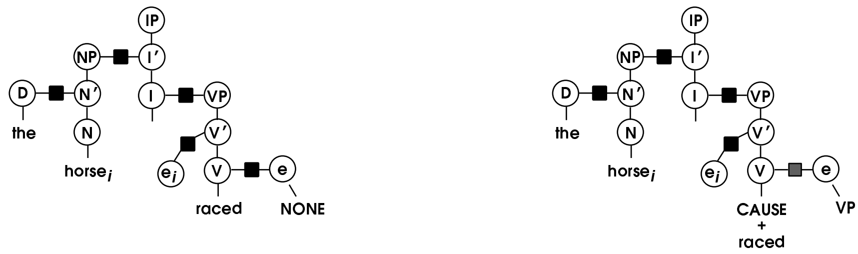

Figure 20: The two best competitors of the set of extended projections within the network corresponding to the main verb reading of *raced*.

matical relations are satisfied in this structure. The causative form, however, is similar to the main verb reading of *found* in having a sister to the V node that cannot remain empty. Since the empty node is not bound, this structure has less activation and is less preferred.

Extended projections are also activated based on the intransitive and causative forms of the past participle reading of *raced*. The intransitive past participle reading decreases in activation, since it can have neither a passive interpretation nor the complex perfect tense interpretation required. The causative past participle interpretation of *raced* yields the two initially activated projections of Figure 21. As in the case with *found*, there is an empty node attached as the sister to the V node that must be bound. However, in the case of *raced*, the empty node is specified to be a VP, according to the CAUSE morpheme with which *raced* has been incorporated. This category specification leads to a situation in which neither of the extended projections can establish a valid grammatical structure. As with the verb *found*, the IP-VP projection has unbound traces and cannot participate in the appropriate predication relation with the noun modified by the relative clause. The CP-IP-VP projection, however, is very different when the verb is *raced* rather than *found*. In the current example, the empty node in [spec, IP] binds the complement of the verb, but this establishes a VP chain. To bind the empty node in [spec, IP], the empty operator in [spec, CP] would inherit the V category of the chain. But, to be in a predication relation with the noun being modified by the reduced relative clause, it is required to have category N. The empty node in [spec, CP] cannot participate in both relations, or an incompatible category specification will be made, as shown in Figure 22(a) (compare the successful binding relations of Figure 22(b), for the RR with *found*). Thus, there is no way to establish grammatical binding relations among the empty nodes in this extended projection. The CP-IP-VP fails to become highly activated, since its $\overline{\mathrm{X}}$ phrases have unlicensed empty nodes. Neither of the possible extended projections for the causative past participle can compete with the highly activated alternative, that of the main verb intransitive reading. The main verb IP-VP structure therefore wins the competition for activation, and the past participle structures become inactive.

It is crucial to note that the outcome for an unergative reduced relative construction in this account is far bleaker than one in which the structure just does not accumulate sufficient activation to compete effectively with the main verb structure. In fact, the parser is unable to activate the empty structure needed for a grammatical analysis of this sentence. Recall our observation that the reduced relative construction with an unergative verb requires that

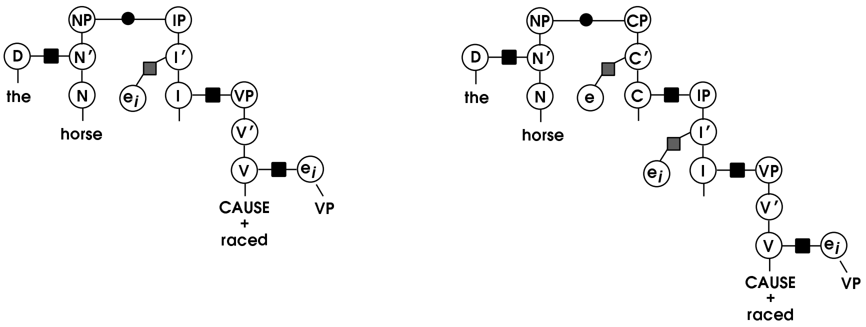

Figure 21: The two extended projections within the network corresponding to the causative past participle reading of *raced*.

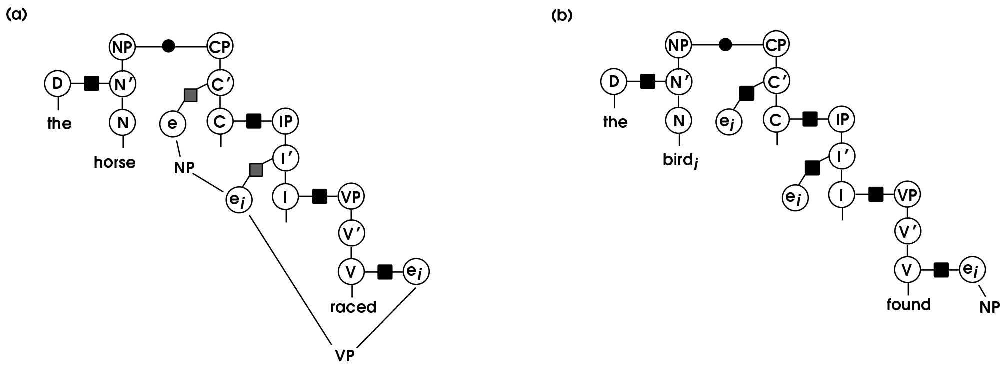

Figure 22: (a) The invalid binding relations established in the reduced relative structure for *raced*. The [spec, CP] must be an NP, and the empty complement must be a VP. To be bound by the [spec, CP] and to bind the empty complement, the empty node in [spec, IP] would have an incompatible category specification—both NP and VP. (b) The successful network representation for the reduced relative structure for *found*, repeated here for comparison.

the verb take an empty VP complement, which must have articulated structure in order to support the right binding relations. This structural requirement arises from the fact that in order to be interpreted as a transitive verb, a syntactic operation is required to combine the unergative verb with the CAUSE morpheme. It is this fact of having an explicit causative structure in the syntax that causes the processing difficulty. The competitive attachment parser has limited abilities to activate empty structure, as explained in Section 3. Here the parser can only activate a single empty node, attached to the V node for the incorporated verb, that stands in the position of the VP complement. The parser cannot activate a full projection for the empty embedded VP, which would be necessary to establish the correct empty nodes for the traces of the NP and V movements. Specifically, the parser is unable to activate an empty projection which has no overt elements or features to attach to it, which is precisely what is required in this case. The CP-IP-VP that it activates, as in Figure 22, is as close as it can come. Thus, it is not the case that the reduced relative structure merely loses the competition for activation. Rather the model predicts a severe garden path, in

Figure 23: The two best competitors of the set of extended projections within the network corresponding to the main verb reading of *melt*.

which a grammatical structure for the reduced relative interpretation is unavailable to the parser within its normal parsing operation.

4.3 Unaccusative and Other Optionally Intransitive Verbs

This account differs from theories that attribute the distinction between easy and hard MV/RR ambiguities to the fact of a verb being obligatorily transitive (like *found*) and optionally intransitive (like *raced*) (e.g., Gibson, 1991; Pritchett, 1992). This distinction is not empirically adequate, as we noted earlier. For example, unaccusative verbs, such as *melt*, are also optionally intransitive, but are easily processed in the reduced relative construction, even when they have an MV/RR ambiguity. Furthermore, the distinction cannot be simply the fact that a transitive unergative verb is causative, since an unaccusative verb is also causative in its transitive alternation. The crucial distinction between the unergatives and unaccusatives is that the unaccusatives undergo causativization in the lexicon, while the unergatives require a syntactic causative operation.

For example, consider *The witch melted in the movie deserved her fate*. Like *raced*, *melted* has both an intransitive reading and a transitive causative reading. Unlike *raced*, the two readings correspond to two different lexical entries for the verb, as shown earlier in Figure 3. When structure is activated based on the main verb entries, the parser ends up with the two competing IP-VP structures shown in Figure 23: the intransitive form in which the verb has an empty node in the verb complement position bound by the subject NP *the witch*, and the causative form in which the V node has an empty complement that needs to be bound. The latter structure is somewhat less preferred because of the unbound trace. The past participle reading is only relevant for the causative form of the verb, since the intransitive form cannot be passivized. This lexical entry therefore leads to one extended projection, the CP-IP-VP shown in Figure 24, which is exactly identical to the reduced relative structure for a simple transitive verb like *found*. This structure is also slightly less preferred than the main verb intransitive reading, because the NP *the witch* needs a predicate. Although there are incomplete relations, however, there is not the problem that arises in the unergative case, in which the parser is unable to complete the necessary structures at all, even with additional input. For each of the unaccusative interpretations, the parser can activate and attach the structures it needs when given additional compatible input.

Although we will not go into the processing details here, this same result holds for other optionally intransitive verbs as well, aside from the unergative class. "Ordinary"

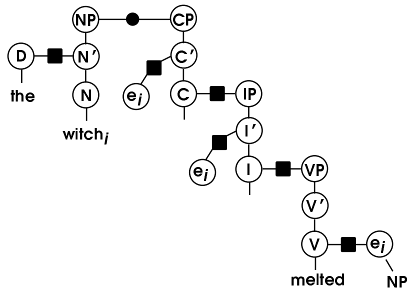

Figure 24: The best competitor of the set of extended projections within the network corresponding to the causative past participle reading of *melt*.

(non-causative) optionally intransitive verbs have the same structure for the RR as the unaccusative verbs, and are therefore also relatively easy to process in the competitive attachment parser.

4.4 Discussion and Additional Consequences

We have shown that, together, our lexical structural analysis and the processing principles of the competitive attachment model account for the severe difficulty of reduced relatives with unergative verbs, as well as for the ease of interpretation of reduced relatives with verbs of any other class that support a transitive alternative. Thus the garden path effect only arises in the very narrow class of unergative verbs with a transitive alternation. Table 3 summarizes the main steps of the processing of the three classes of verbs seen so far.

In our account, the difficulty of an unergative reduced relative arises from the fact that the parser cannot activate the structure needed to represent a grammatical analysis of a syntactically causativized verb in this construction. Every alternative account of which we are aware focuses on the MV/RR ambiguity in accounting for the garden path nature of such well-known examples as *The horse raced past the barn fell* and *The boat floated down the river sank.* Every approach relies on a basic main verb preference, and then pursues one of the following two strategies: that of making the main verb reading so much more preferred that the alternative cannot compete with it, or that of pruning the RR reading and making any reanalysis operations within the parser not applicable in this instance. Our account differs radically in asserting that the difficulty of these infamous garden paths is in fact not a result of a garden path at all, at least not in the usual sense of the term applied to this construction. That is because, in our model, the difficulty of the construction is not a result of the ambiguity.

Note that this explanation of difficulty with the MV/RR ambiguity for unergative verbs like *raced* has the following consequences. First, it predicts that even *unambiguous* reduced relative constructions will be very difficult with unergative verbs, because the structure of the embedded VP cannot be built under the restrictions of the competitive attachment model. Second, it predicts that passives of causative unergatives should also be difficult in this framework. Third, it predicts that unergative RRs should be hard in all contexts. Here we examine each of these potentially problematic issues in turn.

OBLIGATORILY TRANSITIVE: *found*					
The bird found		in the room		died	
MV	RR	MV	RR	MV	RR
VP $\underline{\text{IP-VP}}$ CP-IP-VP	VP IP-VP $\underline{\text{CP-IP-VP}}$	$\underline{\text{IP-VP + PP}}$ equally activated	$\underline{\text{CP-IP-VP + PP}}$		CP-IP-VP + PP + MV only one possible

UNERGATIVE: *raced*					
The horse raced		past the barn		fell	
MV	RR	MV	RR	MV	RR
VP $\underline{\text{IP-VP}_{intr}}$ IP-VP$_{caus}$ CP-IP-VP	VP$_{caus}$ IP-VP$_{caus}$ CP-IP-VP$_{caus}$	$\underline{\text{IP-VP}_{intr}}$		no possible analysis	

UNACCUSATIVE: *melted*					
The witch melted		in the movie		deserved her fate	
MV	RR	MV	RR	MV	RR
VP $\underline{\text{IP-VP}_{intr}}$ IP-VP$_{caus}$ CP-IP-VP	VP$_{caus}$ IP-VP$_{caus}$ $\underline{\text{CP-IP-VP}_{caus}}$	$\underline{\text{IP-VP}_{intr} + \text{PP}}$ IP-VP$_{caus}$ + PP	CP-IP-VP$_{caus}$ + PP		CP-IP-VP$_{caus}$ + PP + MV only one possible

Table 3: Summary of processing steps for different kinds of verbs. The most highly activated competitors are underlined.

First consider the prediction that even unambiguous unergative verbs will pose severe processing difficulty in the reduced relative construction. For example, when processing an unambiguous reduced relative as in *The greyhound run around the track all day was tired*, the parser ends up with the same structure for the reduced relative clause as in the case for *raced*, as shown in Figure 25. The parser also builds the main verb analysis. Agreement features are propagated by the feature passing mechanism from the NP and VP nodes to the attachment node between the NP and the I′, where it is found that the main verb structure has a person/number feature mismatch between the subject and the verb. However, the fact that there is no grammatical main verb alternative for *run* does not lessen the processing difficulty for the parser. Such an ungrammatical structure would "give way" to a valid grammatical structure, even if the latter were more complex. But there is no grammatical reduced relative alternative, so a main verb structure with a feature mismatch is better than none. This prediction is borne out by the intuitive judgments reported in Section 2, in which informants found this and other reduced relatives with an unambiguous unergative verb very difficult.

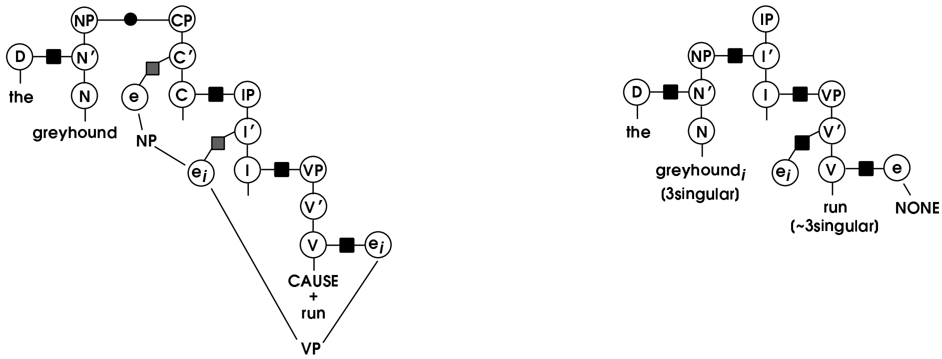

Figure 25: The competing RR and main verb structures for the unambiguous unergative past participle *run*.

Thus, this potential problem appears instead to be an advantage for the model. The same result of processing difficulty will also occur with ambiguous verbs in a disambiguating context. For example, like people, the competitive attachment model will also experience a garden path even when the RR construction occurs in a position that only licenses an NP and not a sentence, as in *Sara laughed at the horse raced past the barn*. It is not clear how other models that rely on there being an MV/RR ambiguity to explain the processing difficulty could account for these unambiguous data. Even a constraint-based account would be problematic since it would have to propose that an ungrammatical main verb structure could be weighted sufficiently more than a perfectly grammatical RR structure to allow it to consistently win the competition between the two.

Now let's look at the syntactic structure of a passive unergative, as in *The horse was raced past the barn*, shown in Figure 26. Here we see again the problematic empty articulated structure within the VP that is the complement to the syntactically incorporated verb *raced*. Turning to the processing model, Figure 27 shows the competitive attachment parser at the point of processing the passive verb. Once again we see that the parser initially activates an empty node in the complement position of the VP, whose category is set to VP by the causative morpheme. The subject *the horse* does not receive a theta role and needs to bind an empty node which does, in order to form a grammatical chain. But it cannot bind the empty VP node. To this extent, the processing of the passive is similar to the processing of the reduced relative, and our model predicts processing difficulty here as well.

However, there is a difference between the network structure here and in the reduced relative construction. In the CP-IP-VP of the RR construction, there are only weak empty nodes which need to bind something or to be bound themselves. By contrast, in the passive, there is an *overt* NP that needs to bind an empty element within the VP. We speculate that this difference may come into play in the following manner. If the parser is in a position in which there is no other possible analysis, we propose that, as a last resort, the overt NP may trigger the activation of a *bound* empty node. Because it would be bound, this empty node in turn would have sufficient activation to initiate an attachment as the specifier of an empty phrase. This empty phrase would attach as the complement of *raced*, and a grammatical passive structure would be obtained, as shown in Figure 28.

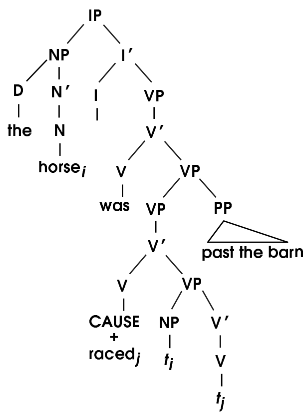

Figure 26: The syntactic structure for a passive construction with an unergative verb.

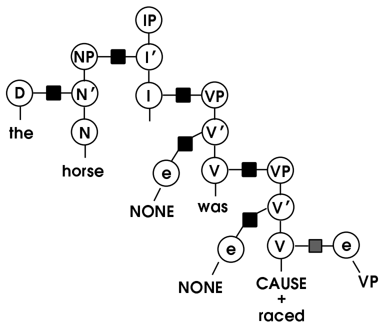

Figure 27: The parser at the point of processing the verb *raced* in the sentence *The horse was raced past the barn.*

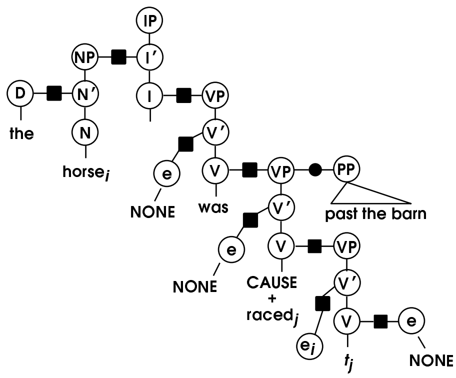

Figure 28: The final network structure for the sentence *The horse was raced past the barn.*

Clearly, however, this type of operation in which a single overt element activates multiple empty elements requires an unusual focusing of activation within the NP, triggered by its need to bind an empty element. Thus we predict some processing difficulty even with the "simple" passive of an unergative verb. While we are unaware of any experimental data that bears on this issue, we do note that at least some unergative verbs are unacceptable in the passive (Jane Grimshaw, p.c.), so we believe this prediction is worth exploring further.

Finally, we turn to a class of apparent counterexamples to our claim that unergative RRs are unparsable, exemplified by the relative ease of sentence (21) (we thank an anonymous reviewer for pointing out this example).

(21) The horse raced in the Kentucky Derby was faster than the one raced at Ascot.

This sentence poses a puzzle, given the observation that unergative RRs are generally not rescuable from a garden path fate. Typically, unergative RRs in unambiguous contexts (e.g., *The stablehand groomed the horse raced past the barn*) or with unambiguous participles (e.g., *The greyhound run around the track all day was tired*) are as difficult as the traditional garden path *The horse raced past the barn fell*, indicating that even an unambiguous context is not sufficient to allay the difficulties inherent in this construction. First, we must point out that our claim is that unergative RRs are not parsable within the normal attachment operation of the parser. It is conceivable that there are special recovery routines that enable the parser to arrive at an interpretation under certain conditions, although it is clear from the empirical data that this ability is limited. One possibility is that a special routine similar to the "last resort" procedure described above for the unergative passive can be exploited.

The question then is why sentence (21) would be amenable to this procedure while the other examples noted above would not. Since the precise nature of a such a procedure awaits further research, we can only speculate here, but it is worth making two observations regarding which properties of example (21) may facilitate recovery. First, the use of "the one" construction is compatible only with a modifying phrase which is predicated of *the one*. This provides unambiguous evidence of a reduced relative (the only possible modifier headed by a verb) for the second occurrence of *raced* in the sentence, and by syntactic parallelism, for the first occurrence as well. Second, sentence (21) contains a lengthy disambiguating region—cf. *was faster than the one raced at Ascot* to *fell*. Although this phenomenon is not entirely understood, it has been noted that a lengthy disambiguating region appears to give the sentence processor the time it needs to revise a misanalysis that it might not otherwise be able to recover from (Frazier and Clifton, 1996). In the competitive attachment parser, the extra length may provide not only extra time but additional activation to lend toward triggering the necessary empty structure. However, the details of this approach are a matter for future research. We simply note in conclusion that, to our knowledge, no account of reduced relative difficulty can explain both the apparent contextual effects of this example, as well as the verb class distinction that we have explored in this paper.

5 General Discussion

Our explanation of the difficulty of a subset of reduced relative clauses relies crucially on the interaction of processing effects and a different grammatical structure for easy and difficult cases. This proposal gives prominence to structural information stored in the lexicon, which determines the syntactic structure of verbal projections. Although it assumes that structure building is one of the causes of processing complexity, it contrasts with previous structure-based proposals which rely on only general syntactic information (e.g., Ferreira and Clifton, 1986; Frazier and Fodor, 1978; Frazier, 1978). The approach here also suggests a finer-grained lexical distinction than the simple obligatorily vs. optionally transitive contrast called on by Gibson (1991) and Pritchett (1992).

Constraint-based models put forward to explain the differential difficulty of reduced relative clauses downplay the role of structure building as a source of complexity, and instead propose that several factors are simultaneously at play in the resolution of syntactic ambiguity (MacDonald, 1994; MacDonald et al., 1994; Trueswell et al., 1994; Trueswell, 1996). In analogy to word recognition, several competitors are activated corresponding to an ambiguous word (the *-ed* form of the verb, for instance, in the MV/RR ambiguity). Linguistic competence determines the number of competitors, depending on the subcategorization frame of the verb and the syntactic constructions in which the ambiguous form of the verb can appear (for instance, active or passive). Frequency of usage of the linguistic features in the alternatives (voice, transitivity, tense, etc.) in the linguistic experience of the speaker determines the strength of the activation of each of these competitors. Probabilistic co-occurrence constraints of phrases can further increase or decrease the activation of the competitors. In these models, structural complexity comes into play indirectly, because the simpler representation usually corresponds to a global bias in the processing network, and it interacts with item-specific biases in the resolution of ambiguity. No constraint is, in principle, qualitatively different from the others. This approach then predicts that any difficulty posed by a constraint can be eliminated, if the opposing constraints are appropriately manipulated.

The current proposal shares the view of constraint-based models that several factors are at play in resolving ambiguity, and that the computation occurs in parallel and is activation based. In fact, we specify a precise constraint-integration algorithm: We provide both a precise description of the representations that enter into the competition, and a numerical formulation of the constraint-integration mechanism, which is competition based, as described above. By contrast to other approaches, however, we believe that not all constraints have the same importance. Structural complexity is given a special status by being defined implicitly in the architectural design, entailing that certain structures cannot be computed. In the current model, structural complexity alone can cause failure to interpret a sentence even when all other factors would help its correct interpretation. This phenomenon is clearly observable when the sentence is not ambiguous. In general, the model predicts categorial processing effects, which are not predicted by other constraint-based models. Proposals that do not make crucial use of computational restrictions incur empirical inadequacies in that they do not readily explain the sharp difficulty for *unambiguous* reduced relatives with unergative verbs. As noted earlier, a constraint based account, where the strength of activation of one analysis depends on the numbers and strength of competitors (MacDonald, 1994; MacDonald et al., 1994), would be problematic, since it would have to propose that an ungrammatical main verb structure could be weighted sufficiently more than a grammatical RR structure to allow it to consistently win the competition between the two.

A model that makes crucial use of architectural restrictions is a first attempt to provide an explanation for frequency differentials. In constraint-based processing, on the other hand, frequency is treated as if it were a primitive of the processing system, a predictor variable not reducible to other factors. Note that if frequency were a primitive—either in the perception or in the production system—one would expect it to vary independently of structural complexity and, in particular, independently of syntactic verb classes. One would expect, for instance, to see transitivity preference scattered across verb types. Corpus counts appear to contradict this position, and instead support the view put forth in Gibson and Pearlmutter (1994), where frequencies are taken to be an indicator of processing complexity. Processing complexity, then, has to be grounded elsewhere, to avoid circularity. We propose it to be grounded in structural complexity.[5]

Many other proposals that have discussed the resolution of MV/RR ambiguity are not equipped to explain the difference between verb classes *within* the same structural ambiguity. For example, referential theory (Crain and Steedman, 1985; Altmann and Steedman, 1988) proposes that RRs are particularly difficult because they require more presuppositions than the main verb alternative. This is due to the restrictive nature of the RR, which presupposes a set of elements in the discourse from which the RR picks one. The suggestion is that the HSPM follows a principle of referential parsimony in preferring the main verb interpretation, which carries fewer presuppositions. We simply notice that this account needs to be extended to deal with the difference between verb classes found in this paper. The same conclusion is drawn regarding "construal theory" (Frazier and Clifton, 1996; Gilboy, Sopena, Clifton, and Frazier, 1995), which proposes an approach for the interpretation of adjuncts. Construal theory determines the time-course of attachment of primary and non-primary (adjunct) relations. It assumes a delay in the attachment of adjuncts, which are "associated" with the current theta domain and interpreted using extragrammatical principles. The MV/RR ambiguity is a case of contrast between primary and non-primary relations, as the main verb role for the VP is primary, while a relative clause is not primary. A strong main verb bias in MV/RR ambiguities is then predicted, as the parser always tries to satisfy a primary relation if it can. No difference between different kinds of verbs would be predicted, as in all cases the primary main verb relation would be preferred.

The current proposal highlights the role of the computation of syntactic structure, as projected from the lexicon. As the syntactic difference is derived from lexical structure, it accounts for differences across *classes* of lexical items and is in contrast with those approaches that want to reduce or eliminate the syntactic level in processing. Within the framework of constraint-based processing, Spivey-Knowlton and Sedivy (1995) propose that PP attachment preferences rely on referential and verb-specific constraints, without any recourse to syntactic knowledge. Crucially, their study shows verb subclass effects (verbs of action vs. psych/perception verbs), with preference for VP attachment for action verbs, and NP attachment for psych and perception verbs. They argue that lexical preferences account for the pattern, but it is never explained *why* action verbs, for instance, should favor an instrument interpretation over a manner interpretation of the ambiguously attached PP.

[5]It has been suggested to us by an anonymous reviewer that differential frequencies come from the production system—i.e., people tend to avoid syntactically complex alternatives. The processing system, being a perception system, is sensitive to such differences. If such a preference exists in production, though, that shows that the linguistic processing faculties are sensitive to structural complexity in general. Simplicity of the theory would lead us to prefer an explanation where the language faculty were sensitive to the same factors both in production and in perception.

More importantly, the conclusion of Spivey-Knowlton and Sedivy that syntactic distinctions are irrelevant to the pattern of attachment preferences seems too hasty, since verb classes are defined in terms of the content of theta roles. Although some researchers believe that theta roles constitute semantic information, there are well attested syntactic corollaries of each different thematically defined subclass of verbs (unergatives, unaccusatives, perception verbs, psych verbs, etc.). In the absence of a precise formulation of the processing model and of the representation of linguistic knowledge, the evidence in Spivey-Knowlton and Sedivy (1995) is inconclusive between a semantic and syntactic account.

Although we propose that theta roles partially determine syntactic structure, they are not directly manipulated in our model. This position contrasts with the proposal put forth in Carlson and Tanenhaus (1988), where thematic roles are directly accessed by the processor, which uses them to integrate the incoming input incrementally into a discourse model. While our account does not address word sense ambiguity and discourse integration, an account in terms of the content of theta roles does not properly explain the difference in difficulty between unergative and unaccusative RRs. The crucial property proposed by Carlson and Tanenhaus to account for difficult RRs is that the role which is assigned to the subject of the intransitive is the same as the role of the object of the transitive. So, in both *the child hurried home* and *the father hurried the child home*, the argument *the child* receives a Theme theta role. When the parser reaches the disambiguating portion of the sentence (the word *fell* in *the child hurried home fell*), the parser realizes that it has incorrectly selected the intransitive theta grid, but it cannot reassign the theta role since the Theme theta role has already been assigned. The difficulty of these sentences contrasts with the ease of reanalysis of a sentence such as *The girl sent the flowers was pleased*, where *the girl* is at first assigned an Agent theta role, which is then easily modified to Goal, a role that is still available at the point of disambiguation.

The Carlson and Tanenhaus (1988) account thus predicts equal difficulty for unergatives and unaccusatives, as in both cases the thematic role that alternates between the subject of the intransitive and the object of the transitive is the same. Thus, the prediction is that unaccusatives should also be hard. Indeed Carlson and Tanenhaus list *the butter melted* as a difficult reduced relative. Our intuitions, supported by those of our informants, suggest that the difference in difficulty between the two subclasses of verbs is substantial. Our account of the difficulty of unergatives rests on the assumption that the object of the transitive variant is not a Theme, but rather an Agent. It might seem odd to have two Agents in the same clause. In order to avoid this situation, it could be proposed that transitive unergative verbs assign two different roles, such as (Agent, Actor). But under this representation, unergatives would be easier than unaccusatives in the Carlson and Tanenhaus (1988) account, as a new theta role could be assigned to the subject of the RR at the point of disambiguation, namely the Actor role. Our account achieves a better explanation of the data by mapping the difference in thematic role assignment in the two verb subclasses onto a structural difference, which in turn causes parsing difficulty in some cases. In this way we are able to account for fine distinctions in the data without positing any parsing mechanism specific to the construction under study.

More recent discussions of the semantic and thematic constraints that affect the ease of resolution of the ambiguity under study are found in Tabossi, Spivey-Knowlton, McRae, and Tanenhaus (1994) and Trueswell et al. (1994). Studies have shown that animacy does not explain the difficulty in garden path RRs (Ferreira and Clifton, 1986; Trueswell et al., 1994). The intuitive judgments reported here confirm the experimental results. Animacy of the modified NP is not necessary to make the sentence hard, since *The model planet rotated on the metal axis fell off the stand* is hard, although *the model planet* is clearly inanimate. Neither is animacy a sufficient condition for processing difficulty, since we found that *The witch melted in the Wizard of Oz was played by a famous actress* was easily understood.

In an attempt to find a better semantic predictor of processing difficulty of a reduced relative clause, Trueswell et al. (1994) suggest using the notion of "degree of semantic fit." The "degree of semantic fit" is a measure of the distance of a nominal lexical item from the prototypical recipient of a given theta role, in the context of a given verb. So, for instance, a *reporter* is a good Agent for a verb like *interview* but a poor or unlikely Patient. In the examples presented here, however, the "degree of semantic fit" of the initial NP with the main verb reading does not precisely correlate with the difficulty of the RR interpretation, as proposed by Trueswell et al. (1994). Trueswell et al. claim that if the mapping between thematic role and grammatical functions permits a main clause interpretation, then a garden path effect will be found. But we observe that this incorrectly predicts a garden path effect for unaccusative RRs, because the initial NP *necessarily* exhibits the same "degree of semantic fit" with either a main verb or RR reading, since the theta role of Theme is equally suited to be an object or a subject in these verbs. Furthermore, although unergatives follow this same pattern of semantic fit, they behave very differently from unaccusatives.[6] In general, we can remark that the subclasses of verbs under study here differ from other classes because both the subject of the active voice and the subject of the passive voice can be equally well an Agent, in the case of unergatives, and a Theme, in the case of unaccusatives. The degree of semantic fit is not informative to resolve the MV/RR ambiguity for these classes of verbs. The manipulation of semantic content of the arguments of these verbs then will have to be studied differently.

In conclusion, we have proposed a lexical structural account for the differential difficulty of reduced relative clauses. We make the novel observation that extreme garden paths occur most of the time with manner of motion/sound emission verbs, a subclass of intransitive verbs with peculiar properties. We have proposed a lexical and syntactic analysis of these verbs that, in conjunction with a specific model of processing, produces the right description of the data. The processing model that is supported by the empirical evidence is the competitive attachment model (Stevenson, 1994b). This approach accounts for the sharp distinctions between unergative RR clauses and RR clauses with other verbs. The analysis achieves a better fit to the empirical data than those proposals which build the differences in parsing difficulty on the verb being optionally or obligatorily transitive (Gibson, 1991; Pritchett, 1992). Both unaccusatives and unergatives are preferentially intransitive verbs, but the difference in processing difficulty is remarkable.

[6]Frequency effects are also not sufficiently strong to explain the difference. On the one hand, we find that unaccusative verbs have a strong preference to be used intransitively overall, thus favoring an association of the theta role Theme with the grammatical function subject and making the wrong prediction if degree of semantic fit were the predominant factor. On the other hand, the reduced relative construction is equally rare for both classes of verbs, entailing that the relative frequency of main verb and past participle use can only be a weak factor in determining the preferred analysis of RRs with these verb classes. See Merlo and Stevenson (1996) for the relevant corpus analysis and discussion.

The proposal achieves its account thanks to an independently justified linguistic analysis and independently justified properties of the processing model. In particular, the parser experiences difficulty in processing reduced relative clauses of unergative verbs because of its limited ability to project empty nodes and to bind them in the structure. Our explanation does not rely on the fact that the RR participle is an ambiguous verb form, and thus it accounts for the observation that complexity effects persist even when there is no ambiguity, a fact not easily explained by other models. For other verbs, such as obligatorily transitive verbs or unaccusatives, where the lexical and syntactic competitors in the analysis are more evenly weighted, there is no processing breakdown, and the model can be influenced by pragmatic and semantic factors. The proposal thus reconciles results that point toward an architecture where several constraints are simultaneously at play, with results that indicate sharp and irrecoverable breakdown.

References

Altmann, G. and M. Steedman (1988). Interaction with context during human sentence processing. *Cognition 30*(3), 191–238.

Baker, M. (1988). *Incorporation.* Chicago, IL: University of Chicago Press.

Baković, E. (1995). Optimality and inversion in spanish. In *Proceedings of the Optimality in Syntax Workshop.* MIT Working Papers in Linguistics and MIT Press. To appear.

Belletti, A. and L. Rizzi (1988). Psych verbs and θ theory. *Natural Language and Linguistic Theory 6*(3), 291–352.

Bever, T. G. (1970). the cognitive basis for linguistic structure. In J. R. Hayes (Ed.), *Cognition and the Development of Language.* New York: John Wiley.

Bresnan, J. (1982). The passive in lexical theory. In J. Bresnan (Ed.), *The Mental Representation of Grammatical Relations*, pp. 3–86. Cambridge, MA: MIT Press.

Brousseau, A.-M. and E. Ritter (1991). A non-unified analysis of agentive verbs. In *West Coast Conference on Formal Linguistics*, Number 20, pp. 53–64.

Burzio, L. (1981). *Intransitive Verbs and Italian Auxiliaries.* Ph. D. thesis, MIT, Cambridge, MA.

Burzio, L. (1986). *Italian Syntax.* Dordrecht: Kluwer Academic Publishers.

Carlson, G. N. and M. K. Tanenhaus (1988). Thematic roles and language comprehension. In W. Wilkins (Ed.), *Syntax and Semantics*, Volume 21, pp. 263–288. San Diego: Academic Press.

Chomsky, N. (1981). *Lectures on Government and Binding: The Pisa Lectures.* Dordrecht: Foris Publications.

Crain, S. and M. Steedman (1985). On not being led up the garden path: The use of context by the psychological parser. In D. R. Dowty, L. Kartunnen, and A. M. Zwicky (Eds.), *Natural Language Processing: Psychological, Computational, and Theoretical Perspectives*, pp. 320–358. Cambridge: Cambridge University Press.

Cruse, D. A. (1972). A note on english causatives. *Linguistic Inquiry 3*(4), 520–528.

Dowty, D. R. (1979). *Word Meaning and Montague Grammar.* Dordrecht: Reidel.

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

Ford, M., J. Bresnan, and R. Kaplan (1982). A competence-based theory of syntactic closure. In J. Bresnan (Ed.), *The Mental Representation of Grammatical Relations*, pp. 727–796. Cambridge, MA: MIT Press.

Frazier, L. (1978). *On Comprehending Sentences: Syntactic Parsing Strategies.* Ph. D. thesis, University of Connecticut. Available through the Indiana University Linguistics Club, Bloomington, IN.

Frazier, L. and C. Clifton (1996). *Construal.* Cambridge, MA: MIT Press.

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

Gibson, E. (1991). *A Computational Theory of Human Linguistic Processing: Memory Limitations and Processing Breakdown.* Ph. D. thesis, Carnegie-Mellon University.

Gibson, E. and N. J. Pearlmutter (1994). A corpus-based analysis of psycholinguistic constraints on prepositional phrase attachment. In C. Clifton, L. Frazier, and K. Rayner (Eds.), *Perspectives on Sentence Processing*, pp. 181–198. Laurence Erlbaum.

Gilboy, E., J.-M. Sopena, C. Clifton, and L. Frazier (1995). Argument structure and association preferences in spanish and english complex nps. *Cognition 54*, 131–167.

Grimshaw, J. (1991). Extended projection. Brandeis University. Waltham, MA.

Grimshaw, J. (1996). Projection, heads and optimality. *Linguistic Inquiry*. To appear.

Hale, K. and J. Keyser (1987). A view from the middle. Lexicon Project Working Papers 10.

Hale, K. and J. Keyser (1993). On argument structure and the lexical representation of syntactic relations. In K. Hale and J. Keyser (Eds.), *The View from Building 20*, pp. 53–110. MIT Press.

Harley, H. (1995). *Subjects, Events and Licensing.* Ph. D. thesis, MIT, Cambridge,MA.

Jackendoff, R. (1972). *Semantic Interpretation in Generative Grammar.* Cambridge,MA: MIT Press.

Kaplan, R. M. and J. Bresnan (1982). Lexical-functional grammar: A formal system for grammatical representation. In J. Bresnan (Ed.), *The Mental Representation of Grammatical Relations*, pp. 173–281. Cambridge, MA: MIT Press.

Koopman, H. and D. Sportiche (1991). The position of subjects. *Lingua 85*, 211–258.

Levin, B. (1993). *English Verb Classes and Alternations.* Chicago, IL: University of Chicago Press.

Levin, B. and M. Rappaport (1986). The formation of adjectival passives. *Linguistic Inquiry 17(4)*, 623–661.

Levin, B. and M. Rappaport Hovav (1995). *Unaccusativity.* Cambridge, MA: MIT Press.

MacDonald, M., N. Pearlmutter, and M. Seidenberg (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review 101*(4), 676–703.

MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes 9*(2), 157–201.

Merlo, P. and S. Stevenson (1996). Integrating statistical and structural information in a distributed architecture for syntactic disambiguation. Technical Report TR-25, Rutgers Center for Cognitive Science, Rutgers University.

Pritchett, B. (1988). Garden path phenomena and the grammatical basis of language processing. *Language 64*(3), 539–576.

Pritchett, B. (1992). *Grammatical Competence and Parsing Performance.* Chicago, IL: University of Chicago Press.

Reggia, J. (1987). Properties of a competition-based activation mechanism in neuromimetic network models. In *Proceedings of the First International Conference on Neural Networks*, pp. II–131–II–138.

Rizzi, L. (1990). *Relativized Minimality.* Cambridge: MIT Press.

Siloni, T. (1994). *Noun Phrases and Nominalizations.* Ph. D. thesis, University of Geneva.

Spivey-Knowlton, M. and J. C. Sedivy (1995). Resolving attachment ambiguity with multiple constraints. *Cognition 55*(3), 227–267.

Spivey-Knowlton, M. J., J. C. Trueswell, and M. K. Tanenhaus (1993). Context effects in syntactic ambiguity resolution: Discourse and semantic influences in parsing reduced relative clauses. *Canadian Journal of Experimental Psychology 47*(2), 276–309.

Stevenson, S. (1993a). A competition-based explanation of syntactic attachment preferences and garden path phenomena. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, pp. 266–273.

Stevenson, S. (1993b). Establishing long-distance dependencies in a hybrid network model of human parsing. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society*, pp. 982–987.

Stevenson, S. (1994a). Competition and recency in a hybrid network model of syntactic disambiguation. *Journal of Psycholinguistic Research 23*(4), 295–322.

Stevenson, S. (1994b). *A Competitive Attachment Model for Resolving Syntactic Ambiguities in Natural Language Parsing.* Ph. D. thesis, Department of Computer Science, University of Maryland. Available as technical report TR-18 from Rutgers Center for Cognitive Science, Rutgers University.

Stevenson, S. (1995). Arguments and adjuncts: A computational explanation of asymmetries in attachment preferences. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, pp. 748–753.

Sturt, P. and M. Crocker (1996). Monotonic syntactic processing: A cross-linguistic study of attachment and reanalysis. *Language and Cognitive Processes.* To appear.

Tabossi, P., M. J. Spivey-Knowlton, K. McRae, and M. K. Tanenhaus (1994). Semantic effects on syntactic ambiguity resolution: Evidence for a constraint resolution process. In C. Umita and M. Moscovitch (Eds.), *Attention and Performance XV*, pp. 589–615. Hillsdale,NJ: Laurence Erlbaum Associates.

Trueswell, J. C. (1996). The role of lexical frequency in syntactic ambiguity resolution. *Journal of Memory and Language 35*, 566–585.

Trueswell, J. C., M. K. Tanenhaus, and S. M. Garnsey (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language 33*, 285–318.