In M. Crocker, M. Pickering and C. Clifton, Eds., Architectures and Mechanisms for Language Processing, 161-188. [Draft version.]

Lexical Syntax and Parsing Architecture

Paola Merlo and Suzanne Stevenson

DRAFT: DO NOT QUOTE OR CIRCULATE

Paola Merlo
Département de Linguistique
Faculté des Lettres
Université de Genève
2 rue de Candolle
CH-1204 Genève
Switzerland
merlo@latl.unige.ch

Suzanne Stevenson
Department of Computer Science,
and Center for Cognitive Science
Rutgers University
New Brunswick, NJ 08903
suzanne@ruccs.rutgers.edu

We thank Charles Clifton, Sven Dickinson, Uli Frauenfelder, Jane Grimshaw, Maryellen MacDonald, Luigi Rizzi, Mike Tanenhaus, John Trueswell, and anonymous reviewers for helpful comments and discussion on earlier drafts of this manuscript. Thanks also are due our informants at Rutgers Center for Cognitive Science (RuCCS) and at the University of

Geneva for volunteering their time. This work was begun under the generous financial and administrative support of RuCCS. It was completed while the first author was a visiting scientist to the Institute for Cognitive Science at the University of Pennsylvania, on Swiss National Science Foundation fellowship 8210-46569, and while the second author was supported by the (US) National Science Foundation, under grant IRI-9702331. The support of these institutions is gratefully acknowledged.

1 Introduction

Traditionally, accounts of cognitive abilities have proposed inherent restrictions on cognitive architecture as the underlying cause of observed patterns of human behaviour. Models of the human sentence processing mechanism (HSPM) have largely fallen within this architectural paradigm, attributing interpretation preferences and processing breakdown to limitations on computational resources or on the interface between cognitive modules (e.g., Crocker, 1992; Frazier and Fodor, 1978; Gibson, 1991; Lewis, 1993; Marcus, 1980). Recently, in conjunction with advances in connectionism, research in cognitive modelling has seen a surge of interest in the role of environmental factors in determining crucial properties of intelligent behaviour. In contrast to a highly constrained underlying architecture, relatively unrestricted cognitive mechanisms are proposed, whose precise behaviour is determined by their adaptation to specific experiences. Within sentence processing research, this approach is exemplified by the constraint-based paradigm (e.g., MacDonald, 1994; MacDonald, Pearlmutter, and Seidenberg, 1994; Spivey-Knowlton, Trueswell, and Tanenhaus, 1993; Trueswell, 1996; Trueswell, Tanenhaus, and Garnsey, 1994; but see also Mitchell, Cuetos, Corley, and Brysbaert, 1995), which proposes that exposure to the linguistic environment, rather than internal architectural properties, is primarily responsible for the specific behaviours of the HSPM. In constraint-based models, individual lexical entries play a central role as the repository of experiential knowledge, and lexical frequencies are a primary determinant of human behaviour. In particular, the operation and outcome of the HSPM are determined almost entirely by the differential frequencies among alternative interpretations.

In contrast to this recent trend, we have argued that restrictive architectural assumptions are necessary to an explanatory account of the HSPM (Stevenson and Merlo, 1997). Our work has focussed on reduced relative clauses, investigating the sharp contrast in acceptability between the following sentences:

(1) The horse raced past the barn fell. [Unacceptable]

(2) The butter melted in the microwave was lumpy. [Easily interpreted]

Until the final verb of each sentence, which is the actual main verb, the first verb (*raced*, *melted*) can be interpreted as either a past tense main verb or as a past participle within a reduced relative construction (e.g., *the horse* [that was] *raced past the barn*). These sentences are thus said to contain a main verb/reduced relative (MV/RR) ambiguity. In general, the main verb interpretation is preferred, but the reduced relative interpretation is necessary for a coherent analysis of the complete sentence in both (1) and (2). In (1), the strength of preference for the main verb interpretation of *raced* leads to breakdown when the verb *fell* is processed. By contrast, although sentence (2) contains a different verb (*melted*) with exactly the same type of MV/RR ambiguity, it is easily understood. The striking effect of individual verbs on the processing of this ambiguity has been argued to support a constraint-based account in which verb-specific frequency information is largely responsible for the variability of processing difficulty in this construction (e.g., MacDonald, Pearlmutter, and Seidenberg, 1994; Trueswell, 1996).

We find the lack of restrictiveness in the frequency-based accounts to be unsatisfactory, since the underlying cause of the differential frequencies themselves remains unexplained. We instead pursue the following two-pronged approach to explaining the differential difficulty in processing these types of sentences. First, we observe that the ease or difficulty of processing a reduced relative clause is correlated with the lexical semantic class of the verb within the construction (e.g., *raced* or *melted* above). Based on lexical and syntactic properties, we propose that verbs of the unergative class (such as *race*) have a more complex transitive structure than other optionally transitive verbs (such as *melt*). The analysis is supported by acceptability judgments which reveal a bi-modal distribution for reduced relative constructions, with the unergative verb class causing processing difficulty, and other optionally transitive verbs yielding easy processing. Second, we present the architectural properties of the competitive attachment parser—a computational model of the HSPM—that are relevant to the analysis of reduced relative constructions. 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. Through this demonstration, we hope to convey the continued importance of investigating restrictive architectural assumptions in order to progress beyond descriptions of human behaviour, to deeper insights into the mechanisms that underlie language processing.

2 Verb Classes and Lexical Syntax

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*, both contain a manner of motion verb (*raced*, *floated*) in its transitive usage. Manner of motion verbs have an intransitive/transitive alternation, in which the intransitive form is classified as unergative (Levin, 1993; Levin and Rappaport Hovav, 1995) and the transitive form is the causative counterpart of the intransitive. We tested our intuitions about the difficulty of unergatives in the reduced relative construction by asking naive informants for acceptability judgments on sentences with reduced relative clauses containing verbs from this class, exemplified in (3).

(3a) The clipper sailed to Portugal carried a crew of eight.
(3b) The troops marched across the fields all day resented the general.
(3c) The greyhound run around the track all day was tired.

Most informants find these sentences very difficult. One conjecture is that the difficulty of unergative verbs in the reduced relative construction derives from the fact that the sentence has a complex causative semantics. As a comparison group, we chose the unaccusative verbs, another well-defined verb class that displays a similar intransitive/transitive alternation related by causativisation (Levin, 1993). Examples are given in (4). Informants found these sentences mostly easy.

(4a) The witch melted in the Wizard of Oz was played by a famous actress.
(4b) The genes mutated in the experiment were used in a vaccine.
(4c) The window broken in the soccer game was soon repaired.

Clearly, judgments on the difficulty of reduced relatives containing unergative versus unaccusative verbs show a patterning of responses along verb class lines: unergative verbs are difficult, while unaccusatives are easy.[1] Looking at the individual verbs, we can summarize

[1]A complete list of the sentences used is reported in Stevenson and Merlo (1997).

the results as follows. In the unergative group,

Advance, *glide*, *march*, *rotate*(I)[2], *sail*, *run* were judged unacceptable.

Shrink, *fly*(I), *withdraw*, *walk* were judged mostly unacceptable.

In the unaccusative group,

Melt, *mutate*, *pour*, *rotate*(N), *begin*, *freeze* were judged slightly degraded.

Reach, *break*, *grow*, *fly*(N) were judged perfect.

First, all of the (mostly or completely) unacceptable verbs belong to the unergative group, while all of the easy verbs (i.e., perfect or only slightly degraded) belong to the unaccusative group.[3] 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*).

The remainder of this section develops our analysis of the lexical structure of unergative and unaccusative verbs, and the ensuing syntactic consequences of the proposed structures.

2.1 Lexical Structure

Sentences (5) and (6) exemplify the manner of motion and sound emission verbs that form the subclass of unergative verbs that occur in an intransitive/transitive alternation.

[2]Two of the verbs, *rotate* and *fly*, were used in two different reduced relatives in which the modified noun differed according to its semantic properties of motion. An I designation means that the verb occurred with an argument with inherent motion (e.g., a bird); an N means that it occurred with an argument with non-inherent motion (e.g., a kite). Classification of verbs as unergative or unaccusative depends on the lexical semantics of the verb and its arguments taken together. Manner of motion verbs with an object of inherent motion are classified as unergative, and with an object of non-inherent motion are classified as unaccusative (Levin, 1993).

[3]Another difference seems to characterize easy and difficult garden paths, namely that easily processed verbs are those that have an adjectival passive form (Bresnan, 1982; Levin and Rappaport, 1986). However, when applying the standard tests for adjectival passives, this distinction did not correctly characterize the differential processing difficulty found here.

(5a) The horse raced past the barn.

(5b) The jockey raced the horse past the barn.

(6a) The lion jumped.

(6b) The trainer jumped the lion through the hoop.

As noted above, the transitive (b) form of each sentence is the causative counterpart of the intransitive (a) form (see Brousseau and Ritter (1991) and Hale and Keyser (1993), but also cf. Levin and Rappaport Hovav (1995)). That is, the syntactic subject of the intransitive form becomes the object of the transitive form, and the transitive subject is the causer of the action.[4] Unaccusative verbs can similarly occur in an intransitive/transitive alternation related by causativisation, as shown in sentences (7) and (8).

(7a) Butter melted in the pan.

(7b) The cook melted the butter in the pan.

(8a) Rice grows in Northern Italy.

(8b) Farmers grow rice in Northern Italy.

However, despite the superficial similarity between the causative alternation in (5)/(6) compared to (7)/(8), there are several crucial differences between unergatives and unaccusatives, which will lead us to postulate a different lexical, and consequently syntactic, structure.

First, the subject of an intransitive unergative, such as *the horse* in sentence (5a), is an Agent, while the subject of an intransitive unaccusative, such as *the butter* in sentence (7a), is a Theme (or Patient). This theta role difference has interesting consequences for the causative alternation of these verbs. Conceptually, causativization involves the insertion of a Causal Agent into the existing argument structure. The causativization of an unaccusative, such as *melt* in sentence (7b), yields a transitive form whose subject is a Causal Agent (*the cook*) and whose object is a Theme (*the butter*). By contrast, the causativization of an unergative, such as *race* in sentence (5b), yields a transitive form whose subject is again the

[4]Sound emission verbs, such as *ring* and *buzz*, show the same linguistic and processing properties of manner of motion verbs. Although we do not discuss them directly, the analysis we develop below applies also to the sound emission verb class.

Causal Agent (*the jockey*), but whose object in this case is itself an Agent as well (*the horse*). That is, although the subject NP *the jockey* is the agent of the causing action, 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 lion through the hoop*, in which the lion is clearly the agent of *jump*.

The transitive form of an unergative, then, is unusual in that it has an argument structure (Causal Agent, Agent), which has two agentive theta roles.[5] The peculiarly agentive nature of the object of manner of motion verbs has been remarked before (Cruse, 1972; Hale and Keyser, 1987), and 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); compare the unergatives in (9) (from Levin and Rappaport Hovav, 1995) with the unaccusatives in (10).

(9a) *The downpour marched the soldiers to the tents.

(9b) *The firecracker jumped the horse over the fence.

(10a) The ice storm froze the pipes.

(10b) 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).

The causative alternation of unergatives and unaccusatives also presents differences in productivity, within the same language and cross-linguistically. Several languages (such as Italian or French) do not permit unergative verbs to transitivize, and those that do, such as English, appear to do so in a limited way. For instance, in English only a subset of unergative verbs can be used transitively. If we consider the previous observation that the object of a transitive manner of motion verb is peculiarly agentive, we can consider this limited productivity as a special case of a more general restriction on the mapping between thematic relations and grammatical functions. Belletti and Rizzi (1988) argue that the Agent

[5]Several accounts describe the object of the verb as a Theme (Levin and Rappaport Hovav, 1995; Harley, 1995). We think that this 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.

theta role seems to have a special status in the mapping of theta roles to syntax. They observe that psych verbs in Italian belong to three different syntactic subclasses, which have all the mappings of the pair (Experiencer, Theme) onto syntactic configurations. Consider the examples below: the Experiencer can be the subject (a), the object (b) or the indirect object (c), and the Theme can be the subject (b,c) or the object (a).

(11a) Gianni teme i serpenti. *Gianni fears snakes*

(11b) L'esame preoccupa Gianni. *The examination worries Gianni*

(11c) A Gianni piace il gelato. *To Gianni pleases ice-cream*

However, there is no verb that has a similar alternation between (Agent, Theme). If there is an Agent, it must be the theta role that is mapped onto the external argument position, i.e. the subject.

Finally, there is a clear morphological difference in the formation of the causative form in those languages that allow both classes of verbs to causativize morphologically, such as Hebrew. As observed in Levin and Rappaport Hovav (1995), unaccusatives show morphological marking of the intransitive member of the pair; moreover, this pattern of causativization is productive and unmarked. On the other hand, unergative causatives show that the transitive form is morphologically derived from the intransitive, which is itself underived. Thus, Levin and Rappaport Hovav conclude that unergative verbs are formed by a different causativization process than the one that applies to unaccusatives.

We can provide a unified explanation of these linguistic facts, by extending Hale and Keyser's (1993) approach to lexical syntax and thematic roles. The main claim of Hale and Keyser's work is that thematic roles do not exist as an independent linguistic entity, but are in fact simply notational shorthands for some structural relations between lexical categories at the lexical level. By reducing thematic roles to the cross-product of lexical categories and structural configurations, they explain why the number and type of thematic roles is so limited and so uniform across languages.

Before presenting the technical details of our proposal, we might want to consider the theoretical import of a purely formal approach, given that semantic factors are clearly relevant to the classification of verbs as unergative or unaccusative (cf. fn.2 regarding the ambiguous verbs *fly* and *rotate*). We follow the view that a linguistic theory that does not make reference to thematic roles directly is more restricted. For example, the lexical structural approach of Hale and Keyser (1993) 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 syntax-semantics 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.

On this view, the apparent sensitivity of certain grammatical processes to thematic content, although empirically correct, is epiphenomenal. Moreover, a syntactic account is more general, as it is supported by a vast body of studies showing that the lexical semantics of unaccusative and unergative verbs entails different syntactic behaviour, 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).

Thus, we follow Hale and Keyser (1993) in assuming that each thematic role corresponds to a given structural configuration within a lexical semantic structure. Consequently, the difference in thematic role assignment noted above 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. The combination operation is one of syntactic head movement within the lexical structure; see Figure 1, top panel.

For example, an intransitive unaccusative verb such as *melt*, which expresses a change of state, consists of an adjectival head *melted* that incorporates into the abstract verbal head BECOME. On the other hand, an intransitive unergative verb such as *race*, which

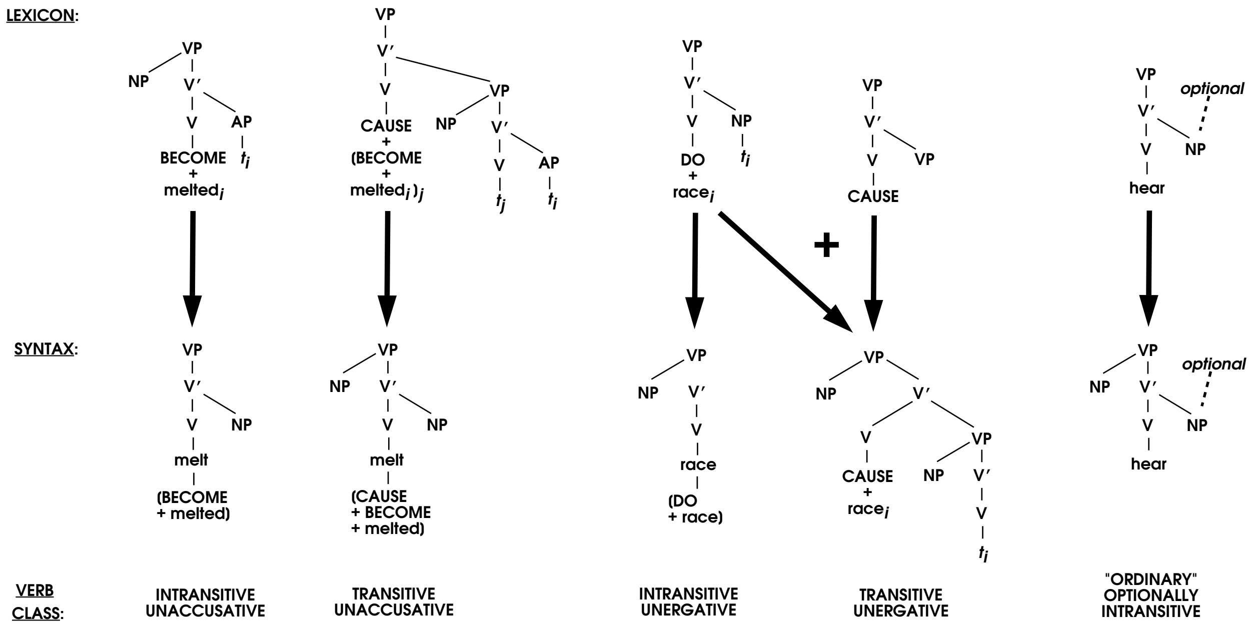

Figure 1: Lexical entries, showing verb incorporation, for various verb types, and their corresponding syntactic structures.

expresses an activity, is generated by the lexical incorporation of a noun *race* into the abstract verbal head DO. Theta roles are assigned in these structures as follows. The specifier position 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 the agent of the activity cannot be in a predication relation with the NP in complement position, as an NP is not a predicate. The agent of the *race* activity must be inserted at the syntactic level, where it is related to the verb by predication; it is therefore not mapped from the lexicon, but is instead projected by a syntactic rule in virtue of the Extended Projection Principle (Chomsky, 1981). It is a true external argument in the sense that it is external to the lexical structure.

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 within the lexicon: because the argument structure of the verb is complete, the entire VP can act as an argument to the CAUSE morpheme. 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 complement of another verb. Figure 1 shows the lexical derivation of the causative form of *melt*. By contrast, an unergative verb like *race* cannot be causativized in the lexicon, because not all its arguments have been assigned. Since the *agent* argument is missing from the lexical structure, the VP for *race* cannot be further embedded as a complement to another verbal head. 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 (5) and (6).

It is therefore necessary to extend the account of Hale and Keyser (1993) to deal with transitive unergatives. We propose 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. 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 a Causal Agent is inserted into its specifier position. The explicit combination of the two verbs in the syntax is responsible for producing the two [spec, VP] positions that are necessary for the two external (agentive) theta roles.[6] As can be seen in Figure 1, bottom panel, the syntactic operation of causativization entails that the syntactic representation of the transitive form of unergatives is more complex than a simple VP.

For verbs that do undergo incorporation within the lexicon, we assume 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.[7] Specifically, a verb that results from lexical incorporation—a sequence of head movements within a nested predicate argument structure—appears in the syntax in a conflated form, within a simple VP. Figure 1 illustrates this relation between the underlying (articulated) lexical structures, and the resulting syntactic structures of various verbs.

[6] The specifier position of an XP will be referred to throughout the paper as "[spec, XP]".

[7] Here again our proposal deviates from that of Hale and Keyser, 1993. The topological approach to the mapping of thematic roles 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.

This linguistic analysis, which posits lexical and syntactic differences between unergative and unaccusative verbs, constitutes the first part of our explanation of the differential difficulty of reduced relative clauses, by entailing that the transitive form of an unergative verb is structurally more complex than that of an unaccusative verb. In order to complete our account, we must exhibit a processing model whose crucial properties are justified on purely computational grounds, and whose interaction with the proposed linguistic representation gives rise to the observed behaviour in processing reduced relative clauses. We do so in the next sections.

3 The Competitive Attachment Architecture

The competitive attachment model is a hybrid connectionist parser that integrates simple symbolic processing with numeric spreading activation, within a network of processing nodes that directly represents syntactic structures. The model dynamically builds a parse tree by activating processing nodes in response to the input, whose attachments within the tree structure are incrementally established through distributed network computation. Symbolic feature processing determines the grammaticality of potential phrase structure attachments, based on Government-Binding theory (Chomsky, 1981; Rizzi, 1990). Numeric spreading activation weighs the relative strengths of the valid attachment alternatives and focuses activation onto the preferred structures. Competition for activation is the crucial mechanism for resolving ambiguity.

The architecture 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 define a precise range of structural hypotheses that can be represented within the network. In particular, the constraints on the parsing architecture generally limit the occurrence of empty nodes, but also, perhaps surprisingly, dictate the presence of empty structure in certain configurations. Because empty structure (empty operators and traces) plays an important role in the syntactic analysis of reduced relative clauses, the constraints on empty structure in the model are of central concern.

Here we will first present the basic network creation and numeric competition processes

Figure 2: An $\overline{\mathrm{X}}$ template and its instantiation as a verb phrase.

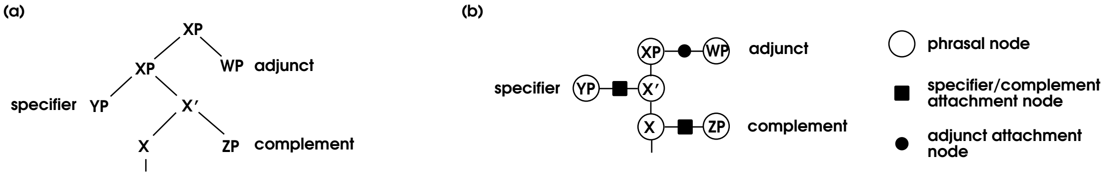

Figure 3: (a) The structure of an $\overline{\mathrm{X}}$ phrase and (b) its representation in the competitive attachment network.

within the parser, then turn to the precise architectural restrictions regarding empty nodes and phrases. We conclude the section with a description of the syntactic projection of verb classes within the model.

3.1 Overview of the Competitive Attachment Process

The network of alternative parse tree structures is dynamically created by activating two types of processing nodes: one corresponding to the phrasal nodes of a standard parse tree representation, and one corresponding to the links or attachments between the phrases. Phrasal nodes are activated from a pool of $\overline{\mathrm{X}}$ templates, and their symbolic features are initialized based on an input token's lexical entry. Figure 2 shows a sample $\overline{\mathrm{X}}$ template and its instantiation. Attachment nodes are established between two phrasal nodes that are potential *sisters* in the parse tree; Figure 3 shows the resulting representation of basic $\overline{\mathrm{X}}$ relations in the network.

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 strength of the corresponding lexical entry. After activating the appropriate phrasal nodes, the parser allocates the attachment nodes to represent the potential attachments between the new phrasal nodes and the existing network.

Figure 4: (a) An example parse tree and (b) its network representation.

Then, each processing node iteratively updates and outputs its symbolic features and numeric activation, until the network settles into a stable state representing the attachment decisions for each of the new phrases. Figure 4 shows the final parse tree for a complete sentence, and the corresponding network structure.

Two types of competitions occur during the network processing described above. An *attachment competition* is a competition at a phrasal node for the best attachment to activate. Since multiple attachments for a phrasal node may be grammatical, a competitive output function (Reggia, 1987) weighs the evidence for an attachment, and proportionately allocates activation from the phrasal node to its attachment nodes. When the network stabilizes, the set of active attachment nodes for each syntactic phrase represents the current parse tree attachments for that phrase; inactive attachment nodes connected to that phrase are deleted. A *lexical competition* arises when an input token has more than one possible syntactic projection. The activation level of each possible phrase is determined in part by the level of activation of its attachments, and in part by the activation it receives from the input token. As in an attachment competition, the input token uses a competitive output function that proportionately divides activation among the recipient nodes—in this case, the possible phrases for the input. This lexical competition determines which of the multiple possible phrases will be most highly activated by the input token.

3.2 Architectural Restrictions and Empty Structure

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. For example, an I′ node must activate a single attachment to a subject phrase, and a V node that represents an obligatorily transitive verb must activate a single attachment to a complement phrase. However, the competitive attachment mechanism is quite limited in its ability to activate a precise number of attachment nodes. Specifically, 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. This ability appears necessary even in the most straightforward cases. For example, in processing a simple sentence with an obligatorily transitive verb, the V node must initially activate zero attachment nodes, since its complement phrase will not have been reached yet in the input. Once the complement phrase is activated, the V node will then need to activate a single attachment node to it. But while the competitive output function can be tuned to effectively activate exactly one attachment, it cannot be adjusted to allow the activation of zero attachment nodes at one point in processing, and one attachment node at another.

Since the competitive function can bring about activation of *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, such a phrasal node automatically posits an attachment to an *empty* node. Each phrasal node that licenses an attachment thus has a "dummy" node to attach to even when no overt phrase is available. Furthermore, to maintain uniformity of numeric processing across all phrasal nodes, *every* phrasal node activates an attachment to an empty node. In the case that an attachment is prohibited at a node, or simply not required, an empty node is interpreted in the network as an explicit representation of the lack of an attachment. Figure 5 shows the empty nodes that are activated with an $\overline{\mathrm{X}}$ 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. For example, in the reduced relative construction, empty nodes are needed to represent both the trace of passive movement, and the empty operator in the specifier of the relative clause. The role of an empty node as an empty category is determined within the

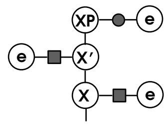

Figure 5: Empty nodes are activated in conjunction with every phrase, at each attachment site.

network by the attachment and binding relations in which the node participates. Note, however, that the activation of empty structure is quite restricted, in that each empty node is a single, "dummy" phrasal node. A empty node cannot have an articulated phrasal structure, which could in turn activate additional empty structure, recursively. This restriction is necessary to curb the potential explosion of nested empty structure, which could quickly exhaust the available processing nodes in the parser. Thus, two fundamental aspects of the competitive attachment architecture—the numeric competition process and the dynamic activation of structure—lead to the necessity of allocating empty structure, which is restricted to being a single processing node.

The status of an empty node as a default phrase to attach to entails that it initially has only minimal activation. If the empty node occurs at a site that does not require an attachment (e.g., as the complement to an intransitive verb), then it represents a valid non-attachment, and quickly becomes fully activated. If the empty node attaches to a site that does require an attachment (e.g., as the specifier of an I′, or the complement of an obligatorily transitive verb), then there are several possible outcomes. Often the default attachment to an empty node is later replaced by an attachment to an overt phrase that is subsequently projected from the input; in this case, the empty node is fully deactivated. However, as noted above, an empty node may participate in the parse tree structure by representing an empty category. The coindexation of an empty node licenses the node as an empty category and instantiates its symbolic features, thereby increasing its activation level. Conversely, if an empty node occurs in a position that requires an attachment and it is *not* coindexed, it will remain only weakly activated.

3.3 The Projection of Lexical Structure

The bottom-up nature of the competitive attachment model entails that syntactic structure is projected from a lexical item, where projection can require the activation of potential extended projections. The activation of a number of extended projections for a verb reflects the potential "stacking" of functional projections over the primary lexical projection of a word (Grimshaw, 1991; Grimshaw, 1997). A verb, for example, will activate three extended projections: a VP, an IP-VP, and a CP-IP-VP. 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 linguistic account presented in Section 2 entails that the internal structure of the VP of the extended projection encodes features from the predicate argument structure of the verb; that is, the VP of each extended projection has the specifier and complement requirements determined by the lexical features of the (possibly incorporated) verb. Initially, the specifier and complement sites will have attachments to empty nodes, as described earlier. The lexical information associated with the verb will put constraints on the interpretation of the empty nodes.

For a conflated verb, such as *melt*, the structure of the VP is based on the position of the verb following incorporation within the lexical structure. The form of *melt* created from the incorporation of BECOME and *melted* projects a VP that has a single internal argument and no external argument. The NP that occurs in the [spec,VP] of the abstract element BECOME is realized in the syntax as a direct object of the conflated 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 node attached to the V$'$ has its category set to NONE. The form of *melt* created from incorporation of CAUSE, BECOME and *melted* also projects a VP with a single internal argument. Again the internal NP argument is realized in the syntax as a direct object of the conflated 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.

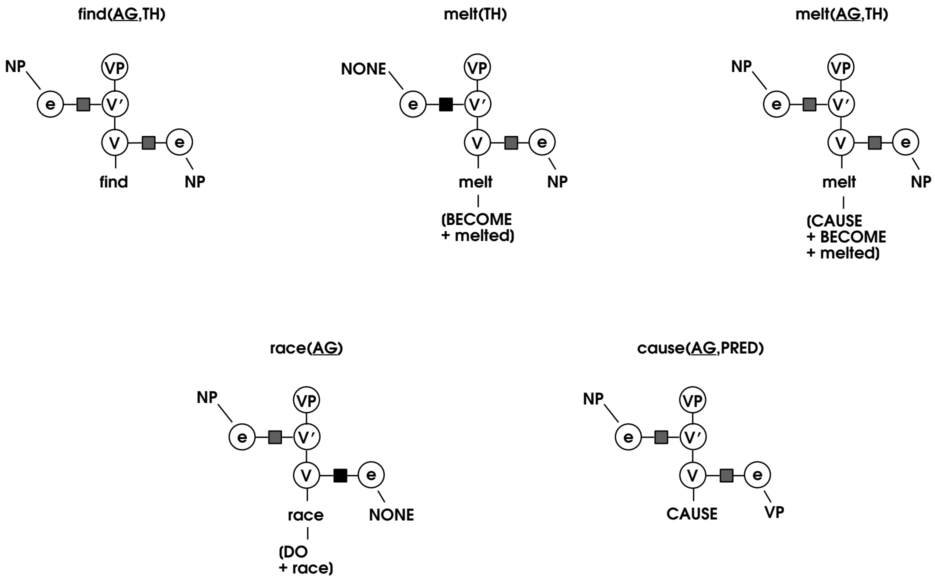

Figure 6: The VP structures initially activated by the parser in response to various types of verbs.

Unergative verbs are distinguished from unaccusative verbs in that the lexical rule of causativization cannot apply, as described in Section 2. The only lexical form available for a verb such as *race* is the intransitive. An unergative verb therefore projects a VP 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 6, 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.[8] The processing consequences of syntactic causativization will be discussed in the next section.

4 Verb Classes and Processing Results

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. It is here we will demonstrate the effects of the architectural restrictions regarding empty nodes and

[8]We assume that CAUSE is linked in the lexicon to the unergative verbs with which it can so combine, since the ability to be causativized, even under the general syntactic operation, is an idiosyncratic property of verbs.

empty phrases in the model. Recall that these constraints are motivated by a convergence of restrictive computational assumptions. The results therefore arise from the interaction of our lexical syntactic proposal with independently justified properties of our parsing architecture.

In what follows, we assume that a verb with a main verb/past participle ambiguity has three lexical entries that are homophonous: the simple past tense, the active past participle, and the passive past participle. Since the active past participle will never be grammatical in the examples considered here, we will only be concerned with the simple past and the passive participle alternatives. 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), and whether the RR is assumed to be a CP, IP, VP, or even a DP is irrelevant for the results here.) To simplify the presentation, only highly activated extended projections will be shown in the figures.

4.0.1 Obligatorily Transitive Verbs

First consider a reduced relative construction with an obligatorily transitive verb, as in *The bird found in the room died*, whose syntactic structure is shown in Figure 7. At the point that *found* is processed, the main verb alternative will highly activate an IP-VP projection, to which the NP *the bird* attaches in subject position. The passive past participle reading will highly activate a CP-IP-VP projection, which has an empty node in [spec, CP] that can participate in a predication relation with the noun *bird*. The empty node in [spec, CP] can also bind the empty node in [spec, IP], which in turn can bind the empty node that is complement to the V node. This establishes the empty node in [spec, CP] as an empty operator, leading to a grammatical reduced relative construction.

Figure 8 illustrates these 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; (2) the CP-IP-VP projection corresponding to the passive past participle lexical entry of *found*, which modifies the NP *the bird*. Neither alternative emerges as a clear winner. The first has an

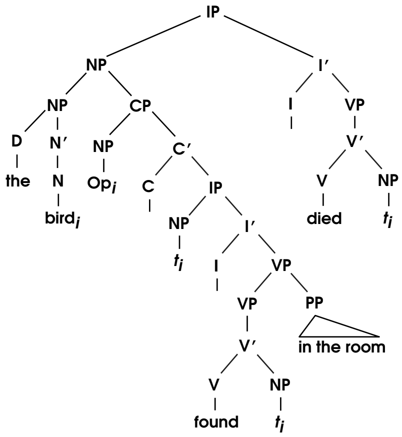

Figure 7: The syntactic structure for a reduced relative construction, in *The bird found in the room died.*

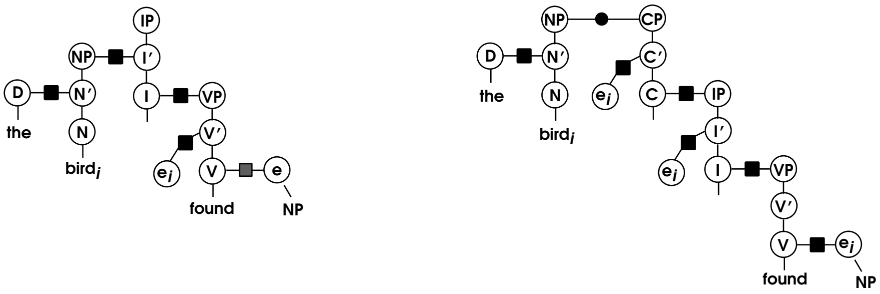

Figure 8: The two best competitors activated in response to the input *the bird found.*

empty node in a required complement position, which cannot become fully activated. The second is waiting on a (presumably required) predicate. (Currently in the model, this expectation for a predicate is implemented by favoring root nodes that are extended projections of a verb.) The closeness of the competition allows context and theta role compatibility to influence preferences throughout the processing of the post-verbal PP (MacDonald, 1994; Trueswell, Tanenhaus, and Garnsey, 1994). But 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 9.

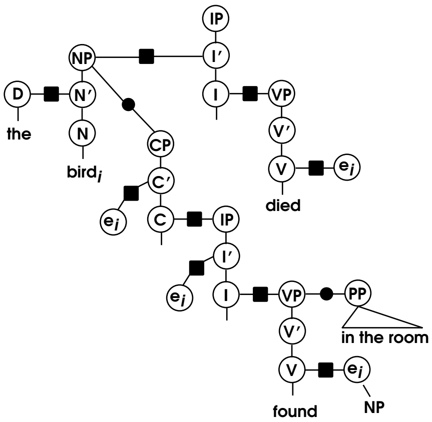

Figure 9: The final network configuration for *The bird found in the room died.*

4.0.2 Unergative Verbs

The ease of a reduced relative construction with an obligatorily transitive verb has received a great deal of attention because of the stark constrast with the prototypical garden path sentence, *The horse raced past the barn fell.* Figure 10 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—in terms of both number of nodes and number of binding relations—than that of a reduced relative with a simple transitive verb such as *found* (compare Figure 10 with Figure 7). The passive participle *raced* requires a VP complement with two traces. One trace within the complement VP, t_i, is the 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]. The other trace within the complement VP, t_j, 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. 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 11. In contrast to the main verb reading of *found*, the empty node

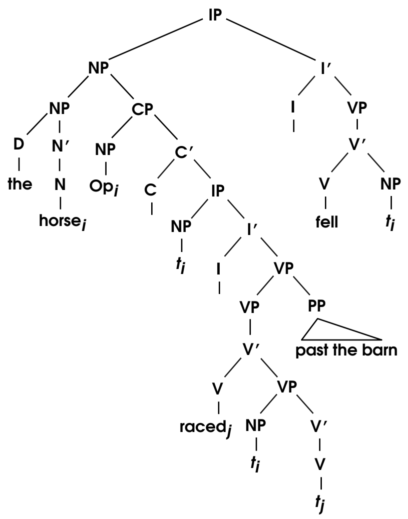

Figure 10: The syntactic structure for *The horse raced past the barn fell.*

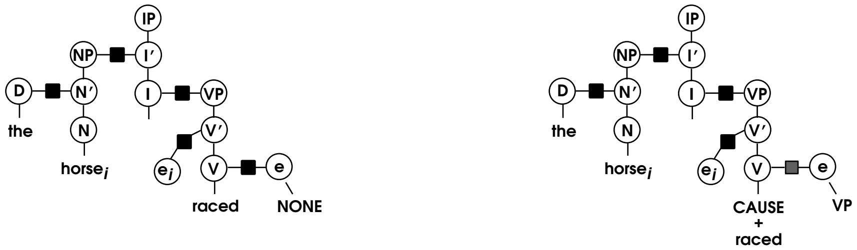

Figure 11: The two best competitors of the set of extended projections within the network corresponding to the main verb reading of *raced.*

that is sister to V in the intransitive form of *raced* is specified to be empty, so it becomes fully activated. 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.

For the past participle reading, the most highly activated extended projection is the CP-IP-VP of the causative form, shown in Figure 12. This structure has empty nodes in the same positions as in the reduced relative form of *found*, but the result is very different with the verb *raced*. In the network for *found* (see Figure 9), the empty node that is sister to the verb can represent the trace of passive movement, thereby successfully participating in a chain headed by the empty node in [spec, CP], leading to a valid reduced relative structure. In the network for raced, however, the empty node that is sister to the verb stands in the

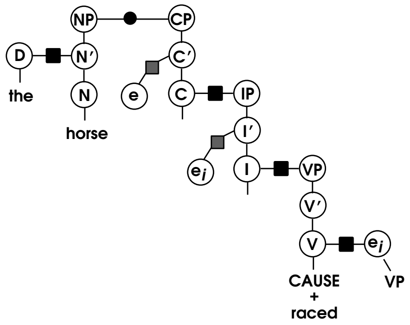

Figure 12: The best competitor of the set of extended projections within the network corresponding to the causative past participle reading of *raced.*

place of a VP in the syntactic structure—one that has articulated empty structure. Since the network is only able to allocate a single empty node in the position in which an articulated empty phrase occurs in the syntactic structure, the network is unable to encode the binding relations that are necessary to successfully represent a reduced relative clause.

Specifically, in Figure 12, the empty node in [spec, IP] binds the complement of the verb, but this establishes a VP chain.[9] 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 13. Thus, there is no way to establish grammatical binding relations among the empty nodes in this extended projection. The CP-IP-VP decreases markedly in activation, since its $\overline{\mathrm{X}}$ phrases contain empty nodes that cannot become fully activated; it is therefore unable to 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.

Although we have described the outcome for the reduced relative structure in this example in terms of a competition with the main verb reading, it is crucial to note that the prognosis for an unergative reduced relative in this account is far bleaker than this may imply. In fact, the parser is simply unable to activate the structure needed for a grammatical analysis of this

[9]Alternatively, the empty node in [spec, IP] may not be able to participate in a VP chain; in either case, a valid chain including the [spec, CP], [spec, IP], and complement of V cannot be formed.

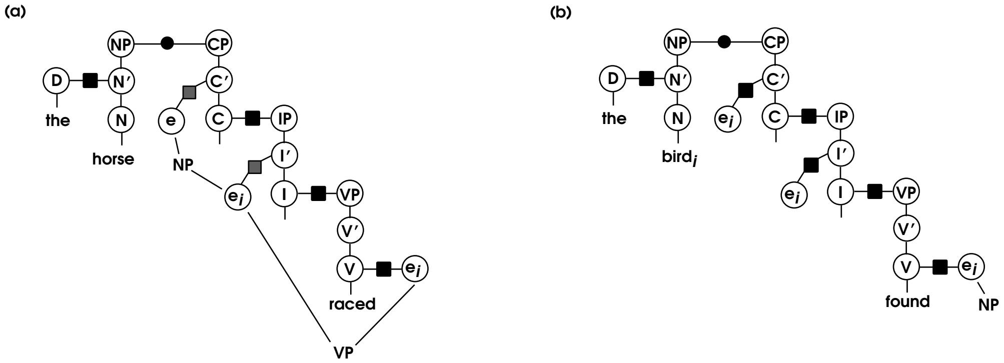

Figure 13: The invalid binding relations established in the reduced relative structure for *raced*. The [spec, CP] must be an NP, and the empty complement must be a VP. To be bound by the [spec, CP] and to bind the empty complement, the empty node in [spec, IP] would have an incompatible category specification—both NP and VP.

sentence. Recall our observation that the reduced relative construction with an unergative verb requires that 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 architecture of the competitive attachment parser limits its 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 has independently motivated constraints on the activation of empty structure that render it 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 13, 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 which a grammatical structure for the reduced relative interpretation is completely unavailable within the normal operation of the parser.

4.0.3 Unaccusative Verbs

This account differs from theories that attribute the difference between easy and hard MV/RR ambiguities to the distinction being verbs that are obligatorily transitive (like *found*) versus optionally intransitive (like *raced*) (e.g., Gibson, 1991; Pritchett, 1992). Such a distinction is not empirically adequate. 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. 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 causative transitive reading. Unlike *raced*, the two readings correspond to two different lexical entries for the verb, as shown earlier in Figure 1. When structure is activated based on the main verb entries, the parser ends up with the two competing IP-VP structures shown in Figure 14: (1) the intransitive form in which the verb has an empty node in its complement position bound by the subject NP *the witch* (the object of an intransitive unaccusative moves to subject position in the syntax since the verb has no external argument); and (2) 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 leads to one extended projection, the CP-IP-VP shown in Figure 15, 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.

Figure 14: The two best competitors of the set of extended projections within the network corresponding to the main verb reading of *melt*.

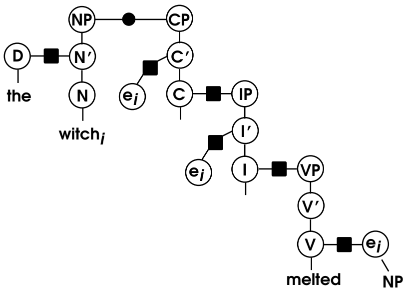

Figure 15: The best competitor of the set of extended projections within the network corresponding to the causative past participle reading of *melt*.

4.1 Discussion

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 simple transitive verbs and unaccusative verbs. A garden path effect only arises in the very narrow class of unergative verbs that have a transitive alternation. Note that this explanation of difficulty in processing the MV/RR ambiguity with 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. That is, the difficulty does not stem from the ambiguity but rather from the inability of the parser to activate the correct reduced relative structure for an unergative verb. This is compatible with the observation that difficulty persists even when the reduced relative occurs in a fully disambiguating context, such as in object position (cf. *I groomed the horse* *raced past the barn*). The account also conforms with the pattern of acceptability judgments presented in Section 2, which reveal that unergative verbs in a reduced relative construction cause processing problems even when the past participle is not ambiguous. Second, our explanation predicts that the passive of a causative unergative should also disrupt normal processing. Since the grammatical structure for a passive unergative includes an empty embedded VP, just as in the reduced relative, the parser should experience the same difficulty in its processing. While the model appears overly strong in this prediction (cf. sentence (7) below), we note that passive unergatives are often awkward (cf. sentence (8)).

(7) The schooner was sailed through the narrow channel with difficulty.

(8) ?? The lions were jumped through the hoop after the elephant act.

The evidence indicates some peculiarity with unergatives in the passive, and encourages future research into the range of difficulty that is exhibited.

5 Conclusion

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 path effects in this construction appear to be limited to manner of motion/sound emission verbs, a subclass of intransitive verbs with peculiar properties. We have proposed a lexical and syntactic analysis of these unergative verbs that in conjunction with a specific model of processing—the competitive attachment model (Stevenson, 1994)—produces the right description of the data. In our account, the linguistic representation is mapped onto an isomorphic network structure built by the parser, in which structure building is restricted. Our explanation of the difficulty of unergative reduced relatives relies crucially on the interaction of computationally justified architectural features of the processor and specific lexical representations, supported by independent principles of linguistic theory.

This proposal gives prominence to structural information stored in the lexicon, which determines the syntactic structure of verbal projections. In this way it is able to make finer grained distinctions than either structure-based proposals which rely on only general syntactic information (e.g.,Ferreira and Clifton, 1986; Frazier and Fodor, 1978; Frazier, 1978), or those accounts based on the distinction between obligatorily and optionally transitive verbs (Gibson, 1991; Pritchett, 1992). 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 (Crain and Steedman, 1985; Altmann and Steedman, 1988; Gilboy, Sopena, Clifton, and Frazier, 1995).

Unlike many constraint-based approaches—which often ascribe subclass phenomena to the interaction of frequency and pragmatic/semantic effects (Spivey-Knowlton and Sedivy, 1995)—our proposal is not purely lexicalist, but exploits structural information which is shared by several lexical items forming a subclass. Moreover, many constraint-based proposals downplay the role of structure building as a source of complexity, and instead capitalize on the differential frequencies of lexical items as the main predictor of processing load (MacDonald, 1994; Trueswell, 1996). 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.

By contrast, we specify a constraint integration algorithm—defining what lexical constraints come into play and how they interact—where not all constraints have the same importance. Structural complexity is given a special status by being defined implicitly by the architectural design: 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.

This approach accounts for the sharp distinctions between unergative reduced relatives and reduced relatives with other verbs: the parser cannot process reduced relative clauses of unergative verbs due to its limited ability to project empty nodes and to bind them in the structure. For other verbs, such as obligatorily transitive verbs, 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 frequency, 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.

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

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.

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.

Crocker, M. W. (1992). *A Logical Model of Competence and Performance in the Human Sentence Processor.* Ph. D. thesis, University of Edinburgh.

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

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

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 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.

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. Manuscript, Brandeis University. Waltham, MA.

Grimshaw, J. (1997). Projection, heads and optimality. *Linguistic Inquiry.*

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. 53–109. Cambridge, MA: 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.

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.

Lewis, R. (1993). *An Architecturally-based Theory of Human Sentence Comprehension.* Ph. D. thesis, Carnegie Mellon.

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.

Marcus, M. (1980). *A Theory of Syntactic Recognition for Natural Language.* Cambridge, MA: MIT Press.

Mitchell, D. C., F. Cuetos, M. M. B. Corley, and M. Brysbaert (1995). Exposure-based models of human parsing: Evidence for the use of coarse-grained (non-lexical) statistical records. *J. of Psycholinguistic Research 24*, 469–488.

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 J. of Experimental Psychology 47*(2), 276–309.

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

Stevenson, S. and P. Merlo (1997). Lexical structure and parsing complexity. *Language and Cognitive Processes 12*(2/3), 349–399.

Trueswell, J. C. (1996). The role of lexical frequency in syntactic ambiguity resolution. *J. 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. *J. of Memory and Language 33*, 285–318.