Title
A Unified Model of Preference and Recovery Mechanisms in Human Parsing

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 16(0)

Author
Stevenson, Suzanne

Publication Date
1994

Peer reviewed

A Unified Model of Preference and Recovery Mechanisms in Human Parsing

Suzanne Stevenson
Department of Computer Science
University of Toronto
6 King's College Road
Toronto, Ontario M5S 1A4
Canada
suzanne@cs.toronto.edu

Abstract

Models of human parsing typically focus on explaining syntactic preferences and garden-path phenomena. This paper explores another aspect of the processing of syntactic ambiguity—the successful revision of previously preferred structure. In the competitive attachment model of parsing, a hybrid connectionist network directly represents the attachment structure among phrasal nodes in a parse tree. A syntactic ambiguity leads to a network of alternative attachments that compete for numeric activation. The winning attachments are determined within a parallel operation that simultaneously revises earlier attachments as needed when initially attaching a new phrase to the developing parse tree. Because of the unique parallel structuring operation, the competitive attachment model provides a unified explanation of human preference and recovery mechanisms in parsing. The paper demonstrates this ability by showing how the model accounts for recency effects in human syntactic processing. In the parsing network, a mechanism of decay, which is independently needed to manage the finite pool of processing nodes, allows more recent phrases to compete more effectively than less recent phrases for new attachments. The effect of decay on the attachment competition underlies a unified account of psycholinguistic observations of recency, both in initial syntactic preferences and in the revision of erroneous attachments.

Introduction

In developing a cognitive model of parsing, a key problem is to elucidate the computational mechanisms underlying the human ability to process and resolve syntactic ambiguities. In order to achieve this, we must understand how the human parser determines which of the possible analyses of a syntactic ambiguity to pursue, and how it responds to later linguistic input that is incompatible with its initially preferred analysis. In general, it is assumed that the human parser adopts a preferred structural interpretation because it is easier or faster for the parser to build. To explain the observed preference behavior, sophisticated processing proposals spell out a set of computational assumptions under which a parsing mechanism will exhibit structural choices matching those of the human parser (e.g., Frazier, 1978; Gorrell, 1987; McRoy & Hirst, 1990). Although people appear to locally optimize structuring operations, this strategy does not always lead them to the best interpretation in the long run; hence, parsing models must also account for the behavior of the human parser when the initially preferred structure turns out to be incorrect. A number of recent models of human parsing have focused on explicating the processing limitations that cause the breakdown of the parser in certain well-known cases in which its local preferences lead it down the garden path (Gibson, 1991; Pritchett, 1992; Weinberg, 1991).

Much progress has thus been made in addressing two long-standing goals in the cognitive modeling of human parsing: to explain people's initial structural preferences when processing syntactic ambiguities, and to account for the breakdown of parsing in garden-path sentences. However, little attention has been given to another aspect of the processing of an ambiguity—the situation that arises when the initially preferred structure turns out to be wrong, but the parser is *not* garden-pathed. The detailed behavior of the parser when successfully revising an initially preferred structure has yet to be explained. Parsing models have till now accounted only for the division of sentences into two discrete classes: those allowing for necessary revisions within some constrained class of parsing operations, and those for which a necessary revision would require an unconstrained recovery mechanism (i.e., garden-path examples) (Abney, 1989; Gibson, 1991; Pritchett, 1992; Weinberg, 1991). These theories cannot account for the range of behavior that the parser exhibits when revising a structural hypothesis. In particular, the human parser has more or less difficulty in adopting a less preferred interpretation of an ambiguity. The ease of doing so is affected by how quickly following the onset of the ambiguity the parser receives disambiguating information (Frazier & Rayner, 1982), or even uncertain but biasing information (MacDonald, in press). It is therefore simplistic to assume that revision is either impossible, as in garden-path sentences, or completely unproblematic.

This paper discusses a competitive attachment model of parsing that provides an explicit computational account of the complex range of human behavior exhibited in the processing of syntactic ambiguities. In the model, ambiguity resolution is formulated as a competition for activation among a set of syntactic attachments within a hybrid connectionist network. Underlying properties of the model explain a number of human structural preferences and account for the status of garden-path sentences, without the use of explicit preference heuristics or revision strategies (Stevenson, 1993a; 1993b; 1994).

This paper will demonstrate how these same properties determine the ease with which the parser is able to successfully revise its attachment decisions. In contrast to other models, the competitive attachment approach yields finer-grained predictions of relative difficulty within the class of possible reanalyses. In fact, the same factors that give rise to the initial structural preferences are instrumental in determining the relative ease of reanalyzing the resulting structures. The model therefore provides a unifying account of the preference and recovery mechanisms involved in human parsing.

Because of length constraints, this paper focuses on how the model explains recency effects in human parsing. Two related results are presented. First, a necessary mechanism for decay of phrasal activation is shown to account for a general preference for recent attachments in parsing, without the use of an explicit structural strategy or preference heuristic as in other models. The single mechanism of decay of activation addresses as well the potential memory load problem that arises in considering multiple attachment sites for a phrase (see, for example, Gibson (1991)). Second, the very same computational factors that give rise to the exhibited recency preferences are shown to account for the effect of length of an ambiguity on the ease of reanalyzing incorrect structure—an observation that has heretofore lacked an adequate explanation. Thus, the paper demonstrates an example of how fundamental properties of the competitive attachment model underlie both structural preferences and ease of recovery, and account for a wider range of psycholinguistic data than has been previously explained.

The paper is organized as follows. The next section presents a brief overview of the model. The following two sections present the preference and recovery results, comparing the competitive attachment account to other proposals. The final section concludes the paper.

Overview of the Model

The competitive attachment model is a hybrid connectionist network of processing nodes that represent syntactic phrases and their attachments within a parse tree.[1] The grammaticality of potential attachments is determined using simple symbolic features and marker passing. Numeric activation encodes the the relative strengths of the attachment alternatives, and a competition mechanism provides a uniform method for determining the preference for one attachment structure over another. The network of alternative attachments is dynamically created by allocating processing nodes in response to the input.

Creating the Network

When an input token is read, the parser activates two types of processing nodes: phrasal nodes representing the current input phrase, and attachment nodes representing the potential attachments between the current input phrase and the existing parse tree. Each attachment node connects to exactly two phrasal nodes that are potential *sisters* in the parse tree; compare the network of Figure 1(a) with the subtree that it represents in Figure 1(b). Once the current phrase is connected to the existing network, each processing node iteratively updates and outputs its symbolic features and numeric activation, until the activation level of each attachment node is either above a certain threshold θ, or is zero.[2] When the network stabilizes, the set of active attachment nodes represents the current parse tree attachments. At this point, the next input token is read and the process is repeated. At the end of a successful parse, the phrases and their attachments form a connected network representing the final parse tree structure.

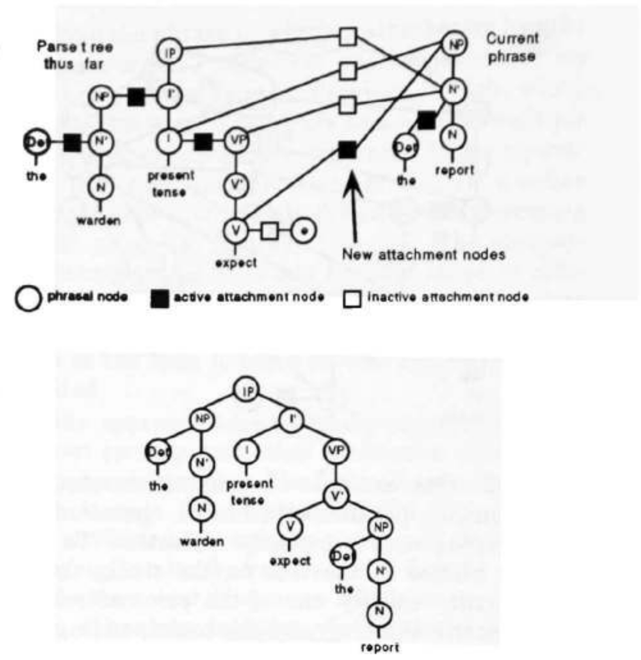

Figure 1: (a) Network and (b) parse tree representation of *The warden expects the report*.

The Competition Mechanism

A syntactic ambiguity leads to a network of grammatical attachment alternatives that cannot all simultaneously participate in the parse tree structure. In order to focus activation in the network onto a consistent set of parse tree attachments, each phrasal node uses a competitive output function (Reggia, 1987) to allocate a percentage of its activation to each of its attachment nodes, proportional to their activation levels. The output function is of the form: $o_{ji} = \frac{a_j w_{ji}}{\sum_k a_k w_{ki}} \cdot a_i$, where o_{ji} is the output from node n_i to node n_j, a_i is the activation of node n_i, w_{ji} is the weight on the connection from node n_i to node n_j,[3] and k ranges over all nodes connected to node n_i. This function brings about indirect inhibition among the attachment nodes, so that only a subset of the nodes is

[1]The parser is implemented in Common Lisp, serially simulating the parallel processing of the network.

[2]The parsing network requires 10–70 iterations to reach a stable state.

[3]Weights encode lexical preference or frequency information.

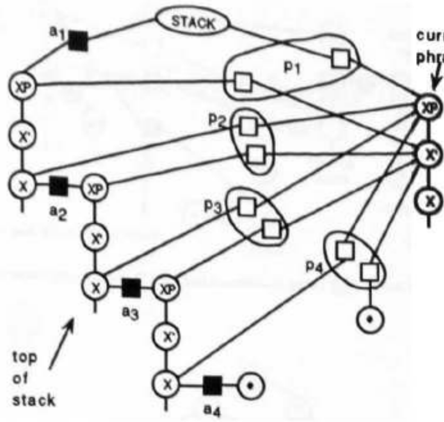

Figure 2: One example of competitive relationships in the primitive parallel attachment operation. Attachments a_1–a_4 were previously activated. To attach the current phrase to the tree on the stack, the following must occur: exactly one of the prior attachments, a_i, must become inactive, and the corresponding pair of attachments, p_i, must become active. This relationship holds for a tree of arbitrary depth on the stack.

able to become fully activated. The competitive function is designed so that in a stable state of the network, each phrasal node activates a number of attachment nodes in accordance with its grammatical properties. For example, since every phrase must have a parent in the parse tree, all XP nodes must activate exactly one attachment node.

Restricting the Network Structure

The competition mechanism is unable to rule out certain mutually exclusive pairs of attachments whose simultaneous activation would create an illegal tree structure. In order to prevent invalid parse trees, the structure of the network of potential attachments must be restricted. New attachment nodes can only represent an attachment between the current phrase and the right edge of the developing parse tree, as in Figure 1. Maintaining the required network structure entails the use of a parsing stack.[4] The only attachment nodes that can compete simultaneously are those in the set of attachments between the current phrase and the tree on top of the stack. The competitive relationships among the allowed attachment nodes completely define the sets of attachment nodes that can be simultaneously active in a stable state of the network. The logical attachment possibilities, an example of which is shown in Figure 2, follow directly from the propagation of local competitions among the attachment nodes.[5] The competitive mechanism of the parser thus determines a circumscribed set of attachment possibilities for both initial and revised attachments in the parser. The result is a unique parallel attachment operation, in which revision of earlier attachments is determined simultaneously with initial attachment of the current phrase. This parallel structuring operation leads the model to a unified explanation of preference and recovery mechanisms in human parsing.

Memory Management

An important effect on the outcome of the attachment operation is the memory management scheme employed in the parser. Since the pool of network processing nodes is finite, and there is no *a priori* limit on the length of a sentence, a mechanism must allow for nodes to be re-used during the parse. Because relatively few attachment nodes are required to represent the parse tree, most attachment nodes that are activated become inactive at the end of each processing loop. The existing competition mechanism thus ensures that most attachment nodes are available for re-use with the very next input phrase. By contrast, phrasal nodes require an additional memory management mechanism. The activation function for a phrasal node incorporates a decay factor that causes the node to decay slowly over time. In the simple case, this decay factor is a function of the difference between the current time and the time at which the phrasal node was activated. This ensures that the least recent portion of the parse tree will become inactive and available for re-use first. However, when an older phrasal node participates in a new attachment, its activation level needs to reflect that more recent participation in the parse tree. Thus the decay factor is actually a function of the difference between the current time and the time of the most recent attachment for a phrase. The degree of decay in the current implementation of the parser is linear over time; however, the only necessary property of a decay function is that it keep a reasonable portion of the developing parse tree active. That is, decay must be fast enough to enable re-use of the nodes when needed, and slow enough to maintain an amount of tree structure that is sufficient to effectively parse most inputs.

The following two sections will demonstrate how the decay of phrasal activation affects the outcome of the competitive attachment operation, accounting for recency effects in attachment preferences and in reanalysis with a single, independently motivated mechanism.

Preference Effects

The human parser shows a strong tendency to attach the current input phrase to more recent syntactic structure. This preference has been labeled Right Association (Kimball, 1973), Late Closure (Frazier, 1978), or Recency Preference (Gibson, 1991); in each of the proposed parsing models, the preference is stated as an explicit processing heuristic. By contrast, the competitive attachment model explains recency effects in parsing as a side-effect of the active memory management techniques that are required to maintain the pool of network processing nodes, thereby providing a more parsimonious account of the data.

[4]The stack is *not* a global control mechanism; it is in fact a single (symbolically degenerate) phrasal node. A phrase pushes itself onto the stack by activating an attachment node between its XP and the stack node.

[5]In over 98% of the approximately 1400 simulations of attachment decisions in the parser, the network stabilized on one of these attachment sets (Stevenson, 1994).

Given the decay of phrasal activation described above, the explanation of recency effects in parsing is straightforward. When a phrase is allocated, there is a competition among the potential attachment nodes connecting it to the right edge of the existing parse tree. The phrases in the tree that were allocated most recently are those lower along the right edge and closer to the current phrase; they will have the highest activation levels because they have decayed less. All other effects on activation being equal, the attachment node between the current phrase and the most recent phrase in the parse tree to which it can attach will win the competition. For example, consider the sentence:

I $[_{VP}$saw the child who was $[_{VP}$sleeping $[_{PP}$on the train].

The PP *on the train* can attach to either verb phrase, modifying *saw* or *sleeping*, although the preference is for attachment to the lower verb phrase, *sleeping*. In the competitive attachment model, the PP does indeed attach to the lower VP. Since the lower verb phrase has been allocated more recently, its activation has decayed less than that of the higher verb phrase; it thus competes more effectively for the PP attachment.[6]

The decay of activation thus gives rise to recency effects that other models must account for by building in explicit structuring strategies. Furthermore, the decay of distributed phrases addresses the issue of the parser potentially exceeding its memory load by having to maintain and process too many attachment possibilities. For discussion surrounding the necessity of "closure" principles for weeding out excessive attachment sites, see Gibson (1991). What proposed closure mechanisms have in common is to remove the less recent of two potential attachment sites (e.g., Gibson, 1991; Kimball, 1973). The competitive attachment model achieves this without an additional pruning mechanism—older attachment sites simply fade out in time. Thus, this single mechanism parsimoniously accounts for what has been previously addressed with two separate heuristics: a recency or locality constraint, and a closure principle.

In addition to providing a simpler account, the competitive attachment model has empirical advantages as well. For example, all previous models of recency effects have addressed only the effect of *relative* recency—that is, they have only been concerned with explaining the choice the parser makes between more and less recent attachment possibilities. The heuristics are always based on comparing two attachment sites. But *absolute* recency effects can also be observed in human parsing—that is, the recency of a phrase can affect the ease of attaching to it, even when there is no competing attachment site. For example, in the sentence *I called the guy a rotten driver*, the NP *a rotten driver* has a single potential attachment site (within the VP). Compare a modified version of the sentence, where the NP in question is separated from the phrase to which it attaches by lengthy intervening material: *I called the guy who smashed my brand new car a rotten driver.*[7] Previous models, which only address relative recency effects arising from multiple attachment sites, predict these sentences to be equivalently easy. However, interpretation of the NP *a rotten driver* appears to be increasingly difficult with increasing length of the prior NP (*the guy who...*). The competitive attachment model accounts for this range of difficulty. A less recently allocated attachment site has less activation to contribute to an attachment node, causing an increase in the time it takes for the node to become fully activated.

Finally, the approach here naturally captures interactions between recency and other preference effects in a way that a simple strategy like Late Closure or Right Association cannot. The activation level of an attachment node is affected not only by recency of the phrases that are trying to attach, but by other factors such as lexical preferences and frequency effects as well. The competitive activation function of the model incorporates all these factors into a continuous function that assigns a "goodness of attachment" value, thereby avoiding the need for additional conflict resolution strategies to mediate between multiple discrete preference heuristics. In fact, this property leads to another prediction of the model that distinguishes it from earlier proposals: in some cases a less recent attachment may win, but the decrease in the activation it receives will cause it to be a weaker attachment.

Recency and Recovery

Because the competitive attachment parser uses a single parallel operation to attach an incoming phrase and simultaneously revise earlier attachments (if needed), the same factors that lead to an initial structural preference in the parser, also affect the ability of the parser to revise that structure. Earlier work has shown how the model accounts for the human parser being unable to recover from the misanalyses of certain ambiguities in well-known garden-path examples (Stevenson 1993b). Here the account is extended by showing how the parser experiences a range of difficulty when successfully reanalyzing structure, focusing in particular on the effects of the decay of activation of phrases over time.

Consider a verb such as *expect* that can take an NP or sentential object, as in: *The warden expects the report soon*, and *The warden expects the report to arrive soon*. The NP *the report* constitutes an ambiguous region in these sentences, since the parser has no way of knowing at that point whether the NP is the direct object of the verb, or the subject of the sentential object of the verb. Studies have revealed that the human parser prefers the interpretation in which the NP is the direct object of the verb (Frazier, 1978; Gorrell, 1987), although people have no trouble in processing verbs such as these when they occur with a sentential object.

[6]The competition between two attachments causes the network to require a larger number of iterations to stabilize than when there is only a single valid attachment for the current phrase. The model therefore predicts that a choice between fairly close attachments will cause a slow down in human parsing.

[7]This example is from Gibson (1991), who predicts no difference in processing the two versions of the sentence.

To learn more about how the human parser processes this type of ambiguity, Frazier & Rayner (1982) conducted eye-movement experiments on sentences with verbs such as these, using preferred (NP object) and non-preferred (sentential object) versions of the sentences.[8] They found that reading times per letter for the non-preferred versions of the sentences were longer than for the preferred versions of the sentences. The experiments also revealed that, in those sentences in which the verb occurred with a sentential object, reading times within the subordinate verb phrase—the disambiguating region of the sentence that indicates the non-preferred structure—were significantly longer than in the prior regions of the sentence. The data indicate that even an unconscious revision of an initial preference slows down the human parser.

The experiment also included versions of the sentences in which the length of the ambiguous region (i.e., the NP whose attachment is ambiguous) was increased from two words to an average of 5–6 words, as in: *The warden expects [the report that the guards slept] to arrive soon.* Note that lengthening the ambiguous region has the effect of delaying the onset of the disambiguating information that indicates a sentential complement. Frazier & Rayner found that the reading times for the disambiguating region in sentences with a 4–8 word ambiguous region were significantly longer than in sentences with a two word ambiguous region. Thus, the parser has increasing difficulty in revising its originally preferred hypothesis, the longer disambiguating information is delayed.

Previous parsing models, serial or parallel, have been unable to provide a straightforward account of this varying degree of difficulty in the successful revision of a preferred analysis. Although proponents of serial models can account for the finding that the human parser takes time to reanalyze structure, they must rely on devising explicit recovery strategies (Frazier & Rayner, 1982; McRoy & Hirst, 1990; Pritchett, 1992). Parallel models avoid this, but are unable to account for the observed revision time, since "reanalysis" consists of the immediate adoption of an available alternative (Gibson, 1991; Gorrell, 1987), or the immediate extension of existing structure (Weinberg, 1991).[9] Neither type of model can explain the increased difficulty in processing disambiguating information following a lengthy ambiguous region.

In contrast to serial models, the unique parallel structuring operation of the competitive attachment model entails that acceptable revision is an integral part of the normal processing operation of the parser—that is, ad hoc revision strategies are unnecessary. On the other hand, in contrast to parallel models, revising a preference involves an explicit reanalysis of attachment relations,

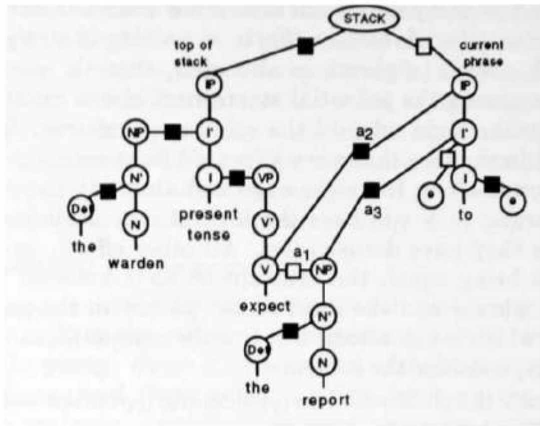

Figure 3: Reanalyzing the initial attachment of a short post-verbal NP.

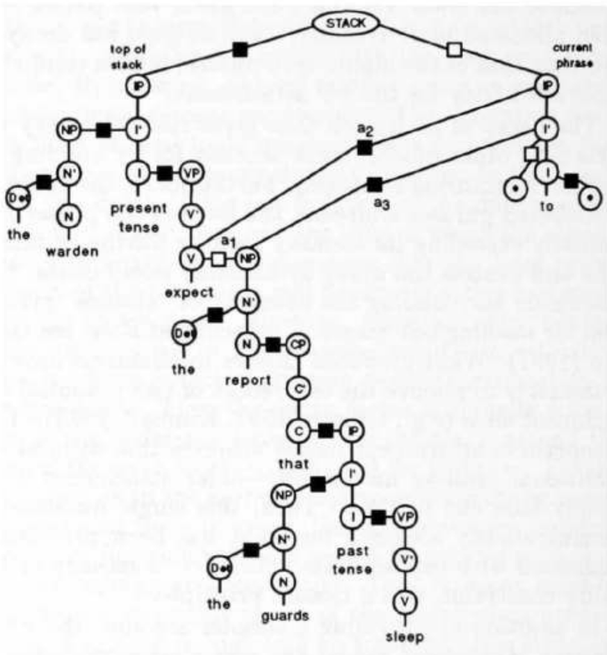

Figure 4: Reanalyzing the initial attachment of a long post-verbal NP.

and the necessary competition for activation takes time. An initial attachment within the network that requires no revision of earlier attachments takes 17 iterations, while an attachment that involves revision takes a minimum of 24 iterations. Furthermore, since the revision of attachments occurs as part of the same competitive operation that determines initial attachments, the same factors that affect the outcome of the competition for an initial attachment are also brought to bear in reanalysis. Thus, the same recency effects that arise in initial attachment preferences will also occur in attempting to revise those attachments, leading to a simple and elegant account of the effect of increased length of an ambiguity on the difficulty of revising its initial analysis.

For example, Figures 3 and 4 show the state of the

[8] Frazier & Rayner used sentences with tensed subordinate clauses. Since the processing results for the competitive attachment parser are the same for tensed and infinitival clauses (Stevenson, 1994), I will continue to use examples with infinitival clauses for ease of display.

[9] Weinberg's serial model employs implicit parallelism in its reliance on underspecification of structure.

parser after processing the infinitival marker *to* in the following example sentences containing a short and a long ambiguous region, respectively: *The warden expects [the report] to arrive soon*, and *The warden expects [the report that the guards slept] to arrive soon*. The IP (an infinitival sentential phrase) must attach to the main verb *expect*, replacing the prior NP as the object of the verb. Note that this revision is one of those allowed under the competitive attachment operation (compare the attachment configurations shown in Figures 3 and 4 to Figure 2). In each of the sentences of Figures 3 and 4, the parser is faced with exactly the same attachment choices; only the distance between the current phrase and its single attachment site differs in the two structures. The length of the NP *[the report that the guards slept]* means that *to* must attach to a less recent phrase than after a short NP like *[the report]*. Not surprisingly, given the decay of phrasal activation, the attachment of the IP takes longer in the network of Figure 4 (31 iterations), than in the network of Figure 3 (24 iterations).[10]

Thus, the memory management techniques of the model that lead to a general recency *preference* also provide an explanation for recency effects in *reanalysis*. The influence of lexical strength can also contribute to making a possible reanalysis more or less difficult for the parser (Stevenson, 1994). The competitive behavior of the model therefore provides an account of the range of difficulty of reanalysis, incorporating precisely the same factors, such as recency and lexical strength, that affect preferences.

Conclusions

The competitive attachment model provides a parsimonious and explanatory account of human preference and reanalysis behavior in the processing of syntactic ambiguities. The single mechanism of decay of activation, which is independently required in the model for managing the finite pool of processing nodes, underlies both relative and absolute recency effects in making initial attachments for a phrase. Furthermore, the competition mechanism that determines the attachment choices can naturally integrate the effects on activation from the decay process with other effects on activation of an attachment, such as lexical preferences and priming. The unique parallel attachment operation of the model entails that the exact same effects are observed in the successful revision of attachments. This allows the model to account for the previously unexplained phenomenon of the length of an ambiguity affecting the ease of revising its initial structural analysis. In fact, because of the model's consistent application of the competitive mechanism to both attachment and coindexation relations, these same recency effects were shown to arise in the domain of filler/gap processing as well (Stevenson, 1993a). The results of the model underscore the importance of determining computational mechanisms that can explain an extensive range of human behavior, rather than simply specifying heuristics that capture individual surface observations.

[10] Note that the increased time to make the revision relies not on a comparison between two available attachments, but from the great distance between the current phrase and its single attachment site. This example is thus quite similar to the absolute recency preference exhibited in the "rotten driver" example of the previous section. Stevenson (1994) in addition demonstrates cases of reanalysis within the model that exhibit relative recency effects, involving a revision of more or less recent attachment sites.

References

Abney, S. (1989). "A Computational Model of Human Parsing." *Journal of Psycholinguistic Research* **18:1**, 129–144.

Frazier, L. (1978). "On Comprehending Sentences: Syntactic Parsing Strategies." Doctoral dissertation, University of Connecticut.

Frazier, L., and K. Rayner (1982). "Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences." *Cognitive Psychology* **14**, 178–210.

Gibson, E. (1991). "A Computational Theory of Human Linguistic Processing: Memory Limitations and Processing Breakdown." Doctoral dissertation, Carnegie-Mellon University.

Gorrell, P. (1987). "Studies of Human Syntactic Processing: Ranked-Parallel versus Serial Models." Doctoral dissertation, University of Connecticut.

Kimball, J. (1973). "Seven Principles of Surface Structure Parsing in Natural Language." *Cognition* **2:1**, 15–47.

MacDonald, M. (in press). "Probabilistic Constraints and Syntactic Ambiguity Resolution." *Language and Cognitive Processes.*

McRoy, S. and G. Hirst (1990). "Race-Based Parsing and Syntactic Disambiguation." *Cognitive Science* **14**, 313–353.

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

Reggia, J. (1987). "Properties of a Competition-Based Activation Mechanism in Neuromimetic Network Models." *Proceedings of the 1st International Conference on Neural Networks*, 131–138.

Stevenson, S. (1993a). "Establishing Long-Distance Dependencies in a Hybrid Network Model of Human Parsing." *Proceedings of the 15th Annual Conference of the Cognitive Science Society*, 982–987.

Stevenson, S. (1993b). "A Competition-Based Explanation of Syntactic Attachment Preferences and Garden Path Phenomena." *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, 266–273.

Stevenson, S. (1994). "A Competitive Attachment Model for Resolving Syntactic Ambiguities in Natural Language Parsing." Doctoral dissertation, University of Maryland.

Weinberg, A. (1991). "A Parsing Theory for the Nineties: Minimal Commitment." Manuscript, University of Maryland.