

A Corpus Analysis of Recency Preference and Predicate Proximity

Edward Gibson

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139
gibson@psyche.mit.edu

Jacob Loomis

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139
jwloomis@athena.mit.edu

Abstract

The recent availability of large on-line parsed corpora makes it possible to test theories of psycholinguistic complexity by comparing the frequency distributions of closely related constructions. In this paper, we use this technique to test the psycholinguistic theory proposed by Gibson et al. (1993), which includes two independently motivated attachment principles: Recency Preference and Predicate Proximity. In order to test this theory, we examined two general classes of attachment ambiguities from the parsed Wall Street Journal corpus from the Penn Treebank: 1) ambiguities which involve three prospective noun phrase attachment sites; and 2) ambiguities which involve three prospective verb phrase attachment sites. Given three prospective noun phrase (NP) sites in English, the theory most naturally predicts a complexity ordering of NP₃ (easiest, most recent), NP₁, NP₂, but a ranking of VP₃, VP₂, VP₁ for verb phrase attachments. Our corpus analyses support both of these predictions.

Introduction

In order to test theories of psycholinguistic complexity, until recently it has been necessary to perform behavioral experiments, in which difficulty measurements (*e.g.* reading times) on systematically varied experimental conditions can be compared. With the advent of large on-line parsed corpora, we now have an additional source of data for testing psycholinguistic complexity theories. In particular, if we assume that the human sentence production mechanism generally prefers to minimize its burden by generating constructions that are less complex, then the frequency of a construction provides an indirect (partial) measurement of its production complexity. That is, if factors such as the function of a construction can be controlled for, and a sufficient quantity of target and control items can be observed, then the frequency of a construction should be an index of its production complexity (Gibson & Pearlmutter, 1994).¹

¹It is important to note that unless the constructions being compared are functionally equivalent, the frequency comparison will not be an indication of complexity. For example, although the declarative and interrogative form of an underlying syntactic structure are closely related syntactically, they are functionally very different, and one would therefore be unsurprised at large frequency differences between the two types of construction.

Furthermore, it is natural to assume that there is some overlap between the production and comprehension complexity metrics used by the human linguistic processor(s), and, regardless of overlap between the production and comprehension systems, most of the text currently available in on-line corpora has been edited at some stage to make it more understandable (*e.g.* newspaper stories, *etc.*). Thus we have good reason to expect that corpus frequencies should reflect human sentence comprehension complexity.

One instance where such a frequency comparison seems reasonable is in the area of ambiguous constituent attachment, *e.g.* prepositional phrases (PPs). If a particular structure contains two prospective attachment sites, and attachment to one of these sites is significantly more frequent than to the other, then the first resolution is probably less complex than the second. (Note that there is no plausible functional difference between these two different attachments.)

Gibson & Pearlmutter (1994) applied this technique to ambiguities involving PP attachment, in order to test the syntactic complexity theory of Gibson et al. (1993), who provided evidence in favor of two parsing principles: Recency Preference (Gibson, 1991; cf. Right Association, Kimball, 1973; and Late Closure, Frazier, 1978; 1987) and Predicate Proximity (cf. Milsark, 1983; Abney, 1989; Frazier, 1990; Gibson, 1991). Recency Preference prefers attachments to structures associated with more recent words in the input. Predicate Proximity, on the other hand, prefers attachments to be structurally as close as possible to the head of a predicate phrase, which is assumed to be a component of the structure associated with a main (non-auxiliary) verb (see Bowers, 1993, and the references there). In ambiguous attachments of a post-head modifier, such as a PP, to the subject noun phrase, these two principles work against each other:

(1) [*S* [*NP*₁ ... *N*₁ [*PP*₁ Prep₁ [*NP*₂ ... *N*₂]]]] PP₂

In (1), Recency Preference favors attachment of PP₂ to the lowest (most recent) site, NP₂. In contrast, Predicate Proximity favors attachment of PP₂ to the highest site, NP₁, because NP₁ is structurally closest to the predicate of the sentence, whether NP₁ is the subject of the verb (as in (1)) or whether it is the object of a preceding verb.² When factors such as lexical prefer-

²In order for Predicate Proximity to apply in cases where the predicate has not yet been encountered in the input string

ences (Ford *et al.* 1982) and argumenthood are equivalent between the two prospective NP attachment sites, attachment to the most recent site is slightly preferred in English (Cuetos & Mitchell, 1988; Gibson *et al.*, 1993). Hence, Recency Preference is assumed to be weighted more strongly than Predicate Proximity in English.

This theory also makes interesting predictions with respect to structures involving three prospective NP attachment sites:

(2) [_S [_{NP₁} ... N₁ [_{PP₁} Prep₁ [_{NP₂} ... N₂ [_{PP₂} Prep₂ [_{NP₃} ... N₃]]]]]] PP₃

In ambiguities such as these, the Recency/Predicate Proximity theory predicts that attachment to the most recent site NP₃ will be the least complex (and hence most frequent), followed by an undetermined order of NP₁ and NP₂.³ Gibson *et al.* (1993) provided behavioral evidence consistent with this prediction in examples like (2), in which the attaching phrase was a relative clause:

- (3) a. NP₁ attachment: [_{NP₁} the computer near [_{NP₂} the models of [_{NP₃} the buildings that was destroyed in the fire
 b. NP₂ attachment: [_{NP₁} the computers near [_{NP₂} the model of [_{NP₃} the buildings that was destroyed in the fire
 c. NP₃ attachment: [_{NP₁} the computers near [_{NP₂} the models of [_{NP₃} the building that was destroyed in the fire

Because of number agreement, the relative clause *that was destroyed in the fire* can only attach to one of the preceding three NP sites. Gibson *et al.* found that attachment to NP₃ was easiest and attachment to NP₁ was next, with attachment to NP₂ much the most difficult. This evidence, together with similar evidence from the processing of Spanish and intuitive evidence in VP attachment preferences in both English and Spanish (see Gibson *et al.*, 1993), supports the Recency/Predicate Proximity theory. Gibson *et al.*'s data also constrained the Recency/Predicate Proximity theory such that the order NP₁, NP₂ is expected. No other accounts of these data have yet been proposed, but it is worthwhile to note that pre-existing theories of complexity are inadequate. Most theories from both the computational and psycholinguistic literature include a preference for recent attachments, as well as an allowance for lexical preferences and argument attachments to outweigh recency

(as in (1)), it must be the case that the parsing algorithm predicts a predicate phrase (*S etc.*) in advance, *e.g.*, at the point when the initial NP is input. We are therefore assuming a partially top-down, predictive parsing algorithm, following Kimball (1973, 1975) and Frazier & Fodor (1978) among many others, who assume a predictive parsing algorithm for independent complexity reasons. See Gibson (1991) for a predictive parsing algorithm that is compatible with the assumptions under consideration here.

³No prediction is made about the relative complexity of attachments to NP₁ and NP₂ because, while attachment to NP₁ has Predicate Proximity in its favor, it has Recency against it. Similarly, attachment to NP₂ has Predicate Proximity working against it, but it is a more recent attachment site.

(see *e.g.* Frazier, 1987; Hobbs & Bear, 1990; Pritchett, 1988; Schubert, 1984; Taraban & McClelland, 1988; Whittemore *et al.*, 1990; Wilks *et al.*, 1985; Woods, 1973). However, none of these theories allows for the possibility that a less recent attachment (to NP₁) will be favored over a more recent attachment (to NP₂) if lexical and argument preferences are factored out.⁴

Using the corpus analysis methodology described above, Gibson & Pearlmutter (1994) observed the same complexity ranking as in the behavioral studies, in an analysis of constructions of the form in (2) (*i.e.* PP, not relative clause, attachment) from the one-million-word Brown corpus (Kučera & Francis, 1967). Table 1 lists the results for the 105 examples in which 1) lexical preferences did not play a role; and 2) attachment was unambiguously to one of NP₁, NP₂, or NP₃.

Table 1: Unambiguous PP₃ modifier attachments, Brown corpus (Gibson & Pearlmutter, 1994).

Attachment Site	Number of Tokens (Percentage)
NP ₁	27 (26%)
NP ₂	10 (10%)
NP ₃	68 (65%)
Totals	105 (100%)

χ^2 tests show that the frequencies of the attachments to each of the three sites are significantly different, resulting in the predicted complexity ranking of NP₃, NP₁, NP₂.

In this paper, we provide further corpus-analysis tests of this psycholinguistic complexity theory. These tests were performed on a second corpus, the Penn Treebank Wall Street Journal corpus consisting of roughly one million words of parsed text. Because this corpus is fully parsed — not just tagged with part of speech information as was the Brown corpus that Gibson & Pearlmutter used — it is much more appropriate for the task at hand. Thus, although much work still must be done by hand to filter out attachments that are not directly relevant, this hand-parsing is far less than the hand-parsing that is required when using a tagged corpus.

First, we analyzed patterns which match the form in (2), thus attempting to replicate Gibson & Pearlmutter's findings in the new corpus. Second, we examined patterns closely related to (2), in which a clause (labeled SBAR or SBARQ in the Penn Treebank) is the attaching item rather than a prepositional phrase. Motivation for this test is provided by the fact that Gibson *et al.*'s

⁴One potential account of the data is provided by the linguistic tuning framework of Cuetos & Mitchell (1988) and Mitchell & Cuetos (1991). Under this framework, parsing preferences in ambiguous constructions are tuned by imitating the resolution patterns of similar ambiguities observed in the input. This hypothesis is compatible with the preferences observed by Gibson *et al.* (1993), because less frequent patterns are more complex, as predicted. However, no particular version of the tuning framework has been specified in any detail, so it is difficult to test it empirically yet. See Gibson *et al.* (1993) for further comments on the tuning framework.

behavioral experiments involved the attachment of a relative clause, not a prepositional phrase, so that Gibson & Pearlmutter's earlier corpus-analysis tests are actually only indirect tests of the relationship between frequencies and complexity, relying on the assumption that all modifier attachments (e.g., at least PPs and SBARs) behave in the same way. The analysis of clause attachments in patterns like (2) provides a direct test of the original hypothesis, together with an indirect test of this assumption.

Third and most interestingly, we examine how the preferences differ when the prospective attachment sites are VPs rather than NPs. The Recency/Predicate Proximity theory predicts a different preference ordering among three VP sites from that which is predicted among three NP sites. In particular, this theory predicts a preference ordering from most to least recent for VP sites: VP₃, VP₂, VP₁. This prediction is made because, unlike the case with NP sites, Predicate Proximity prefers all three VP attachment sites equally: All are equally close to a predicate phrase (i.e. attachment is as close as possible in each case). Thus Recency preference dominates, and a strong recency effect is predicted between succeeding less recent attachment sites.

This VP attachment test also serves to distinguish Predicate Proximity from a primacy preference principle, in which the earliest attachment sites of a given category might be uniformly favored. Such an explanation would predict a similar pattern of complexity in VP attachments as in NP attachments (VP₃, VP₁, VP₂), and thus could lead to a partial collapse of these attachment preferences with serial position effects from the short-term memory literature (see, for example, Murdock, 1962; Tulving, 1968; and the references in each). The VP attachment data is therefore important for distinguishing these theories.⁵

The corpus analyses

All instances of each of the three target patterns in the corpus were obtained using programs developed specifically for this task by the second author. Following Gibson & Pearlmutter (1994), we then hand-filtered the following general classes of items:

1. Phrases that were mis-parsed in some way in the Penn Treebank.⁶
2. Items in which one of the three potential sites was unlikely to be an attachment site for principled reasons, such as:
 - (a) Items in which one of the sites which was tagged as an NP was part of a complex verb or idiomatic expression, such as *in spite of*, *in connection with*, or *is contingent upon*.

⁵See also Gibson et al. (1993) for reasons from the preferences in two-site ambiguities that the primacy explanation is probably not correct.

⁶In a few instances (no more than three per analysis set), the mis-parse was such that the corrected structure was actually an instance of one of the other attachments in question. Such instances were not filtered: they were included in their appropriate site counts.

- (b) Items in which any of the prospective attachment sites was part of a complex proper name that includes at least one prepositional phrase, such as *the State University of New York* and *Committee for the Scientific Branch of the Paranormal*.
 - (c) Items in which one of the NP sites was a quantifier such as *one of* or *some of*, which are unlikely attachment sites.
 - (d) Items in which the presence of intervening punctuation such as quotation marks or dashes made at least one of the attachment sites less likely.
3. Items in which the attaching item (PP, SBAR, or AdvP) is strongly preferred to attach to one of the three prospective sites by reasons of argument structure or lexical preferences. It is well-known that argument attachment are generally preferred over modifier attachments, often even in cases in which recency favors the non-argument attachment. For example, Frazier's (1978, 1987) Minimal Attachment principle favors argument attachments over modifier attachments because argument attachments involve fewer phrase structure nodes. Thus only cases involving potential modifier attachments are considered here. Examples of prepositional arguments of NPs that have been filtered here include *resignation from*, *reduction of*, and *heart of*.
 4. Items in which the attaching item could attach in more than one location without distinguishable differences in meaning (cf. Hobbs & Bear, 1990; Hindle & Rooth, 1993).

Prepositional phrase attachment to three NP sites

In order to replicate the work of Gibson & Pearlmutter (1994) in the Wall Street Journal corpus, we collected all patterns matching (2) and filtered these according to the steps listed above. Examples of the remaining attachments are given in (4), one for each of the possible attachment sites:

- (4) a. NP₁: [_{NP₁} the introduction of [_{NP₂} a major new family of [_{NP₃} computers [_{PP₃} in April]
- b. NP₂: [_{NP₁} shareholder support for [_{NP₂} a sale to [_{NP₃} a labor-management group [_{PP₃} at the last price discussed by that group]
- c. NP₃: [_{NP₁} potential problems with [_{NP₂} the construction of [_{NP₃} two big cruise ships [_{PP₃} from Finland]

After all of the less relevant items had been removed, 88 matching instances remained, all of which involved unambiguous attachment of a modifying PP to one of the three prospective NP sites. The breakdown of these 88 items by attachment site is presented in Table 2.

Table 2: Unambiguous PP₃ attachments (modifiers only).

Attachment Site	Number of Tokens (Percentage)
NP ₁	20 (23%)
NP ₂	10 (11%)
NP ₃	58 (66%)
Totals	88 (100%)

When compared with the data in Table 1, it seems clear that we have a very similar pattern of attachment preferences here. Statistics confirm these intuitions. A χ^2 analysis including all three attachment types reveals a significant effect ($\chi^2(2) = 43.73, p < .001$). Furthermore, attachments to NP₃ are the most frequent (vs. NP₁: $\chi^2(1) = 17.55, p < .001$; vs. NP₂: $\chi^2(1) = 32.48, p < .001$), and attachments to NP₁ are marginally more frequent than attachments to NP₂ ($\chi^2(1) = 2.70, p = .10$).⁷ Thus the Wall Street Journal corpus analysis provides an additional confirmation of the complexity ordering NP₃, NP₁, NP₂ predicted by the theory of Recency Preference and Predicate Proximity, and observed earlier by Gibson & Pearlmutter in the Brown corpus.

Clause attachment to three NP sites

Because the original behavioral experiments of Gibson *et al.* (1993) tested the complexity of a relative clause attachment rather than a prepositional phrase attachment in patterns like (2), it is worthwhile to perform a corpus analysis on relative clause attachments in order to verify that there are no unexpected differences. Part of the reason that Gibson & Pearlmutter (1994) did not analyze relative clause attachments was that cases of such attachments are hard to identify in an unparsed corpus. Since the corpus that we are analyzing here is parsed, the appropriate patterns were relatively easy to identify.

Because relative clauses are not tagged differently from other modifying clauses in the Penn Treebank corpus, our analyses include all modifying clauses. Furthermore, because some relative clauses were tagged with the category SBAR and others with the category SBARQ, our analysis includes both such categories. All instances matching the general pattern in (2) with SBAR(Q) as the attaching item were gathered and we then hand-filtered this set as described above. Examples of the remaining 88 items are as follows:

- (5) a. NP₁: [_{NP₁} the 20% decline in [_{NP₂} shares of [_{NP₃} the bank [_{SBAR} since the firm lowered its rating in early October]
- b. NP₂: [_{NP₁} public outrage over [_{NP₂} Mr. Mack's violent attack on [_{NP₃} a young woman [_{SBAR} when he was 19 years old]
- c. NP₃: [_{NP₁} the cloak of [_{NP₂} outrage at [_{NP₃} a war [_{SBAR} they characterized as "immoral"]

⁷All χ^2 tests with exactly one degree of freedom in this paper have Yates' correction for continuity applied (Hays, 1988).

The breakdown of attachments of the 88 SBAR(Q) attachments are given in Table 3.

Table 3: Unambiguous SBAR and SBARQ attachments (modifiers only).

Attachment Site	Number of Tokens (Percentage)
NP ₁	7 (8%)
NP ₂	2 (2%)
NP ₃	79 (90%)
Totals	88 (100%)

This pattern is quite similar to the complexity rankings observed earlier. First, there is obviously an effect of attachment site ($\chi^2(2) = 126.57, p < .001$). Furthermore, attachments to the most recent site NP₃ are more common than to either NP₁ ($\chi^2(1) = 58.62, p < .001$) or NP₂ ($\chi^2(1) = 71.30, p < .001$). Although attachments to NP₁ are numerically more common than those to NP₂, this difference is not significant, probably because our sample size is not large enough yet to observe this difference. Note, however, that the difference is in the predicted direction and we thus have additional support for the complexity ranking of NP₃, NP₁, NP₂, which is predicted by the Recency/Predicate Proximity theory.

Although the distributions of PP attachments and SBAR(Q) attachments have the same general preference ordering among the sites, the two distributions are not identical ($\chi^2(2) = 14.81, p < .001$). The tendency to attach low in SBAR(Q) attachment ambiguities seems to be stronger (90%) than in PP attachment ambiguities (66%). We do not yet know the cause of this difference, so we leave this issue for future work.

Adverbial attachment to three VP sites

A crucial test of the Recency/Predicate Proximity theory comes from attachments to succeeding VP sites. Unlike the prediction in multiple NP sites, the Recency/Predicate Proximity theory predicts a preference ordering which obeys strict recency in VP attachments, because Predicate Proximity is factored out in such instances. In contrast, a theory which explains the NP attachment pattern in terms of a competing preference for the earliest available attachment site predicts a preference ordering in VPs which is the same as in NPs, *i.e.*, VP₃, VP₁, VP₂. In order to test these competing hypotheses, we extracted all sentences with three consecutive VPs and a following adverbial modifier tagged ADVP, which was a prospective modifier for each of the three VP sites. As above, we filtered items that could ambiguously attach without a noticeable difference in meaning, as well as all argument attachments (*e.g.*, many particles are tagged as adverbs in the Penn Treebank), idiomatic expressions and mis-parsings. The breakdown of the remaining 95 items is presented in Table 4, with an example of each attachment first.

- (6) a. VP₁: No applicable examples in the corpus
 b. VP₂: Composer Marc Marder ... has [VP₁ prepared an exciting, eclectic score that [VP₂ tells you what the characters are [VP₃ thinking and feeling]] [AdvP far more precisely than intertitles, or even words, would.]]
 c. VP₃: Texaco Inc. [VP₁ reported an 11% increase in third-quarter earnings, which it [VP₂ attributed partly to the company's massive restructuring after it [VP₃ emerged from bankruptcy-law proceedings] [AdvP 18 months ago.]]]

Table 4: Unambiguous adverbial attachments to VPs (modifiers only).

Attachment Site	Number of Tokens (Percentage)
VP ₁	0 (0.0%)
VP ₂	8 (8.4%)
VP ₃	87 (91.6%)
Totals	95 (100%)

These results provide clear evidence in favor of the Predicate Proximity theory over the primacy theory. There is a main effect of attachment site ($\chi^2(2) = 146.04$, $p < .001$), with VP₃ attachments much the most frequent (vs. VP₁: $\chi^2(1) = 85.01$, $p < .001$; vs. VP₂: $\chi^2(1) = 64.04$, $p < .001$). Furthermore, VP₂ attachments are much more frequent than VP₁ attachments ($\chi^2(1) = 6.13$, $p = .01$), as predicted by the Predicate Proximity theory, but not by the primacy theory. Thus we have striking evidence in favor of the independently motivated Recency/Predicate Proximity theory.

Summary and conclusions

The three corpus analysis tests of the Recency Preference/Predicate Proximity theory strongly support this theory. First, it was found that attachments to the most recent of three sites of the same category (NP or VP) are much more common than attachments to less recent sites. Second, it was found that the least recent site is preferred over intermediate sites when the category of the site is an NP, as predicted by the Recency/Predicate Proximity theory. And third, and most strikingly, it was found that attachment preferences to VP sites do not follow the same pattern as observed in NP sites, but rather they follow a strict recency preference, with more recent sites preferred over more distant ones. This result is just as predicted by the Recency/Predicate Proximity theory, and disconfirms a potential explanation of the data in terms of primacy.

Acknowledgements

We would like to thank Don Mitchell, Neal Pearlmuter, and three anonymous reviewers for their comments on an earlier draft of this work.

References

Abney, S.P. (1989). A computational model of human parsing. *Journal of Psycholinguistic Research*, 18, 129-144.

Bowers, J. (1993). The syntax of predication. *Linguistic Inquiry*, 24, 591-656.

Cuetos, F. & Mitchell, D.C. (1988). Cross-linguistic differences in parsing: Restrictions on the use of the Late Closure strategy in Spanish. *Cognition*, 30, 73-105.

Ford, M., Bresnan, J., & Kaplan, R. (1982). A competence-based theory of syntactic closure. In J. Bresnan (Ed.), *The mental representation of grammatical relations* (pp. 729-796). Cambridge, MA: MIT Press.

Frazier, L. (1978). On comprehending sentences: Syntactic parsing strategies. Unpublished doctoral dissertation. Storrs, CT: University of Connecticut.

Frazier, L. (1987). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and performance XII*, (pp. 559-586). Hillsdale, NJ: Lawrence Erlbaum Associates.

Frazier, L. (1990). Parsing modifiers: Special purpose routines in the human sentence processing mechanism? In D.A. Balota, G.B. Flores d'Arcais & K. Rayner (Eds.), *Comprehension processes in reading*, (pp. 303-330). Hillsdale, NJ: Lawrence Erlbaum Associates.

Frazier, L., & Fodor, J.D. (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. Unpublished doctoral dissertation. Pittsburgh, PA: Carnegie Mellon University.

Gibson, E. & Pearlmuter, N. (1994). A corpus-based analysis of psycholinguistic constraints on prepositional phrase attachment. In C. Clifton Jr., L. Frazier, & K. Rayner (Eds.) *Perspectives in Sentence Processing* (pp. 181-198). Hillsdale, NJ: Lawrence Erlbaum Associates.

Gibson, E., Pearlmuter, N., Canseco-Gonzalez, E., & Hickok, G. (1993). Cross-linguistic attachment preferences: Evidence from English and Spanish. Manuscript. Cambridge, MA: Massachusetts Institute of Technology.

Hays, W. (1988). *Statistics* (4th ed.). Orlando, FL: Holt, Rinehart and Winston.

Hindle, D. & Rooth, M. (1993). Structural ambiguity and lexical relations. *Computational Linguistics*, 19, 103-120.

Hobbs, J.R. & Bear, J. (1990). Two principles of parse preference. *Proceedings of the Thirteenth International Conference on Computational Linguistics* (vol. 3, pp. 162-167). Helsinki: University of Helsinki.

Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition*, 2, 15-47.

Kučera, H. & Francis, W.N. (1967). *Computational analysis of present day American English*. Providence, RI: Brown University Press.

Milsark, G.L. (1983). On length and structure in sentence parsing. *Cognition*, 13, 129-134.

- Mitchell, D.C. & Cuetos, F. (1991). The origins of parsing strategies. In C. Smith (Ed.), *Conference proceedings: Current issues in natural language processing* (pp. 1-12). Austin, TX: University of Texas at Austin.
- Murdock, B., Jr. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, 64, 482-488.
- Pritchett, B. (1988). Garden path phenomena and the grammatical basis of language processing. *Language*, 64, 539-576.
- Schubert, L.K. (1984). On parsing preferences. *Proceedings of the Tenth International Conference on Computational Linguistics* (pp. 247-250).
- Taraban, R. & McClelland, J.R. (1988). Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. *Journal of Memory and Language*, 27, 597-632.
- Tulving, E. (1968). Theoretical issues in free recall. In T.R. Dixon & D.L. Horton (Eds.), *Verbal behavior and general behavior theory* (pp. 2-36). Englewood Cliffs, NJ: Prentice-Hall.
- Whittemore, G., Ferrara, K., & Brunner, H. (1990). Empirical study of predictive powers of simple attachment schemes for post-modifier prepositional phrases. *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics* (pp. 23-30).
- Wilks, Y., Huang, X., & Fass, D. (1985). Syntactic preference and Right Attachment. *Proceedings of the International Joint Conference on Artificial Intelligence 1985* (pp. 779-784).
- Woods, W.A. (1973). An experimental parsing system for transition network grammars. In R. Rustin (Ed.), *Natural language processing* (pp. 111-154). New York, NY: Algorithmics Press.