

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Architecture and Experience in Sentence Processing

Permalink

<https://escholarship.org/uc/item/53j2t766>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Authors

Stevenson, Suzanne

Merlo, Paola

Publication Date

1997

Peer reviewed

Architecture and Experience in Sentence Processing

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

Paola Merlo
University of Geneva,
and IRCS, University of Pennsylvania,
3401 Walnut Street
Philadelphia, PA 19104-6228
merlo@linc.cis.upenn.edu¹

Abstract

Models of the human sentence processing mechanism have traditionally appealed to innate architectural restrictions to explain observed patterns of behavior. Recently, a number of proposals have instead emphasized the role of linguistic experience in guiding sentence interpretation, suggesting that various frequency measures play a crucial role in ambiguity resolution. What has been lacking thus far is a detailed analysis of the linguistic and computational properties that could explain why those particular aspects of experience are effective in shaping behavior. In this paper, we present a linguistic analysis that reveals restrictions on the representational ability of the sentence processor, explaining its sensitivity to particular factors in the linguistic environment. The proposal receives strong support from a large-scale corpus analysis.

Introduction

Traditional theories of cognition propose that innate constraints on cognitive architecture are responsible for crucial aspects of human behavior. With advances in connectionism, an alternative view has arisen which instead emphasizes the role of experience in shaping relatively unrestricted cognitive mechanisms. In this view, it is the external influences on a cognitive mechanism, rather than the inherent restrictions within it, that determine human behavior. Models of the human sentence processing mechanism (HSPM) have largely fallen within the architectural paradigm, attributing interpretation preferences and processing breakdown to built-in limitations on linguistic representations or computational resources (e.g., Frazier and Fodor, 1978; Gibson, 1991; Lewis, 1993; Marcus, 1980). However, the contrasting experiential approach has recently gained prominence, proposing exposure to the linguistic environment as a primary determinant of the behavior of the HSPM (e.g., MacDonald, 1994; MacDonald, Pearlmutter, and Seidenberg, 1994; Mitchell, Cuetos, Corley, and Brysbaert, 1995; Spivey-Knowlton, Trueswell, and Tanenhaus, 1993; Trueswell, 1996; Trueswell, Tanenhaus, and Garnsey, 1994). Specifically in these approaches, the frequency of lexical features, words, or even grammatical constructs, plays a critical role in the sentence interpretation process, guiding the resolution of linguistic ambiguities

by favoring those aspects that have been encountered more frequently.

The challenge this view faces is not only to demonstrate which of the infinite possible frequency values the HSPM is actually sensitive to, but to explain *why* it is sensitive to those particular features of the environment. We believe that progress in this area relies on recognizing that the generally accepted dichotomy between architectural and experiential models of sentence processing is a false one. In order to predict which aspects of the environment the HSPM records and uses in processing, an experiential theory must make substantive assumptions regarding the internal properties of the underlying mechanism. In short, the interesting question is what representational and/or computational restrictions make the HSPM sensitive to some environmental distinctions and not others, returning the emphasis once more to architectural concerns.

In this paper, we explore the underlying representational basis for the role of certain environmental variables in ambiguity resolution. We investigate the main verb/reduced relative (MV/RR) ambiguity because it has been the focus of a number of experiments demonstrating the effects of frequencies in sentence processing. The following classic “garden-path” example demonstrates the severe processing difficulty that can be associated with the MV/RR ambiguity (Bever, 1970):

(1) The horse raced past the barn fell.

Problems arise here because the verb *raced* can be interpreted as either a past tense main verb, or as a past participle within a reduced relative clause (i.e., *the horse [that was] raced past the barn*). Because *fell* is the main verb of (1), the reduced relative interpretation of *raced* is required for a coherent analysis of the complete sentence. But the main verb interpretation of *raced* is so strongly preferred that the sentence processor breaks down at the verb *fell*, unable to integrate it with the interpretation that has been developed to that point.

Over the last several years, the contrast in acceptability of (1) with sentences like (2) and (3) has been noted:

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

(3) The referee kicked in the soccer game was angry.

Although all three sentences contain an MV/RR ambiguity, (1) is completely unacceptable, while (2) and (3) are relatively easy to understand. Proposals within the

¹The second author was supported jointly by the Swiss National Science Foundation advanced researcher fellowship 8210-46569 and by IRCS.

experiential paradigm have claimed that a crucial difference in the processing of various MV/RR examples arises from verb-specific frequency information that differentially supports the RR interpretation (MacDonald, 1994; MacDonald, Pearlmuter, and Seidenberg, 1994; Trueswell, 1996; Trueswell, Tanenhaus, and Garnsey, 1994). However, these approaches have pursued a relatively shallow level of frequency analysis—determining the properties that are superficially related to the different resolutions of the ambiguity (e.g., past participle versus main verb usage of a verb), and counting the occurrences of words in those uses. We observe that the ambiguous verbs in the examples above are from three different lexical semantic classes, and instead hypothesize that deeper properties of these verb classes are the basis for the differing behavior of the HSPM in response to sentences (1) through (3).

In support of this view, we present acceptability judgments and corpus data demonstrating that the frequencies previously proposed as factors in resolving the MV/RR ambiguity do not match the pattern of human preferences. The data further reveals consistent behavior across lexical semantic verb classes. Based on a detailed linguistic analysis of these classes, we propose a different frequency factor as the crucial environmental variable. Furthermore, we explain the sensitivity of the HSPM to this particular factor in terms of restrictions on its representational ability that follow from the syntactic analysis. We find remarkably strong corroboration in the corpus data for this proposal. There are two corollary results to this finding. First, the environmental factor that we propose models the preference behavior with a single frequency value, eliminating the need for a complex model of frequency combination. Second, the feature whose count best matches the data is one that, surprisingly, is not itself an element of the ambiguous linguistic construction, and whose potential relevance is only revealed by a deeper analysis of the lexical and syntactic properties of verbs.

Verb Classes and Judgment Data

The sentences in (1)–(3) above use ambiguous verbs from three different lexical semantic classes of optionally transitive verbs. Sentence (1) uses a manner of motion verb, *raced*. In English, these verbs form a subclass of unergative verbs (Levin and Rappaport Hovav, 1995), intransitive action verbs that may appear in a transitive form:

(4a) The horse raced past the barn.

(4b) The rider raced the horse past the barn.

The transitive form of an unergative (4b) is the causative counterpart of the intransitive form (4a), in which the subject of the intransitive becomes the object of the transitive (Brousseau and Ritter, 1991; Hale and Keyser, 1993; Levin and Rappaport Hovav, 1995). Sentence (2) uses an unaccusative verb, *melt*. Unaccusatives are intransitive change of state verbs which also have a causative transitive form:

(5a) The butter melted in the pan.

(5b) The cook melted the butter in the pan.

Finally, sentence (3) uses an object-drop verb, *kicked*; these verbs have a non-causative transitive/intransitive alternation, in which the object NP is optional:

(6a) The player kicked the referee.

(6b) The player kicked.

We crucially observe that the prototypically difficult reduced relatives—i.e., those in which the verb induces a very strong MV bias, as in sentence (1)—contain a manner of motion verb (an unergative). We investigated the relative difficulty of unergatives and unaccusatives by asking naive informants for acceptability judgments on sentences with RRs containing these verbs. We found that unergative verbs uniformly led to a severe garden path in the RR construction, while unaccusative verbs were overwhelmingly judged completely fine in the RR, with a few responses of them being slightly degraded (Stevenson and Merlo, 1997). Although we did not ask for judgments on object-drop verbs, our intuition was that they are readily interpretable in a RR. Support for this view comes from experiments in MacDonald (1994) including object-drop verbs, that showed that RRs are relatively easy to understand given a context that is not strongly biased toward a main verb reading. Thus, the difficulty of the RR interpretation patterns along verb class lines, with unergatives difficult, and unaccusatives and object-drop verbs relatively easy.²

Corpus Analysis

The fact that acceptability of the RR clearly patterns with verb classes leads us to believe that the difficulty of unergatives is a grammatical phenomenon. That is, we hypothesize that the critical behavioral and frequency differences between the above verbs in the MV/RR ambiguity are grounded in the differing lexical syntactic properties of the verb classes. In this section, we investigate whether the lexical semantic classification of these verbs and their differing behavior in the RR is reflected in different distributions of pertinent corpus frequencies. Specifically, we explore the following hypothesis across the three classes of verbs:

H₀: differences in the difficulty of the RR interpretation correspond to differences in the frequency of α .

This hypothesis was tested for several values of α —the RR construction itself, transitivity, passive voice, and past participle use. These properties were chosen because a verb in the RR is a passive (and therefore transitive) past participle. It has been claimed that the frequency of these lexical features correlates with ease of interpreting a RR (MacDonald, 1994; Trueswell, 1996; Trueswell, Tanenhaus, and Garnsey, 1994). For hypothesis H₀ to be confirmed, unergatives must exhibit a significantly lower

²Related proposals concerning the MV/RR ambiguity have relied only on coarse classification of verbs as obligatorily transitive or not (Gibson, 1991; Pritchett, 1992), while our approach is founded on a finer-grained lexical semantic distinction.

frequency of α compared to unaccusative and object-drop verbs. Since there is no sharp contrast in acceptability between unaccusative and object-drop verbs, support for the hypothesis also requires there to be little to no difference in α across these two verb classes.

Materials and Method

We chose a set of 10 verbs from each class, based primarily on the classification of verbs in Levin (1993): the unergatives are manner of motion verbs (*jumped, rushed, marched, leaped, floated, raced, hurried, wandered, vaulted, paraded*), the unaccusatives are verbs of change of state (*opened, exploded, flooded, dissolved, cracked, hardened, boiled, melted, fractured, solidified*), and the object-drop verbs are unspecified object alternation verbs (*played, painted, kicked, carved, reaped, washed, danced, yelled, typed, knitted*). Each verb presented the same form in the simple past and in the past participle, as in the MV/RR ambiguity. These verbs were chosen because they did not show massive departures from the intended verb sense; for example, *run* was eliminated because it occurs most often in phrases such as *run a meeting*, where it is not a manner of motion use. All verbs can occur in the transitive and in the passive. Verbs in the three sets were matched in frequency, and their logarithmic frequency varies between 2 and 4 inclusive.

Counts were performed on the tagged version of the Brown Corpus and on the portion of the Wall Street Journal distributed by the ACL/DCI (years 1987, 1988, 1989), a combined corpus in excess of 65 million words. Four pairs of counts were collected. First, each verb was counted in its main verb (i.e., simple past) and past participle uses, based on the part of speech tag of the verb in the corpora. Second, active and passive uses of the verbs were counted; cases in which usage could not be determined by a simple pattern search were classified by hand. The third count also required manual intervention: verbs were initially classified as transitive or intransitive according to a set of regular search patterns, then individual inspection of verbs was carried out to correct item-specific errors. In the fourth count, uses of the verb form as main verb or as reduced relative were collected. Reduced relatives were counted by hand after extracting from the corpus all occurrences of the past participle preceded by a noun.

Data and Analysis

Recall that the goal is to determine whether differences in the counts correspond to differences in the ease of interpreting a RR. Since MV/RR preferences follow verb class lines, we tested for differences in the distribution of the corpus counts across the verb classes. The counts within each verb class were pooled, and a χ^2 test was used to compare each pair of distributions, for each property counted.³

³Since we could not assume a normal distribution of the underlying data, the small sample size (10 verbs per group)

	RR	MV	TOTALS
UNERGATIVE	7	4910	4917
UNACCUSATIVE	21	3321	3342
TOTAL	28	8231	8259

$$\chi^2(1) = 12.508; p = .000405$$

	RR	MV	TOTALS
UNACCUSATIVE	21	3321	3342
OBJECT-DROP	202	2316	2518
TOTAL	223	5637	5860

$$\chi^2(1) = 212.447; p < .000001$$

Table 1: Frequencies of RR and main verb uses.

	PP	MV	TOTALS
UNERGATIVE	647	4910	5557
UNACCUSATIVE	1476	3321	4797
TOTAL	2123	8231	10354

$$\chi^2(1) = 576.623; p < .000001$$

	PP	MV	TOTALS
UNACCUSATIVE	1476	3321	4797
OBJECT-DROP	1939	2316	4255
TOTAL	3415	5637	9052

$$\chi^2(1) = 209.622; p < .000001$$

Table 2: Frequencies of past participle and main verb uses.

The contingency tables along with the resulting χ^2 values and significance levels are shown in Table 1 through Table 4. All the data sets show the same pattern. The properties related to the RR—transitivity, passive voice, and past participle use, as well as the RR construction itself—were less frequent for unergatives than for unaccusatives, and less frequent for unaccusatives than for object-drop verbs. The differences between the unergative and unaccusative distributions for each property are highly significant, as are the differences between the unaccusative and object-drop distributions.⁴

Discussion

The most striking finding is that the lexical semantic classes of the verbs correspond to significant and sizable differences in their distributions. This important result confirms that the frequency distributions to which the HSPM is exposed are themselves shaped by deeper linguistic properties of words. However, we also found that the hypotheses of correspondence between frequency and ease of interpretation of the RR is disconfirmed: for all properties, the corpus data indicates a three-way distinction between the verb classes, while preference data indicates a binary distinction, with only the unergative class

prevented us from using a parametric test, such as a t-test. In no case was the expected frequency for any of the cells less than 10, so it was unnecessary to apply Fisher's exact test.

⁴We also directly compared the unergative and object-drop distributions, which were significantly different—with $p < .000001$ —for each property counted.

	PASS	ACT	TOTALS
UNERGATIVE	139	5330	5469
UNACCUSATIVE	717	3930	4647
TOTAL	856	9260	10116

$$\chi^2(1) = 537.049; p < .000001$$

	PASS	ACT	TOTALS
UNACCUSATIVE	717	3930	4647
OBJECT-DROP	1339	3074	4413
TOTAL	2056	7004	9060

$$\chi^2(1) = 286.088; p < .000001$$

Table 3: Frequencies of passive and active uses.

	TRANS	INTR	TOTALS
UNERGATIVE	463	5065	5528
UNACCUSATIVE	2402	2359	4761
TOTAL	2865	7424	10289

$$\chi^2(1) = 2251.868; p < .000001$$

	TRANS	INTR	TOTALS
UNACCUSATIVE	2402	2359	4761
OBJECT-DROP	3355	922	4277
TOTAL	5757	3281	9038

$$\chi^2(1) = 762.189; p < .000001$$

Table 4: Frequencies of transitive and intransitive uses.

strongly disfavoring the RR.

Since each of the frequencies associated with the RR are significantly less for unergatives than for the other verbs, one might propose that the preference data simply results from a more complex frequency model taking multiple factors into account (e.g., as in Juliano and Tanenhaus, 1994; Collins, 1996), or from a simple threshold effect. We pursue an alternative approach of investigating whether additional linguistic information can lead to an individual frequency measure that more closely models the bimodal pattern of observed preferences.

Our Linguistic Analysis

We return to the observation that both unergatives and unaccusatives have a causative transitive form. We note, however, that they differ in theta role assignments. Theta roles are labels that identify the semantic relation of an argument to a verb. In an intransitive unergative—an action verb—the subject NP is an *agent*, and in an intransitive unaccusative—a change of state verb—the subject is a *theme*. The theta role assignments to the corresponding semantic arguments of the transitive forms—i.e., the direct objects—are the same, with the addition of a *causal agent* as subject in both cases. This leads to an unusual situation for a transitive unergative because it assigns two agent theta roles—the subject is the agent of causation, and the object is the agent of the action expressed by the verb. For example, in *The trainer jumped the lion through the hoop*, the trainer causes the lion to jump, but the lion

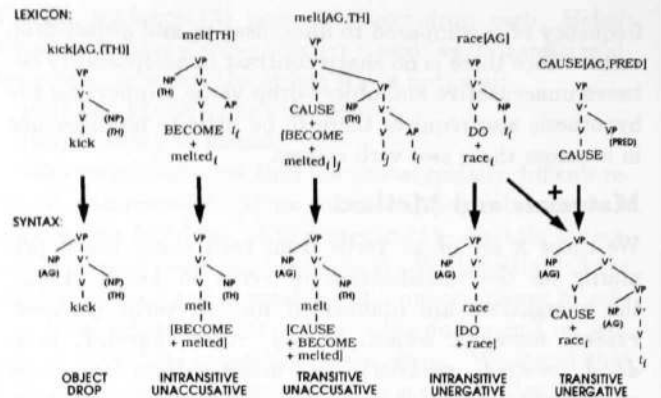


Figure 1: Lexical entries and the corresponding syntactic structures for various verb types. AG=agent; TH=theme.

is the agent of the actual jumping.

We follow Hale and Keyser (1993) in assuming that verb forms are created through the structural combination of lexical items with abstract verbal morphemes, such as DO, BECOME, and CAUSE; see Figure 1. Theta roles correspond to particular structural positions in the resulting lexical entries. Consequently, the difference in theta role assignment noted above is evidence of a difference in the underlying lexical structures of unergatives and unaccusatives. In this approach, the agent role has a distinguished status because it must be inserted into a subject position in the syntax. Since an intransitive unaccusative verb has no agent role, it can combine with CAUSE in the lexicon, allowing the causal agent to be inserted into the single subject position in the syntax. By contrast, a transitive unergative must be formed by a syntactic process in which the abstract verbal morpheme CAUSE takes the unergative verb as its complement in the syntax. This creates two syntactic subject positions, one in the embedded verb phrase for the agent of the unergative verb, and one in the main verb phrase for the agent of CAUSE (see Figure 1).⁵

This analysis has important lexical and syntactic consequences. An unaccusative verb has two distinct entries in the lexicon, both a transitive and intransitive form. By contrast, an unergative is underlyingly *intransitive only*. An unergative verb has a single (intransitive) lexical entry, as its transitive counterpart can only be formed through a syntactic process of verb incorporation. An object-drop verb differs from both unergatives and unaccusatives in having a single transitive lexical form, with its direct object argument marked as optional. However, the crucial distinction among the verb classes is that unergatives have only an intransitive form in the lexicon, while unaccusative

⁵ This analysis of the transitive form of unergatives is supported by, for example, observations on their semantic restrictions (Cruse, 1972), and by cross-linguistic facts related to causative formation in Hebrew (Levin and Rappaport Hovav, 1995). For further details and justification of the analysis, see Hale and Keyser (1993), Stevenson and Merlo (1997).

UNERG	UNACC	OBJ-DROP
race[ag]	melt[th]	melt[ag,th] kick[ag,(th)]

Table 5: Sample lexical entries for verbs of each class.

and object-drop verbs both have a transitive form in the lexicon. Table 5 shows the lexical entries for a sample verb in each class.

Our proposal is that the lexical nature of the transitive unaccusative and object-drop verbs is what makes those verb classes more readily interpretable in the RR construction. The reasoning is as follows. The RR entails the passive use of the past participle as a post-nominal modifier. A related passive use of the past participle is as a prenominal adjective (Levin and Rappaport, 1986). Our linguistic analysis has the perhaps surprising result that there is no transitive unergative lexical item, and that, therefore, unergative verbs have no passive (transitive) past participle in the lexicon. Furthermore, the process of syntactic causativization that forms the transitive unergative is inapplicable in the adjectival use, as it involves verb incorporation. The consequence is that the unergative past participle, in contrast to the other verb classes, cannot be used adjectivally (cf. Bresnan, 1982; Levin and Rappaport, 1986). We suggest that the inability of the unergative past participle to occur as a prenominal adjective makes difficult its interpretation in the similar use as a post-nominal modifier in the RR. By contrast, the prenominal adjective use of the unaccusative and object-drop verbs aids their interpretation in the RR.

Our conjecture then is that frequency of adjectival use should correspond to ease of interpretability of the RR construction, as both involve the use of the past participle as a nominal modifier. Specifically, we suggest the following hypothesis:

H'_0 : Unergative past participles occur significantly less often in an adjectival use than unaccusative and object-drop past participles, which are not significantly different in frequency.

If this hypothesis is borne out, then the adjectival use makes the right cut of verb classes in relation to the MV/RR preference.

Materials and Method

The same corpora were used as in the earlier counts. None of the verb forms are explicitly marked as adjectives in these corpora. To determine the counts of adjectival uses, we simply divided the verb occurrences labeled with the past participle part of speech tag into prenominal and other uses.

Data and Analysis

We again tested for differences in the distributions of corpus counts across verb class lines, using a χ^2 test. The

	ADJ PP	NON-ADJ PP	TOT
UNERG	21	626	647
UNACC	155	1321	1476
TOTAL	176	1947	2123

$$\chi^2(1) = 30.199; p < .000001$$

	ADJ PP	NON-ADJ PP	TOT
UNACC	155	1321	1476
OBJ-DROP	220	1719	1939
TOTAL	375	3040	3415

$$\chi^2(1) = 0.528; p = .467297 \text{ (not significant)}$$

Table 6: Frequencies of adjectival and non-adjectival past participle uses.

contingency tables along with the resulting χ^2 values and significance levels are shown in Table 6. The difference between unergative and unaccusative past participles in the adjectival use was highly significant,⁶ with unergative adjectival uses far less frequent than unaccusative adjectival uses. In contrast, the unaccusative and object-drop adjectival distributions are not significantly different.⁷

Discussion

The hypothesis is confirmed that unergative past participles are used as a prenominal adjective less frequently than unaccusative and object-drop past participles, which are not distinguished in frequency. The two-way distinction that emerges from the data is the first grouping of the verbs that corresponds to the difficulty lines indicated by the acceptability judgments on the RR construction, in which only the unergatives are strongly dispreferred. Our intuition that adjectival and RR uses are related, in that they are both nominal modifiers, is clearly supported. Thus, our linguistic analysis has led to an easily-computed frequency factor—the prenominal use of the past participle—that has the potential to explain human preferences in the MV/RR ambiguity.

However, since our claim is that unergatives do not have a transitive past participle that could be used adjectivally, we must question why unergative adjectives occur at all. Interestingly, we found that the unergative adjectival uses in the corpus were restricted to only 2 verbs: *hurried* occurred 20 times prenominally, and *rushed* once. Examination revealed that these were not the causative use of the

⁶The difference between unergative and object-drop participles was also highly significant, $p < .000001$.

⁷We also examined the frequency of adjectival uses compared to all other uses of the verb form, not just other past participle uses. These comparisons revealed a significant three-way split among the verb classes. Although we performed these statistical tests for completeness, we do not believe they represent the right comparison. We assume that frequency of adjectival uses is stored with the past participle, and thus the frequency differential that affects interpretation is the distribution of adjective versus non-adjective past participle use, and not the frequency of adjectival use versus all other uses of the verb.

verb. Recall from our linguistic analysis that the causative form of an unergative has an agentive direct object, corresponding to the agentive subject of the intransitive form. However, in the prenominal uses of *hurried* and *rushed* in the corpus, the modified noun (which corresponds to the direct object of the verb) is not an agent. For example, in *hurried head count*, the head count is not itself an agent of the hurrying action; this is particularly clear when we consider that *head count* cannot be the subject (agent) of the intransitive unergative: **The head count hurried*. The adjectival uses of these verbs thus correspond to a non-causative transitive form that has a *patient* or *theme* direct object (such as *head count*). This use does not have the “double agent” problem of the causative use, which prevents the lexicalization of the latter. The non-causative form of *hurry* and *rush* therefore occurs in the lexicon like any other “normal” transitive verb, and it is this form that is being used adjectivally in the corpus.

Conclusions

Recent work has demonstrated a range of environmental factors that influence sentence interpretation. What has been lacking thus far is a detailed analysis of the linguistic and computational properties that could explain why those particular aspects of linguistic experience are effective in shaping behavior. In this paper, we have shifted the focus from the environment itself, to what the environment reveals about the restrictions on knowledge and processing that are internal to the HSPM.

In investigating this issue, we focused on the varying acceptability of the reduced relative reading of verbs that exhibit a main verb/past participle ambiguity. Interestingly, the frequency of linguistic features directly associated with the reduced relative construction do not match the human preference data. The best predictor of human judgments of reduced relatives is the frequency with which the past participle of the verb is used as a prenominal adjective; our proposal is that this usage aids in the interpretation of the verb as a modifier within the reduced relative construction. The proposed explanation for the sensitivity of the HSPM to this particular experiential factor is grounded in independently justified theoretical assumptions. Specifically, our claim is that the differing frequency distributions among the verbs arise from underlying grammatical differences between lexical semantic verb classes, which restrict the representational ability of the HSPM. We conclude that, although the environment shapes our use of knowledge, an explanatory account of behavior requires a deeper understanding of the highly restricted form of that knowledge.

References

- Bever, T. G. (1970). The cognitive basis for linguistic structure. In J. R. Hayes (Ed.), *Cognition and the Development of Language*. New York: John Wiley.
- Bresnan, J. (1982). The passive in lexical theory. In J. Bresnan

- (Ed.), *The Mental Representation of Grammatical Relations*. 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.
- Collins, M. J. (1996). A new statistical parser based on bigram lexical dependencies. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics*, pp. 184–191.
- Cruse, D. A. (1972). A note on English causatives. *Linguistic Inquiry* 3(4), 520–528.
- 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.
- 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*. Cambridge, MA: MIT Press.
- Juliano, C. and M. Tanenhaus (1994). A constraint-based lexicalist account of the subject/object attachment preference. *J. of Psycholinguistic Research* 23(6), 459–471.
- 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. Pearlmuter, 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.
- 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. and P. Merlo (1997). Lexical structure and parsing complexity. *Language and Cognitive Processes*.
- 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.