

Evidence for Gradient Saliency: What Happens with Competing Non-salient Referents during Pronoun Resolution?

Ralph L. Rose

Faculty of International Communication
Gunma Prefectural Women's University
Gunma, Japan
rose@gpwu.ac.jp

Abstract

The necessity of a gradient approach to saliency ranking of referents introduced in a discourse is evaluated by looking at (unbound) pronoun resolution preferences when there are competing non-salient referents. The study uses a sentence-completion technique in which participants had to resolve pronouns (“John sprayed the paint on the wall and then it ...”). Results suggest that a gradient saliency model is necessary. Syntactic and semantic prominence effects on pronoun resolution were also compared with results showing that semantic prominence (i.e., agent > patient) determined the saliency ranking of competing referents.

1 Introduction

A pervasive theme in theories of discourse coherence is the concept of saliency. It has proven to be a useful means of explaining how particular entities seem to receive some preferential treatment in both the production and perception of a discourse. For instance, it has long been observed that entities realized in certain structural positions (e.g., grammatical subject or first-mention) are preferred entities for topic continuation (Givón, 1983)—that is, they are preferentially referred to in the subsequent utterance. Similarly, pronominal reference to entities realized in certain structural positions (again, subject position, for example) is preferred to reference by repeated name (Gordon et al., 1993). In order to account for these observations, it has often been theorized that in the speaker's and hearer's mental representation of the discourse, these entities are salient (similar terms include focused or given).

To illustrate this line of thinking, consider (1). The pronoun in the second clause is preferentially interpreted as referring to LUKE rather than MAX. This has been observed in numerous psycholinguistic investigations (cf., Hudson-D'Zmura and Tanenhaus (1997); Mathews and Chodorow (1988); inter alia). In a simple saliency-based account, it is hypothesized that LUKE is a salient entity after the first clause and that reference to salient entities should be pronominal.

- (1) Luke_i hit Max_j and then he_{i/#j} ran home.

While many studies have investigated differences between pronominal reference to salient and non-salient entities, I have found no studies that have looked explicitly at what happens when a salient entity is not compatible with the pronoun, but more than one non-salient entity is. This is one of the main themes of the present study. Putting it as a question, what happens when there is competition among non-salient entities for pronoun interpretation? The answer to this question has some wider implications for how saliency is to be understood. In particular, the answer to this question leads to conclusions about whether theoretical models require a gradient model of saliency ranking or whether a categorical model is sufficient. In the following background section I will discuss this primary question further and introduce two related questions which must also be addressed. This will be followed by description of the experiment performed in this study. Briefly, results of the experiment are consistent with a gradient model of saliency ranking. Implications of these findings are discussed in the final section.

2 Background

2.1 Pronoun Reference Resolution

In this paper I will be focusing on the interpretation of unbound pronouns. Much has been written on this area of anaphora resolution and only a cursory overview is possible in this paper (see Hirst (1981) for a comprehensive overview of earlier work and Mitkov (2002) for an overview of more recent work). In this section, I will describe a generalized model of pronoun resolution and how salience plays a role in this process as well as discuss in some detail how salience is determined.

When interpreting pronouns in discourse, readers search a list of previously evoked entities in memory. Following Karttunen (1976) and Heim (1982), I will call these *discourse referents* (or just *referents*, for short). The list of discourse referents is ranked according to salience.

Two basic approaches may be taken to salience ranking: a categorical approach in which at most one referent is salient and all others are then, by definition, not salient; or a gradient approach in which referents are ranked along a salience continuum. In computational implementations of pronoun resolution algorithms, a gradient approach is often used, perhaps by necessity (cf., Lappin and Leass (1994)). However, psycholinguistic studies are often not so explicit about the approach taken and the results of most studies can be explained in terms of a categorical salience ranking. For instance, Gernsbacher and Hargreaves (1988) present a model of comprehension in which order-of-mention determines salience ranking, but their experimental evidence only compares first and second mentioned entities. In another case, Hudson-D’Zmura and Tanenhaus (1997) seek to verify the basic predictions of Centering Theory (Grosz and Sidner, 1986; Grosz et al., 1995), one aspect of which is a syntactic hierarchy: *subjects* > *objects* > *others*. However, their experimental evidence really only demonstrates a categorical ranking: *subjects* > *others*.

The difference between the categorical and gradient approaches is important in the present study because if salience is categorical, then pronominal reference should show no preference among non-salient entities. On the other hand, if salience is gradient, then it should be possible to observe preferences even among non-salient (or perhaps more accurately here, less salient) entities.

I should note however, that while it is clear that

those who take a gradient point of view must rule out a categorical point of view, I do not intend to imply that those studies that have a categorical approach (implied or otherwise) necessarily rule out a gradient approach. Many of those investigators may in fact be amenable to it. However, actual evidence of the necessity of a gradient approach in theory remains somewhat scarce. It is hoped that the present study will add to this evidence.

2.2 Computing Salience

Returning then to the model of pronoun resolution, the list of referents is first pruned to remove incompatible referents based on morphosyntactic features (Arnold et al., 2000; Boland et al., 1998). Then search for a referent should proceed with respect to the ranking (either categorical or gradient) of the referents in the list. But what determines this ranking? One of the most dominant factors has been shown to be syntactic prominence. However, in Rose (2005), I have argued that (in English, at least) syntactic and semantic information are often conflated. I will therefore discuss both of these factors below. In addition, another significant factor is coherence relations, also discussed this further below.

2.2.1 Syntactic Prominence

Several ways of determining the syntactic prominence of evoked entities have been discussed in the literature. These include left-to right ordering in which discourse referents introduced earlier are more prominent than those introduced later (Gernsbacher and Hargreaves, 1988), depth-first hierarchical search in which referents introduced higher and leftward in the syntactic tree are more prominent (Hobbs, 1978), and grammatical role in which referents introduced as subjects are more prominent than those introduced in other roles (Grosz et al., 1995). These different approaches typically make the same predictions when dealing with syntactically simpler constructions (i.e, no subordinate clauses or complex noun phrases), but may make different predictions with more complex constructions. The stimuli used in the experiment described below are all relatively syntactically simple and so in this paper I will not evaluate the differences among these various approaches.¹

¹See Rose (2005) for detailed discussion of these different approaches and a psycholinguistic and corpus linguistic comparison of hierarchical and grammatical role approaches.

However, for expository purposes, I will use grammatical role labels in discussion below.

When ranking referents according to the grammatical role in which they have been introduced, a categorical ranking predicts that the referent introduced as subject is salient and that any other referent is non-salient. This is the point of view implicitly taken in Stevenson et al. (2000), for example, in which they argue that a pronoun should refer to the referent introduced as subject of the preceding utterance.² A gradient salience approach, however, requires a more detailed prominence hierarchy such as that in (2). This is the point of view taken in most studies using the Centering framework (Grosz and Sidner, 1986; Grosz et al., 1995). An even more detailed gradient salience approach may be taken in which the points on the prominence hierarchy carry different (relative) weights. This is the approach taken in many practical applications such as the pronoun resolution algorithm of Lappin and Leass (1994).

(2) *subject > object > oblique > others*

2.2.2 Semantic Prominence

In English, syntactic role and semantic role are often conflated. That is, syntactic subjects are often semantic agents while syntactic objects are often semantic patients, and so on. Thus, it could be that the kind of pronoun resolution preferences previously observed and usually attributed to syntactic prominence effects might actually be attributed to semantic prominence effects. In other words, perhaps subject-preference is actually agent-preference.

In order to investigate this question, in Rose (2005), I used argument-reordering constructions: constructions which allow the order of arguments to vary—hence effecting a different relative syntactic prominence of discourse referents—with no (or minimal) change in their semantic role. For instance, so-called *psychological*-verbs have the alternate forms shown in (3)-(4).

(3) The audience admired the acrobats.

(4) The acrobats amazed the audience.

²More precisely, Stevenson et al. (2000), using the Centering framework (Grosz and Sidner, 1986; Grosz et al., 1995), argue that the backward-looking center, **Cb**, should refer to the subject of the preceding utterance. This is a simplification of the original Centering proposal in which it was suggested that the **Cb** refer to the highest-ranking member of the set of forward-looking centers in the previous utterance which is realized in the current utterance.

In (3), AUDIENCE is realized in a more syntactically prominent position than is ACROBATS: that is, in subject position. The reverse is true in (4). On the other hand, the semantic roles remain the same in both: ACROBATS is the stimulus while AUDIENCE is the experiencer. If syntactic prominence is most important, then a subsequent pronoun *they* should pick out AUDIENCE in (3) and ACROBATS in (4). On the other hand, if semantic prominence is most important, then a subsequent pronoun should pick out the same discourse referent in both alternatives. Assuming experiencer is higher on a semantic prominence hierarchy than stimulus (cf., thematic hierarchies in Jackendoff (1972); Speas (1990); inter alia), then this would be AUDIENCE.

In Rose (2005), I compared the effects of syntactic and semantic prominence on the salience of discourse referents in psycholinguistic experiments using two argument-reordering constructions: *tough*-constructions and *spray/load*-constructions. Results show that both syntactic and semantic prominence contribute to the salience of discourse referents. This suggests that experiments of this sort should carefully control for both syntactic and semantic prominence. In the present experiment, I do so by using an argument-reordering construction for the test stimuli.

2.2.3 Coherence Relations

Several investigators have theorized and observed that pronoun interpretation preferences differ when there is a causal connection between utterances compared to when there is a narrative connection (Hobbs, 1978; Kehler, 2002; Stevenson et al., 2000). For instance, in the narrative relation shown above in (1) (repeated below as (5)), the preference is for the pronoun to refer to LUKE. However, in (6) in which the utterances are related by a causal connection, the preference is for the pronoun to refer to MAX.

(5) Luke_i hit Max_j and then he_{i/#j} ran home.

(6) Luke_i hit Max_j because he_{#i/j} ran home.

Therefore, when investigating pronoun resolution, it is also necessary to take into account the influence of coherence relations by either controlling for these relations or making them another point of investigation. In the present study, I will take the latter course of action in order to see how coherence relations might influence pronoun resolution to competing non-salient entities. Previ-

ous accounts of the effects of coherence relations on pronoun resolution have taken the view that the kind of relation shifts attention to different aspects of the event being described (Stevenson et al., 1994; Stevenson et al., 2000). If an event has, for example, a start-state and an end-state, then a narrative relation will shift attention toward the start-state while a causal relation will shift attention toward the end-state. Subsequent pronominal reference will therefore prefer referents associated with these respective states, as illustrated in (5)-(6). Based on this argumentation, the prediction would be that pronominal reference might favor one non-salient referent over another if it is associated with that part of the event to which attention has been shifted by the coherence relation.

3 Experiment

Before describing the experiment, I'll review the primary and secondary questions which this experiment is designed to test. First, there is the question of what happens during pronoun resolution processes when there are competing non-salient referents. Answers to this question should provide evidence toward either a categorical or a gradient model of salience ranking. Furthermore, because investigating this question requires controlling for syntactic versus semantic prominence as well as coherence relation effects, two other secondary questions are also investigated. First, which is a more important factor in pronoun resolution: syntactic or semantic prominence? Second, what effect do coherence relations have on pronominal reference to non-salient entities?

3.1 Design

The research questions described above were investigated in this study using the well-known *spray/load*-constructions which exhibit the locative alternation (Levin, 1993) as shown in (7) and have synonymous alternative forms.³

- (7) a. John sprayed some paint on a wall.
b. John sprayed a wall with some paint.

³There is some difference of opinion on whether the two forms of *spray/load*-constructions are actually synonymous. One central point of contention is whether the totality effects on the direct object (i.e., the judgment that the entity in direct object position is totally used up in the event) are consistent across both forms. In the judgment of Rappaport and Levin (1988), the totality effect applies only with the *with*-variant. In contrast, it is my judgment (Rose, 2005) and also that of Tenny (1994, see her data items (100) and (102)) that the effect applies across both forms.

According to prominence hierarchies in which the syntactic subject or the semantic agent is most prominent, then JOHN should consistently be regarded as the (most) salient referent while PAINT and WALL should be regarded as less or non-salient referents in these sentences. Thus, subsequent pronominal reference with the third-person singular pronoun, *it*, allows a test of the three different questions outlined above.

First, if a categorical approach to salience is sufficient, then there should be no overall preference for either PAINT or WALL. But if gradient salience is necessary for ranking, then it might be possible to observe a difference between the two.

The nature of this difference, however, might be more complex depending on the way salience ranking is determined. If syntactic prominence is the only relevant factor, then preferences should consistently favor the object (i.e, PAINT in (7a), WALL in (7b)) according to the well-established syntactic prominence hierarchy in (2) above. But if semantic prominence is the only factor, then preferences should favor either the theme (PAINT) or the location (WALL) depending on how the semantic prominence hierarchy is ordered. One prediction might be based on proposed thematic hierarchies (cf., Larson (1988), Speas (1990)) which place theme above location. According to such a hierarchy, PAINT should be consistently preferred. This is what I observed in Rose (2005).

Other differences may result from the kind of coherence relation used. However, for *spray/load*-constructions, this is a little difficult to predict. The two non-salient entities are both arguably a part of the end-state of the event—that is, together, they are the product of the agent's work. Thus, any motivation to distinguish between the two with respect to the coherence relation must come from some other feature of the event or its participants. I will address the possibility in the discussion section below.

3.2 Method

3.2.1 Participants

The participants in this experiment included 36 undergraduate students at Morehead State University in Kentucky. Students were recruited through fliers and classroom announcements and received five dollars for their participation.

3.2.2 Materials

Twenty-four stimulus items were prepared using *spray/load* verbs as the matrix verb. The agent/subject was a commonplace proper name (12 male and 12 female) and the themes and locations were all inanimate nouns presented in indefinite form. The *spray/load* sentence was then followed by one of two connectives: *and then* to force a narrative relation or *because* to force a causal relation. These connectives were then followed immediately by *it*. Each stimulus was then followed by a blank line for participants to fill in a completion for the sentence. The experiment was therefore a 2×2 design pitting ORDER of entities (theme-location or location-theme) against coherence RELATION (narrative or causal). (8) shows an example of the four variants of one stimulus item.

- (8) a. John sprayed some paint on a wall
and then it _____
(theme-location, narrative)
b. John sprayed a wall with some paint
and then it _____
(location-theme, narrative)
c. John sprayed some paint on a wall
because it _____
(theme-location, causal)
d. John sprayed a wall with some paint
because it _____
(location-theme, causal)

Stimulus items were placed into twelve different tests such that each test contained only one variant of each item but conditions were balanced across all tests. The order of the items was pseudo-randomized such that consecutive items were not from the same experimental condition. The 24 items were combined with 101 items from an unrelated experiment to make a total of 125 items. Tests were printed in landscape orientation allowing every stimulus item to be followed by a blank line of at least three inches—ample space for participants to write their continuations.

3.2.3 Procedures

Participants were given the test forms and were asked to complete each sentence in the way that seemed most natural to them. Participants' responses were then analyzed and marked with one of four designators: If their completion showed that they interpreted the pronoun unambiguously as the theme of the *spray/load* verb then the re-

sponse was marked THEME. Similarly, if they interpreted the pronoun as the location, then the response was marked LOCATION. If the response was ambiguous as to the participant's interpretation, then it was marked INDETERMINATE. Finally, if the response indicated pronominal reference to some other entity, or the pronoun was taken as an empty pronoun, then the response was marked OTHER.

3.3 Results

In total, there were 836 usable responses (23 responses were left blank and 5 were ungrammatical). 130 responses were judged INDETERMINATE and 65 responses were judged OTHER. Only the remaining 641 responses are therefore used in the analysis below.

In order to evaluate the results, it is useful to look at the participants' pronoun resolution preferences. However, there are two ways of looking at these preferences: syntactically or semantically. Thus, while it is somewhat more laborious for the reader, I will present the results from these two perspectives for the sake of completeness. The results are therefore presented in terms of object-preference as well as theme-preference. Object preference is calculated as the total number of choices for the object minus the total number of choices for the oblique. Theme-preference, on the other hand is calculated as the total number of choices for the theme minus the total number of choices for the location. These results by subjects and by items are shown in Table 1 and Table 2, respectively.

The results show that there was an overall preference for the location (i.e., *wall*) in both variants. This can be most readily seen by noting the consistently negative theme-preference values in Table 2. This is underscored by the significant main effect for ORDER in the object-preference results in contrast with the nonsignificant main effect for ORDER in the theme-preference results. This contrast also indicates that in this experiment, participants' pronoun resolution processes were guided by a salience ranking determined by semantic prominence and not syntactic prominence.

As for the main question of categorical versus gradient salience, the results point toward a gradient model of salience ranking. Participants showed a clear, consistent preference for one non-salient entity (location) over another (theme).

Table 1: Overall Results for Object-preference by subjects

RELATION	ORDER	
	theme-location	location-theme
narrative	-1.75	2.50
causal	-0.89	1.14
Variant	$F(1, 35) = 58.2$	$p < 0.001$
Relation	$F(1, 35) < 1.0$	n.s.
Variant*Relation	$F(1, 35) = 8.4$	$p < 0.01$

RELATION	ORDER	
	theme-location	location-theme
narrative	-2.62	3.75
causal	-1.33	1.71
Variant	$F(1, 23) = 18.3$	$p < 0.001$
Relation	$F(1, 23) < 1.0$	n.s.
Variant*Relation	$F(1, 23) = 3.2$	$p = 0.085$

Table 2: Overall Results for Theme-preference by subjects

RELATION	ORDER	
	theme-location	location-theme
narrative	-1.75	-2.50
causal	-0.89	-1.14
Variant	$F(1, 35) = 2.8$	$p = 0.10$
Relation	$F(1, 35) = 8.4$	$p < 0.01$
Variant*Relation	$F(1, 35) < 1.0$	n.s.

RELATION	ORDER	
	theme-location	location-theme
narrative	-2.62	-3.75
causal	-1.33	-1.71
Variant	$F(1, 23) = 2.2$	$p = 0.15$
Relation	$F(1, 23) = 3.2$	$p = 0.085$
Variant*Relation	$F(1, 23) < 1.0$	n.s.

Finally, the results pertaining to coherence relations are somewhat inconclusive. In order to discuss this, it is better to refer to the theme-preference results because semantic prominence has proven to be the dominating factor here. While there is a significant main effect of RELATION by subjects, the effect is, at best, marginal by items. It is possible that a more thorough investigation with more items could yield a clear, significant result. On the other hand, even if the current effect is somehow real, it is actually quite weak. Note that the theme-preference values, which are negative in the narrative condition, are merely less negative in the causal condition—not enough to flip-flop resolution preferences. So, it seems difficult to make the case here that coherence relations shift these preferences in any meaningful way.

4 Discussion

In the present study, there were three questions under investigation. Let me review these three questions in turn and what the results say about them. First there was the primary question of categorical versus gradient approaches to salience ranking. The results here are not consistent with a categorical approach and clearly suggest a gradient approach. In this respect, the study lends psycholinguistic support to the many implementations of pronoun resolution algorithms which incorporate a gradient ranking of candidate referents for resolution (e.g., Kennedy and Boguraev (1996); Lappin and Leass (1994)).

However, just how fine-grained an approach is necessary is not conclusive from this investigation since competition among only two non-salient referents was tested. A more thorough study with stimuli including a large number of referents would be necessary to draw further conclusions about the necessity of a fine-grained gradient model of salience ranking.

The second question in this study was the question of whether syntactic prominence or semantic prominence is more important for determining the salience of referents. Results quite clearly point toward semantic prominence. These results contrast with those of Rose (2005) in two ways. The psycholinguistic results in that study suggest first that *both* syntactic and semantic prominence play a role in determining salience and second that theme is higher than location on the semantic prominence hierarchy.

The first contrast might be attributed to differences in experimental technique. The fact that participants in the present experiment had to compose a completion to each sentence means that they may have spent more time focusing on the semantic representation of the situation. This may have inflated the semantic prominence effects while attenuating syntactic prominence effects.

The second contrast, however, is somewhat more difficult to resolve. But once again, it may be useful to appeal to differences in the experimental technique. In the present study, the process of composing a sentence continuation for the events described by *spray/load* verbs would have required visualizing the event in a more vivid way than might be required for mere reading comprehension. If this visualization process were to require participants to anchor their visualizations through some fixed objects in the representation, this might naturally lead them toward paying closer attention to the location than the theme. Further testing will be required to evaluate this hypothesis and disambiguate these contrasting results.

Finally, the third question in this study dealt with the influence of coherence relations on pronoun resolution to competing non-salient referents. The present study did not test this in a manner comparable to previous studies since unlike those studies, both target referents were associated with the end-state of the event. Nonetheless, results showed a weak (but inconclusive) tendency to shift resolution preferences from location toward (but not to) theme. While more evidence would be necessary to confirm this to be a real effect, if it does turn out to be real then it would be a very interesting result. Assuming for the sake of argument that it is, then this might suggest that participants do not see the theme argument of *spray/load* verbs as part of the end-state of the event. To illustrate how this might be so, consider a couple of examples. Once a bucket of paint has been sprayed onto a wall, it takes on certain properties of the wall—for instance, its texture and size. Similarly, hay loaded onto a cart also takes on certain properties of the cart such as its size and shape. It might then be the case that the end-state of a *spray/load* event is more centrally focused on the location argument than on the theme argument because it is the location which determines many of the main properties of the end-state.

Before concluding, I would like to suggest some applications of these findings. Computational implementations of resolution algorithms that use an explicit salience ranking mechanism can be adapted to incorporate semantic prominence information as one of the contributors to a candidate's overall salience index (e.g., as in Lappin and Leass (1994)). However, even implementations that do not have an explicit salience ranking mechanism might still incorporate semantic prominence information. The coreference resolution system described in Soon et al. (2001) and its more knowledge-rich extension in Ng and Cardie (2002) classify NP pairs as coreferent or not based on constraints learned from an annotated corpus. These constraints are based on a number of features. While the Ng and Cardie system does incorporate a syntactic role feature (i.e., whether or not either NP in a pair is a subject), neither system incorporates a semantic prominence feature. It would be interesting to see if any further gains could be made in these systems by incorporating such a feature in future work.

5 Conclusion

The main aim of this paper has been to explore the question of whether a gradient model of salience ranking for candidate referents in pronoun resolution is necessary, or if a categorical model is sufficient. In this endeavor, two other questions have been addressed along the way: the influence of syntactic and semantic prominence on salience ranking of referents and the influence of coherence relations on pronoun resolution preferences. Results point toward the necessity of a gradient model of salience in which salience ranking is primarily determined by semantic information. Results were inconclusive regarding the influence of coherence relations. However, further work is necessary to confirm that this is the case.

Acknowledgments

I am indebted to Stefan Kaufmann and Michael Dickey for their guidance on this research project and also to Katy Carlson for her help on the experimental portion of this study. I also express my gratitude to three anonymous reviewers whose constructive comments on an earlier draft were most helpful.

References

- Jennifer Arnold, Janet Eisenband, Sarah Brown-Schmidt, and John Trueswell. 2000. The rapid use of gender information: evidence of the time course of pronoun resolution from eyetracking. *Cognition*, 76:B13–B26.
- Julie Boland, M. Acker, and L. Wagner. 1998. The use of gender features in the resolution of pronominal anaphora. Cognitive Science Technical Report 17, The Ohio State University Center for Cognitive Science.
- Morton Ann Gernsbacher and D. Hargreaves. 1988. Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language*, 27:699–717.
- Talmy Givón. 1983. Topic continuity in discourse: An introduction. In Talmy Givón, editor, *Topic Continuity in Discourse: A Quantitative Cross-Language Study*, pages 1–42. John Benjamins, Amsterdam.
- Peter Gordon, Barbara Grosz, and Laura Gilliom. 1993. Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, 17:311–347.
- Barbara Grosz and Candace Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12:175–204.
- Barbara Grosz, A. Joshi, and S. Weinstein. 1995. Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21:203–225.
- Irene Heim. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- Graeme Hirst. 1981. *Anaphora in Natural Language Understanding: A Survey*. Springer-Verlag, Berlin.
- Jerry Hobbs. 1978. Resolving pronoun references. *Lingua*, 44:311–338.
- Susan Hudson-D’Zmura and Michael Tanenhaus. 1997. Assigning antecedents to ambiguous pronouns: The role of the center of attention as the default assignment. In M. Walker, A. Joshi, and Ellen Prince, editors, *Centering Theory in Discourse*, pages 199–226. Clarendon Press, Oxford.
- Ray Jackendoff. 1972. *Semantic Interpretation in Generative Grammar*. MIT Press, Cambridge, MA.
- Lauri Karttunen. 1976. Discourse referents. In James McCawley, editor, *Syntax and Semantics, Vol. 7: Notes from the Linguistic Underground*, pages 363–385. Academic Press, New York.
- Andrew Kehler. 2002. *Coherence, Reference, and the Theory of Grammar*. CSLI Publications, Stanford University, CA.
- Chris Kennedy and Branamir Boguraev. 1996. Anaphora for everyone: Pronominal anaphora resolution without a parser. In *Proceedings of the 16th International Conference on Computational Linguistics (COLING ’96)*, pages 113–118, Copenhagen, Denmark.
- Shalom Lappin and Herbert Leass. 1994. An algorithm for pronominal anaphora resolution. *Computational Linguistics*, 20:535–561.
- Richard Larson. 1988. On the double object construction. *Linguistic Inquiry*, 19:335–392.
- Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press, Chicago, IL.
- Alison Mathews and Martin Chodorow. 1988. Pronoun resolution in two-clause sentences: Effects of ambiguity, antecedent location, and depth of embedding. *Journal of Memory and Language*, 27:245–260.
- Ruslan Mitkov. 2002. *Anaphora Resolution*. Longman, London.
- Vincent Ng and Clarie Cardie. 2002. Improving machine learning approaches to coreference resolution. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 104–111, Philadelphia, PA.
- Malka Rappaport and Beth Levin. 1988. What to do with θ -roles. In Wendy Wilkins, editor, *Syntax and Semantics 21: Thematic Relations*, pages 7–36. Academic Press, New York.
- Ralph Rose. 2005. *The Relative Contribution of Syntactic and Semantic Prominence to the Salience of Discourse Entities*. Ph.D. thesis, Northwestern University.
- Wee Meng Soon, Hwee Tou Ng, and Daniel Chung Yong Lim. 2001. A machine learning approach to coreference resolution of noun phrases. *Computational Linguistics*, 27:521–544.
- Margaret Speas. 1990. *Phrase Structure in Natural Language*. Kluwer Academic Publishers, Dordrecht, Netherlands.
- Rosemary Stevenson, Rosalind Crawley, and David Kleinman. 1994. Thematic roles, focus and the representation of events. *Language and Cognitive Processes*, 9:519–548.
- Rosemary Stevenson, Alistair Knott, Jon Oberlander, and Sharon McDonald. 2000. Interpreting pronouns and connectives: Interactions among focusing, thematic roles and coherence relations. *Language and Cognitive Processes*, 15(3):225–262.
- Carol Tenny. 1994. *Aspectual Roles and the Syntax-Semantics Interface*. Kluwer Academic Publishers, Dordrecht, Netherlands.