Lexical and Grammatical Role Constraints in Resolving *Other*-Anaphora

Natalia N. Modjeska
Division of Informatics
University of Edinburgh
natalian@cogsci.ed.ac.uk

Abstract

Computational research on reference resolution has mostly focused on pronouns and, to a lesser extent, on definite descriptions. We present a computational treatment of *other*-anaphora, i.e., noun phrases modified by "other" and "another". The approach presented here is a lexical method LEX that finds antecedents on the basis of information in WordNet. LEX is compared with the state-of-the-art pronoun resolution method based on Centering Theory. A performance comparison between the two algorithms indicates that LEX is superior to the Centering method (48% vs. 37% success rate). Furthermore, an analysis of examples that the Centering method resolved correctly and LEX did not suggests the following: to improve the performance of LEX requires a better treatment of metonymy, bridging, implicit antecedents and non-NP antecedents. Grammatical role, however, does not seem to play a significant part in resolving *other*-anaphora, other than with pronominal antecedents.

1 Introduction and scope

The primary focus of reference resolution systems has been on pronouns and, to some extent, definite NPs; other anaphoric devices, notably NPs modified by comparative modifiers, e.g., "other", "another", "such", "similar", and "same", have not been addressed, despite their high frequency in text and speech. For instance, a corpus study of "other" and "another" in the British National Corpus (http://info.ox.ac.uk/bnc) showed that comparative modifiers are in the top 200 most frequent words, ranging from rank 75 for "other" (tagged as adjective) to rank 159 for "another" (tagged as determiner) (Modjeska, 2000). Bierner (2001) quotes frequency data for a number of comparative words in tutorial dialogues. In 269 such dialogues, each dialogue contained, on average, 3.65 comparative phrases. "Other" was by far the most frequent comparative word, with 686 occurrences of the total 983 comparative phrases (70%).

That comparative modifiers are *anaphoric*[1] is clear from the following example:

(1) Fujitsu and NEC said they were still investigating, and that knowledge of more such bids could emerge ... **Other major Japanese computer companies contacted yesterday** said they have never made such bids.

In (1), the NP "other major Japanese computer companies" (and the subsequent pronoun "they") refers to a set of entities characterised as "major Japanese computer companies", excluding "Fujitsu" and "NEC". Thus, in order to resolve the reference of "other major Japanese computer companies", one must identify the antecedent of the *other*-NP (here, the split antecedent "Fujitsu" and "NEC") and *exclude* it from the referential scope of the phrase. Without the antecedent, the NP "other major Japanese computer companies" cannot be correctly interpreted.

An account of the semantics of comparative modifiers, cast in terms of *alternative sets* (Rooth, 1992) was presented by Bierner (2001). Our work focuses on design, implementation and empirical evaluation of a computational algorithm to resolve them. In particular, we focus on the modifiers "other" and "another", specifically, on their anaphoric uses, excluding idiomatic expressions (e.g., "the other week"), reciprocal "each other" and "one another", and the discourse connectives "on the other hand" and "in other words". Also excluded are the *easy* anaphoric cases: elliptic constructions "one X ... the other(s)" and "one X ... another"; substitutions "the other one" and "another one"; comparatives "Xs other than Ys" and "other Xs than Ys"; and conjoined *other*-NPs (joined by "and", "or", "but", "as well as", and "along with"). In these constructions, the antecedent of "(an)other" is either available structurally, as is the case with *other than* constructions and conjoined *other*-NPs[2], or there are clues that point to the antecedent, e.g., "one" in elliptic constructions.[3]

[1] In the sense of Carter (1987), who views anaphora as "the special case of cohesion where the meaning (sense and/or reference) of one item in a cohesive relationship (the **anaphor**) is, in isolation, somehow vague and incomplete, and can only be properly interpreted by considering the meanings of the other item(s) in the relationship (the **antecedent(s)**)."

[2] There are examples of conjoined "other" in which the left conjunct is not its proper antecedent, e.g., "Most dogs live for about 10 years on average, and during their lives they will come into contact with possibly hundreds of people and **other dogs** ...". But such examples are rare.

[3] For a computational treatment of conjoined "other" and other constructions with structurally given antecedents in the context of an information retrieval system, see (Bierner, 2001).

2 The rationale

The experiments presented in this paper were inspired by two remarks. First, it has been noted (see, for instance (Strube and Hahn, 1999)) that nominal anaphora is far more constrained by conceptual criteria than pronominal anaphora. Second, even in pronoun resolution, researchers have claimed that resolving pronoun references on the basis of syntactic information only is naive, and that pronouns would ideally be resolved by a combination of syntax and semantics (Tetreault, 2001). For a full set of factors known to affect pronoun resolution, see chapter 18 in (Jurafsky and Martin, 2000). Implemented pronoun resolution systems tend to rely heavily on recency and grammatical role (Brennan et al., 1987; Lappin and Leass, 1994; Tetreault, 2001) or informational status of discourse referents (Strube and Hahn, 1999).

Current approaches to *nominal* anaphora, e.g., Vieira and Poesio (2000) and Harabagiu and Maiorano (1999), operate on a set of *heuristics* that combine semantic information from WordNet (Fellbaum, 1998) with recency, information about named entity types, and rules for certain syntactic constructions (e.g., appositions).

To our knowledge, no empirical work has been done to *systematically* examine the extent to which some factors — i.e., lexical semantics and grammatical role — *independently* contribute to resolving either pronominal or nominal anaphora. An attempt was made by Harabagiu (1998), who first resolved definite descriptions with a Centering algorithm, and then applied a WordNet-based lexical approach to examples that the Centering algorithm resolved incorrectly. Combining the two knowledge sources resulted in improved resolution of definite NPs in her corpus. Vieira and Poesio (2000), having examined the performance of their approach on bridging definite descriptions, concluded that a focusing mechanism would be needed to constrain the search. Our work fills this gap, focusing on *other*-anaphora. We confront (i) our lexical method that finds antecedents on the basis of information in WordNet with (ii) the state-of-the-art pronoun resolution algorithm based on grammatical role ranking. Both algorithms also use recency and syntactic constraints on the antecedent of "other", which we identify in the next section.

2.1 Syntactic constraints on the antecedent of *other*

Syntactic positions that can NOT serve as antecedents of "(an)other":

Apposition: (a) NP preceeding an appositive, if the appositive contains "(an)other"; (b) an appositive clause of the *other*-NP:

(2) a. Mary Elizabeth Ariail, **another social-studies teacher**

b. **The other social studies teacher**, Mary Ariail, (both not "other than Mary Elizabeth Ariail")

Copular sentences: (a) subject NP of a copular clause, if "other" is the predicative NP; (b) the predicative NP if *other*-NP is the subject:

(3) a. The reputed wealth of the Unification Church is **another matter of contention**.

b. **The other matter of contention** is the reputed wealth of the Unification Church.
(both not "other than the wealth")

Possessives S/OF: (a) the possessor NP, if "other" realises the possessed entity; (b) the posessive PP complement of an *other*-NP:

(4) a. Koito's **other shareholders**

b. **other shareholders** of Koito
(both not "other than Koito")

Constructions with locative and temporal "there":

(5) a. In London, there are **other locations** where we could meet. (not "other than London")

b. On Tuesday, there are **other times** when we could meet. (not "other than Tuesday")

3 The LEX algorithm

Our previous work (Modjeska, 2000) identified five types of systematic lexical relations between anaphoric *other* and its antecedents:
Same predicate, i.e., the same head noun is used to evoke both the antecedent and the anaphor:

(6) Employers can pay the subminimum for 90 days ... to workers with less than six months of job experience, and for **another 90 days** if the company uses a government-certified training program for the young workers.

Hypernymy, e.g., (1) and (7):

(7) Mr. Stoll draws his title from the cuckoo's habit of laying eggs in the nests of **other birds.**

Re-description is a relation such that a class description evoked by "other" associates the antecedent with a different (but compatible) class than the one to which it is known to belong:

(8) A marketing study indicates that Hong Kong consumers are the most materialistic in the 14 major markets where the survey was carried out. The study by the Backer Spielvogel Bates ad agency also found that the colony's consumers feel more pressured than those in any of **the other surveyed markets** ...

Metonymy (9) and **metaphor** (10):

(9) Moscow has settled pre-1917 debts with **other countries** in recent years at less than face value.

(10) The human memory, in common with **every other store**, has to be positively consulted before it will function.

Bridging:

(11) While this court ruling was only on Hammersmith, it will obviously be very persuasive in **other cases of a similar nature**.

In (11), the antecedent is the NP "this court ruling", while the entity to be excluded is "the (legal) case in which this court ruling was made". The interpretation of "other cases" requires a *contextual parameter* which is derived from the antecedent (Webber et al., submitted).

The LEX approach presented here handles *same predicate* and *hypernymy* relations. It consists of three modules: LEX1, LEX2 and NEM (Named Entity Module). LEX1 matches the head noun of the anaphor with the head noun of all other NPs in the sentence, and it reports a match if the anaphor and antecedent candidate are evoked by the same predicate. LEX2 extracts from WordNet a hyponym tree for the anaphor and matches this tree with the antecedent candidates. NEM asumes that all named entities have been identified and classified. It then uses simple heuristics to find the antecedent of "other". For this experiment, named entities were annotated manually, with one of the following tags:

- TIME, e.g., "April 30, 1956", "Wednesday";
- PERSON, e.g., "Wilbur Ross Jr.";
- ORG, e.g., "IBM";
- PRODUCT, e.g., "Thunderbird", "Leche Fresca";
- MONEY, e.g., "$ 101 million";
- NUM (neither money nor time), e.g., "45", "3 1/5";
- LOC, e.g, "Aslacton";
- MISC, e.g., "Great Depression"; "the Republicans".

The named entity classification schema we used follows MUC-7 Named Entity Task Definition (Chinchor, 1997), with several exceptions, of which we mention two. Unlike MUC, we allow for nested expressions, e.g., "U.S.A." (tagged LOC) in "Campbell U.S.A." below, needed to resolve such examples.

(12) The way that we've been managing Campbell U.S.A. can hopefully spread to **other areas of the company**. (i.e., "other areas than the U.S.A. branch of Campbell")

We annotated titles and roles, and non-organizational entities (with the tag MISC), and adjectival forms of locations, e.g, "Japanese", (with the tag LOC), to resolve examples with adjectival antecedents:

(13) He chastised the media for paying such close attention to Japanese investment when **other foreign countries** ... are acquiring more American assets.

The heuristics that NEM uses to identify the antecedent of "other" are as following:

- If the head noun of the *other*-NP is the word "year", "month", "week", "day", "time", in singular or plural, propose as antecedent an NP tagged TIME;
- If the head noun is "product", "wares", "merchandise", "goods", "commodity" or "service" (or their plural forms), propose an NP tagged PRODUCT;
- If the head noun is "million", "thousand", "dollar", "yen", "pound", propose an NP tagged MONEY;

With persons, organisations, and locations, it is hard to come up with a full list of predicates. Instead, we used the WordNet lexical hierarchy and the following rules:

- If the head noun of *other*-NP has synset "location" among its hypernyms, suggest as antecedent an NP tagged LOC;
- If the synset is "person, individual", suggest an NP tagged PERSON;
- If the synset is "organisation", the antecedent is a named entity tagged ORG.[4]

The LEX method is *incremental*; it attempts to resolve the anaphoric references intra-sententially, as soon as the anaphor is encountered, before considering entities in the previous sentence. The three resolution modules are applied in the following order: LEX1 on current sentence, NEM on current sentence, LEX1 on previous sentence, and NEM on previous sentence. LEX2 (hypernymy) applies only if LEX1 and NEM both failed on both current and previous sentences. Like the modules before it, LEX2 applies first to the sentence containing the anaphor, and then to the sentence preceding it.

4 The Centering algorithm

The Centering algorithm we used is an adaptation of Tetreault's (2001) Left-Right Centering, currently the state of the art in pronoun resolution. LRC is built upon Centering Theory's constraints and rules (Grosz et al., 1995) as implemented by Brennan et al. (1987), hereafter BFP, with a few modifications. In particular, BFP makes no provision for incremental resolution of pronouns, while it is well established that humans process utterances one word at a time. This lack of incrementality is amended in LRC, which first searches the current sentence, and if no antecedent is found, then the previous sentences are searched left-to-right. The second difference is that LRC does not use Rule 2 of Centering Theory, pertaining to ranking of transition states.

Our implementation of LRC differs from the original algorithm in two aspects. LRC keeps track of all utterances processed so far (and their forward- and backward-looking centers), and, if the antecedent is not found in the current sentence, it searches all previous Cf-lists, one sentence at time, starting with the previous one. We restrict the search space to two sentences (current and previous), and ignore criteria such as gender and number, which do not appear relevant to resolving *other*-NPs.

[4] If the head noun of the *other*-NP has both "location" and "organisation" as hypernyms, e.g., "country" in sense 1 and 2, NEM proposes as antecedents entities tagged ORG and entities tagged LOC.

5 The corpus

We used 189 two-sentence samples of *other*-anaphora that we extracted from the WSJ corpus (Penn Treebank release 2, directories 00-02).

The samples were preprocessed in the following fashion. We extracted all nested NPs and their head nouns from each sentence of each sample. To avoid deciding which nouns are used as modifiers and which are part of a compound noun, we treated as compound nouns all strings with proper and common noun tags. A WordNet-lookup module (in the LEX experiment) would then first look up the whole string "term1 term2 term3". If the string was not found, the script recursively stripped off the leftmost term ("term1") and looked up both "term1" (to resolve examples such as "a woman ringer ... another woman") and the remaining string ("term2 term3"). In the next stage of corpus preprocessing, we removed NPs that did not conform to syntactic constraints described in section 2.1. For the LEX experiment, named entities were classified according to the scheme described in section 3. The order of NPs was randomized within each sentence, to prevent surface order from enhancing LEX's performance, i.e., in sentences which, in addition to the antecedent, also contain a distractor entity, e.g., (16).

The algorithms's performance was compared with human judgements in a gold standard corpus, annotated by the author. We used the same samples for algorithm development and evaluation, and to create the gold standard. A large-scale blind evaluation of LEX is being planned.

6 Comparing LEX and Centering

The scores we present are for the 124 cases with explicitly evoked NP antecedents within a two-sentence window, the two conditions that both algorithms are designed to handle. The remaining 65 instances, discussed in section 6.3, involve (i) non-NP antecedents, (ii) antecedents outside the window scope, and (iii) examples classified as undecidable in the gold standard corpus.

6.1 Results for LEX

LEX resolved correctly 45 to 59 instances of "(an)other" (35% to 48%). The first score is a strict score, the second is a lenient score. When scoring leniently, we counted as success the following cases:

- Split antecedents, if the algorithm found at least one antecedent (and the others were evoked outside the two-sentence window).[5]
- Cases in which the correct antecedent is a pronoun and subject of a copular sentence, and LEX resolved the antecedent to the predicative NP:

  (14) This is a company that has invested in capacity additions more aggressively than **any other company in the industry** ... (i.e., "any other company than this")

- Cases with the antecedent most recently mentioned in the current sentence, which LEX resolved to the same discourse referent in the previous sentence:

  (15) The computer can process 13.3 million calculations called floating-point operations every second. The machine can run software written for **other Mips computers**, the company said.

- Cases that require full text understanding and reasoning, in order to distinguish the correct antecedent from a distractor (16).

  (16) Integra, which owns and operates hotels, said that Hallwood Group Inc. has agreed to exercise any rights that aren't exercised by **other shareholders**. (Both "Integra" and "Hallwood" qualify as antecedent, because both are ORG)

Of the 65 samples which LEX resolved incorrectly or not at all, 17 are metonymies (though some metonymies, e.g., (9), were resolved successfully); at least four are bridging examples. Nine errors are due to WordNet omissions ("markdown", "Walkman", also examples (17) and (18)), or anaphor sense ambiguity (7), sometimes in combination with the application order of the resolution modules or the heuristics of NEM.

(17) designer's age ... **other risk factors** (age is not recorded in WordNet as a risk factor)

(18) cotton and corn ... **other crops**

The noun "crop" has three senses in WordNet: (i) the yield from plants in a single growing season (ii) the stock or handle of a whip; and (iii) a pouch in many birds and some lower animals that resembles a stomach for storage and preliminary maceration of food. In (18), "crops" is used in the same sense as "plants".

Example (7) is a classic example of hypernymy, but sense 3 of "bird" is "dame, doll, wench, skirt, chick, bird (informal terms for a (young) woman)", and because (i) this sense has a hypernym "person, individual", and (ii) we resolve named entities before hypernymy, LEX resolved "other birds" to "birds other than Mr. Stoll".

Of the eight samples with pronominal antecedents in our corpus, five were resolved successfully, if lenient scoring was applied. The remaining three were resolved incorrectly. This is a principled error; pronouns do not contain enough information for a lexical algorithm.

Another type of principled error is shown by (19):

(19) the question of investors' access to the U.S. and Japanese markets ... **other important economic issues**

The noun "issue" (as in "economic issues") is a very general concept that could refer to a variety of discourse entities. Other such general concepts are "thing", "alternative", and "factor" , nine in total in our corpus. Additionally, four "other" are phrases "among other things"; they should, perhaps, be treated as idioms:

[5] LEX successfully resolves split antecedents if they are given within the two-sentence window.

(20) That designation would, **among other things**, provide more generous credit terms under which the Soviets could purchase grain.

Some errors stem from the heuristics of NEM, sometimes combined with the named entity annotation:

(21) The ruling could lead to the cancellation of huge bank debts the London Borough of Hammersmith and Fulham ran up after losing heavily on swap transactions. As many as 70 U.K. and international banks stand to lose several hundred million pounds should the decision be upheld and set a precedent for **other municipalities**.

In (21), "U.K." in "U.K. and international banks" was tagged LOC, which forced NEM to choose that NP as the antecedent over "the London Borough of Hammersmith and Fulham" in the preceding sentence.

Example (22) illustrates errors caused by the application order of the resolution modules:

(22) PaineWebber Inc., for instance, is forecasting growth in S&P 500 dividends of just under 5% in 1990, down from an estimated 11% this year. In **other years** in which there have been moderate economic slowdowns — the environment the firm expects in 1990 — the change in dividends ranged from a gain of 4% to a decline of 1% ...

Because we first try to resolve the anaphor by looking for the same predicate, and only if none is found do we apply NEM, the antecedent of "other years" was resolved to "this year", while the correct choice should have been "1990". Note that the relative clause "in which there have been moderate economic slowdowns — the environment the firm expects in 1990" would help to resolve this example correctly, but that would require deeper text understanding techniques than those currently available.

Finally, there are three samples with constructions "other types/kinds/forms of X", in which the most informative component is the head noun of the prepositional phrase. During corpus preprocessing, we extracted only nested NPs (e.g., "other types" rather than "other types of watches"), and this information was lost. Related to this type of error are errors caused by bugs in the various preprocessing and resolution modules, which account for 7 unresolved or incorrectly resolved cases of "(an)other".

6.2 Results for the Centering algorithm

The Centering method successfully resolved 45 to 46 occurencies of "(an)other" (36% to 37% of all cases with explicitly evoked NP antecedents).

6.3 What neither algorithm could handle

First, 11 samples of "other" in our corpus were classified as *undecidable* in the gold standard corpus. (The annotator could not unambiguously identify the antecedents of these expressions.) Since a human annotator could not resolve these examples, we do not expect that a computational method would be able to resolve them either.[6]

Of the remaining errors, 30 examples could not be resolved because both methods focus on explicitly realised NP antecedents: 19 samples with implicitly realised antecedents, e.g., (23), and 11 samples in which the antecedents were evoked by a VP (24), sentence, clause (25), text segment, or verbal form (26):

(23) "What's he doing?" hissed my companion, who was **the only other English-speaking member of the convention** ...

(24) ... they consented to findings that they had inaccurately represented the firm's net capital, maintained inaccurate books and records, and made **other violations**.

(25) "The fact of the matter is, I am a marketer. That's **another reason** [for the Backer Spielvogel job]".

(26) Such individuals ... having borrowed from the bank, they may continue to use it [the bank] for **other services**.

Neither algorithm could handle 20 cases of "(an)other" whose antecedents are evoked outside of the two-sentence window. Extensions of both methods could perhaps handle some of these cases; in 15 cases the antecedent is given within a five-sentence window. (Vieira and Poesio (2000), for example, used a five-sentence window in resolving references of *bridging* definite descriptions.[7])

Since both methods attempt to resolve "other" as soon as possible, cataphoric example such as (27) cannot be handled by the current versions of the algorithms.

(27) What is **another name** for the Roman numeral IX?

Finally, three samples contained more than one "other"; such examples are currently out of scope for both methods; and one example was an idiom that hasn't previously been excluded — "put another way".

6.4 LEX vs. Centering

Together, LEX and the Centering method correctly resolved 77 occurencies of "(an)other", or 63% of all cases with explicitly evoked NP antecedents within a two-sentence window. Of the samples that both methods resolved correctly, LEX resolved 77%. Samples that the Centering approach resolved correctly and LEX did not involve metonymy, bridging, and pronominal antecedents. This, in combination with performance scores, suggests two observations. First, the algorithms

[6] Some of these examples are expressions "in other news" and "in other commodity markets", which seem to function as genre-specific discourse connectives.

[7] They define bridging as "definite descriptions that either (i) have an antecedent denoting the same discourse entity, but using a different head noun (as in "house ... building") or (ii) are related by a relation other than identity to an entity already introduced in the discourse." Please note that Vieira and Poesio's definition of bridging would cover *all* uses of *other*-anaphora; we reserve the term for examples like (11).

overlap to some extent. Second, since the Centering method resolves the antecedent to an entity with the highest grammatical role ranking (usually the subject) and fails to produce a correct antecedent in almost 2/3 of all "other" cases, the antecedent of "other" is evoked by a less salient entity in 2/3 of the samples in our corpus (if saliency is mainly determined by the grammatical role of a discourse referent). A resolution algorithm that relies primarily on grammatical role ranking will therefore tend to show poor performance (or inefficiency, if supplemented, e.g., with lexical constraints, as pronoun resolution algorithms use gender, person, and number agreement), because grammatical role does not seem to play a significant part in resolution of *other*-anaphora. The performance scores for LEX, however, show that a lexical method can efficiently and successfully resolve antecedents of "other". To improve the performance of LEX, we need to incorporate a systematic treatment of metonymy, bridging, examples with implicitely evoked antecedents, non-NP antecedents, and to correct the errors listed in section 6.1.

7 Conclusion and further directions

We presented a lexical method LEX, which resolves anaphoric references of comparative modifiers "other" and "another" on the basis of information in WordNet. The algorithm is further augmented with recency and syntactic constraints on antecedent realisation, identified in this paper. LEX is the first computational method for resolving *other*-anaphora.

The performance of LEX was compared with the state-of-the-art pronoun resolution algorithm based on Centering Theory. The Centering algorithm we used employs grammatical role as a major determinant in reference resolution. The Centering algorithm showed performance inferior to LEX (37% vs. 48% success rate).

By comparing the two algorithms, we were able to examine the extent to which two factors known to affect salience of discourse entities — grammatical role and lexical semantics — contribute to the resolution of *other*-anaphora. Our results seem to show that grammatical role does not play a significant part in the resolution of "other". Lexical semantics, on the other hand, does.

In the next step, we intend to improve the LEX method by incorporating a principled treatment of metonymies, bridging examples, implicit and non-NP antecedents, and to run a large-scale blind evaluation.

Acknowledgements

The author is supported by grant No. GR/M75129 from the Engineering and Physical Sciences Research Council to the University of Edinburgh. I would like to thank Bonnie Webber, Graeme Hirst, and Katja Markert for valuable comments on the ideas, experimental design, and the contents of this paper; and Mark Chignell at the University of Toronto, who provided a good work environment in the Interactive Media Lab.

References

G. Bierner. 2001. Alternative phrases and natural language information retrieval. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*, Toulouse, France, July.

S. E. Brennan, M. W. Friedman, and C. J. Pollard. 1987. A centering approach to pronouns. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, pages 155–162.

D. Carter. 1987. *Interpreting anaphors in natural language texts*. Ellis Horwood Limited, John Wiley and Sons, Chichester.

N. Chinchor. 1997. MUC-7 named entity task definition. In *Proceedings of the 7th Conference on Message Understanding*, Washington, DC.

C. Fellbaum, editor. 1998. *WordNet: An Electronic lexical database*. The MIT Press, Cambridge, Mass.

B. J. Grosz, A. K. Joshi, and S. Weinstein. 1995. Centering: A framework for modeling the local coherehce of discourse. *Computational Linguistics*, 21(2):203–225.

S. Harabagiu and S. Maiorano. 1999. Knowledge-lean coreference resolution and its relation to textual cohesion and coherence. In *Proceedings of the ACL-99 Workshop on the relation of discourse/dialogue structure and reference*, Univ. of Maryland, June.

S. Harabagiu. 1998. *WordNet-based inference of textual context, cohesion and coherence*. Ph.D. thesis, University of Southern California.

D. Jurafsky and J. H. Martin. 2000. *Speech and language processing: an introduction to natural language processing, computational linguistics, and speech recognition*. Prentice Hall, Upper Saddle River, N.J.

S. Lappin and H. J. Leass. 1994. A syntactically based algorithm for pronominal anaphora resolution. *Computational Linguistics*, 20(3):535–561.

N. N. Modjeska. 2000. Towards a resolution of comparative anaphora: A corpus study of "other". Talk presented at PAPOCOL, 1st Postgraduate Linguistic Conference, Pavia, Italy, 24-25 November.

M. Rooth. 1992. A theory of focus interpretation. *Natural Language Semantics*, 1:75–116.

M. Strube and U. Hahn. 1999. Functional centering — grounding referential coherence in information structure. *Computational Linguistics*, 25(3):309–344.

J. R. Tetreault. 2001. A corpus-based evaluation of centering and pronoun resolution. *Computational Linguistics*, 27(4):507–520.

R. Vieira and M. Poesio. 2000. An empirically-based system for processing definite descriptions. *Computational Linguistics*, 26(4):539–593.

B. Webber, A. Joshi, M. Stone, and A. Knott. (submitted). Anaphora and discourse semantics.