Recognizing Referential Links: An Information Extraction Perspective

Megumi Kameyama

Artificial Intelligence Center

SRI International

333 Ravenswood Ave., Menlo Park, CA 94025, U.S.A.

megumi@ai.sri.com

Abstract

We present an efficient and robust reference resolution algorithm in an end-to-end state-of-the-art information extraction system, which must work with a considerably impoverished syntactic analysis of the input sentences. Considering this disadvantage, the basic setup to *collect*, *filter*, then *order by salience* does remarkably well with third-person pronouns, but needs more semantic and discourse information to improve the treatments of other expression types.

Introduction

Anaphora resolution is a *component technology* of an overall discourse understanding system. This paper focuses on reference resolution in an information extraction system, which performs a partial and selective ‘understanding’ of unrestricted discourses.

Reference Resolution in IE

An information extraction (IE) system automatically extracts certain predefined target information from real-world online texts or speech transcripts. The target information, typically of the form “who did what to whom where when,” is extracted from natural language sentences or formatted tables, and fills parts of predefined template data structures with slot values. Partially filled template data objects about the same entities, entity relationships, and events are then merged to create a network of related data objects. These template data objects depicting instances of the target information are the raw output of IE, ready for a wide range of applications such as database

updating and summary generation.¹

In this IE context, reference resolution takes the form of *merging* partial data objects about the same entities, entity relationships, and events described at different discourse positions. Merging in IE is very difficult, accounting for a significant portion of IE errors in the final output. This paper focuses on the referential relationships among *entities* rather than the more complex problem of *event* merging.

An IE system recognizes particular target information instances, ignoring anything deemed irrelevant. Reference resolution within IE, however, cannot operate only on those parts that describe target information because anaphoric expressions within target linguistic patterns may have antecedents outside of the target, and those that occur in an apparently irrelevant pattern may actually resolve to target entities. For this reason, reference resolution in the IE system needs access to *all* of the text rather than some selective parts. Furthermore, it is reasonable to assume that a largely *domain-independent* method of reference resolution can be developed, which need not be tailored anew each time a new target is defined.

In this paper, I discuss one such entity reference resolution algorithm for a general geo-political business domain developed for SRI’s FASTUSTM system (Hobbs et al., 1996), one of the leading IE systems, which can also be seen as a representative of today’s IE technology.

¹The IE technology has undergone a rapid development in the 1990s driven by the series of Message Understanding Conferences (MUCs) in the U.S. government-sponsored TIPSTER program (<http://www.tipster.org>).

The Input to Reference Resolution

Multiple top-scoring sites working on IE have converged on the use of finite-state linguistic patterns applied in stages of smaller to larger units. This finite-state transduction approach to IE, first introduced in SRI's FASTUS, has proven effective for real-world texts because full parsing is far too ambiguous, slow, and brittle against real-world sentences. This means that we cannot assume correct and full syntactic structures in the input to reference resolution in a typical IE system. The input is a set of (often overlapping or discontinuous) finite-state approximations of sentence parts. We must *approximate* fine-grained theoretical proposals about referential dependencies, and adapt them to the context of sparse and incomplete syntactic input.

The input to reference resolution in the theoretical literature is assumed to be fully parsed sentences, often with syntactic attributes such as grammatical functions and thematic roles on the constituents (Webber, 1978; Sidner, 1979; Hobbs, 1978; Grosz, Joshi, and Weinstein, 1995). In implemented reference resolution systems, for pronoun resolution in particular, there seems to be a trade-off between the completeness of syntactic input and the robustness with real-world sentences. In short, more robust and partial parsing gives us wider coverage, but less syntactic information also leads to less accurate reference resolution. For instance, Lappin and Leass (1994) report an 86% accuracy for a resolution algorithm for third-person pronouns using fully parsed sentences as input. Kennedy and Boguraev (1996) then report a 75% accuracy for an algorithm that approximates Lappin and Leass's with more robust and coarse-grained syntactic input. After describing the algorithm in the next section, I will briefly compare the present approach with these pronoun resolution approaches.

Algorithm

This algorithm was first implemented for the MUC-6 FASTUS system (Appelt et al., 1995), and produced one of the top scores (a recall of 59% and precision of 72%) in the MUC-6 Coreference Task, which evaluated systems' ability to recog-

nize coreference among noun phrases (Sundheim, 1995). Note that only *identity* of reference was evaluated there.²

The three main factors in this algorithm are (a) accessible text regions, (b) semantic consistency, and (c) dynamic syntactic preference. The algorithm is invoked for each sentence after the earlier finite-state transduction phases have determined the best sequence(s) of nominal and verbal expressions. Crucially, each nominal expression is associated with a set of template data objects that record various linguistic and textual attributes of the referring expressions contained in it. These data objects are similar to *discourse referents* in discourse semantics (Karttunen, 1976; Kamp, 1981; Heim, 1982; Kamp and Reyle, 1993), in that anaphoric expressions such as *she* are also associated with corresponding anaphoric entities. A pleonastic *it* has no associated entities. Quantificational nominals such as *each company* are associated with entity objects because they are 'anaphoric' to group entities accessible in the context. In this setup, the effect of reference resolution is *merging* of multiple entity objects. Here is the algorithm.

1. INPUT: Template entities with the following textual, syntactic, and semantic features:
 - (a) determiner type (e.g., DEF, INDEF, PRON)
 - (b) grammatical or numerical number (e.g., SG, PL, 3)
 - (c) head string (e.g., *automaker*)
 - (d) the head string sort in a sort hierarchy (e.g., *automaker*→*company*→*organization*)
 - (e) modifier strings (e.g., *newly founded*, *with the pilots*)
 - (f) text span (the start and end byte positions)
 - (g) sentence and paragraph positions³
 - (h) text region type (e.g., HEADLINE, TEXT)

²Other referential relationships such as subset and part-whole did not reach sufficiently reliable interannotator agreements. Only identity of reference had a sufficiently high agreement rate (about 85%) between two human annotators.

³Higher text structure properties such as subsections and sections should also be considered if there are any. Exact accessibility computation using complex hierarchical text structures is a future topic of study.

2. FOR EACH potentially anaphoric entity object in the current sentence, in the left-to-right order, DO

(1) COLLECT antecedent entity objects from the *accessible text region*.

- For an entity in a HEADLINE text region, the entire TEXT is accessible because the headline summarizes the text.
- For an entity in a TEXT region, everything *preceding* its text span is accessible (except for the HEADLINE). Intrasentential cataphora is allowed only for first-person pronouns.
- In addition, a *locality* assumption on anaphora sets a (soft) window of search for each referring expression type—the entire preceding text for *proper names*, narrower for *definite noun phrases*, even narrower for *pronouns*, and only the current sentence for *reflexives*. In the MUC-6 system, the window size was arbitrarily set to ten sentences for definites and three sentences for pronouns, ignoring paragraph boundaries, and no antecedents beyond the limit were considered. This clearly left ample room for refinement.⁴

(2) FILTER with *semantic consistency* between the anaphoric entity E_1 and the potential antecedent entity E_2 .

- *Number Consistency*: E_1 's number must be consistent with E_2 's number—for example, twelve is consistent with PLURAL, but not with SINGULAR. As a special case, plural pronouns (*they*, *we*) can take singular organization antecedents.
- *Sort Consistency*: E_1 's sort must either EQUAL or SUBSUME E_2 's sort. This reflects a *monotonicity* assumption on anaphora—for example, since *company* subsumes *automaker*, *the company* can take a *Chicago-based automaker* as an antecedent, but it is too risky to allow *the automaker* to take a *Chicago-based company* as an antecedent. On the other hand, since

automaker and *airline* are neither the same sort nor in a subsumption relation, *an automaker* and *the airline* cannot corefer. (The system's sort hierarchy is still sparse and incomplete.)

- *Modifier Consistency*: E_1 's modifiers must be consistent with E_2 's modifiers—for example, *French* and *British* are inconsistent, but *French* and *multinational* are consistent. (The system doesn't have enough knowledge to do this well.)

(3) ORDER by *dynamic syntactic preference*. The following ordering approximates the relative *salience* of entities. The basic underlying hypothesis is that intrasentential candidates are more salient than intersentential candidates as proposed, for example, in Hobbs (1978) and Kameyama (in press), and that fine-grained syntax-based salience fades with time. Since fine-grained syntax with grammatical functions is unavailable, the syntactic prominence of subjects and left-dislocation is approximated by the left-right linear ordering.

- i. the preceding part of the *same sentence* in the **left-right** order
- ii. the *immediately preceding sentence* in the **left-right** order
- iii. other preceding sentences within the 'limit' (see above) in the **right-left** order

3. OUTPUT: After each anaphoric entity has found an ordered set of potential antecedent entities, there are destructive (indefeasible) and nondestructive (defeasible) options.

- (a) *Destructive Option*: MERGE the anaphoric entity into the preferred antecedent entity.
- (b) *Nondestructive Option*: RECORD the antecedent entity list in the anaphoric entity to allow reordering (i.e., *preference revisions*) by event merging or overall model selection.

The MUC-6 system took the destructive option. The nondestructive option has been implemented in a more recent system.

These basic steps of "COLLECT, FILTER, and ORDER by salience" are analogous to Lappin and

⁴For example, in a more recent FASTUS system, paragraphs are also considered in setting the limit, and at most one candidate beyond the limit is proposed when no candidates are found within the limit.

Leass's (1994) pronoun resolution algorithm, but each step in FASTUS relies on considerably poorer syntactic input. The present algorithm thus provides an interesting case of what happens with extremely poor syntactic input, even poorer than in Kennedy and Boguraev's (1996) system. This comparison will be discussed later.

Name Alias Recognition

In addition to the above general algorithm, a special-purpose *alias recognition* algorithm is invoked for coreference resolution of proper names.⁵

1. INPUT: The input English text is in mixed cases. An earlier transduction phase has recognized unknown names as well as specific-type names for **persons**, **locations**, or **organizations** using name-internal pattern matching and known name lists.
2. FOR EACH new sentence, FOR EACH unknown name, IF it is an alias or acronym of another name already recognized in the given text, MERGE the two—an *alias* is a selective substring of the full name (e.g., *Colonial* for *Colonial Beef*), and *acronym* is a selective sequence of initial characters in the full name (e.g., *GM* for *General Motors*).

Overall Performance

The MUC-6 FASTUS reference resolution algorithm handled only coreference (i.e., identity of reference) of proper names, definites, and pronouns. These are the 'core' anaphoric expression types whose dependencies tend to be constrained by surface textual factors such as locality. The MUC-6 Coreference Task evaluation included coreference of bare nominals, possessed nominals, and indefinites as well, which the system did not handle because we didn't have a reliable algorithm for these mostly 'accidental' coreferences that seemed to require deeper inferences. Nevertheless, the system scored a recall of 59% and precision of 72% in the blind evaluation of thirty newspaper articles.

⁵In addition, a specific-type name may be converted into another type in certain linguistic contexts. For instance, in a *subsidiary* of *Mrs. Field*, *Mrs. Field* is converted from a person name into a company name.

Expression Type	Number of Occurrences	Correctly Resolved
Definites	61	28(46%)
Pronouns	39	24(62%)
Proper Names	32	22(69%)
Reflexives	1	1(100%)
TOTAL	133	75(56%)

Table 1: Core Discourse Anaphors in Five Articles

Grammatical Person	Intra/Inter-S Antecedent	Number of Occurrences	Correctly Resolved
3rd person	intra-S	27	21(78%)
3rd person	inter-S	6	2(33%)
<i>that</i>	inter-S	1	0(0%)
1st/2nd person	inter-S	5	1(20%)
reflexive	intra-S	1	1(100%)

Table 2: Pronouns in Five Articles

Table 1 shows the system's performance in resolving the core discourse anaphors in five randomly selected articles from the development set. Only five articles were examined here because the process was highly time-consuming. The performance for each expression type varies widely from article to article because of unexpected features of the articles. For instance, one of the five articles is a letter to the editor with a text structure drastically different from news reports. On average, we see that the resolution accuracy (i.e., recall) was the highest for proper names (69%), followed by pronouns (62%) and definites (46%). There were not enough instances of reflexives to compare.

Table 2 shows the system's performance for pronouns broken down by two parameters, grammatical person and inter- vs. intrasentential antecedent. The system did quite well (78%) with third-person pronouns with intrasentential antecedents, the largest class of such pronouns.

Part of the pronoun resolution performance here enables a preliminary comparison with the results reported in (1) Lappin and Leass (1994) and (2) Kennedy and Boguraev (1996). For the third-person pronouns and reflexives, the performance was (1) 86% of 560 cases in five computer manuals and (2) 75% of 306 cases in twenty-seven Web page texts. The present FASTUS system correctly resolved 71% of 34 cases in five newspaper arti-

cles. This progressive decline in performance corresponds to the progressive decline in the amount of syntactic information in the input to reference resolution. To summarize the latter decline, Lapin and Leass (1994) had the following components in their algorithm.

1. INPUT: fully parsed sentences with grammatical roles and head-argument and head-adjunct relations
2. Intrasentential syntactic filter based on syntactic noncoreference
3. Morphological filter based on person, number, and gender features
4. Pleonastic pronoun recognition
5. Intrasentential binding for reflexives and reciprocals
6. Saliency computation based on grammatical role, grammatical parallelism, frequency of mention, proximity, and sentence recency
7. Global saliency computation for noun phrases (NPs) in equivalence classes (with seven saliency factors)
8. Decision procedure for choosing among equally preferred candidate antecedents

Kennedy and Boguraev (1996) approximated the above components with a poorer syntactic input, which is an output of a part-of-speech tagger with grammatical function information, plus NPs recognized by finite-state patterns and NPs' adjunct and subordination contexts recognized by heuristics. With this input, grammatical functions and precedence relations were used to approximate 2 and 5. Finite-state patterns approximated 4. Three additional saliency factors were used in 7, and a preference for intraclausal antecedents was added in 6; 3 and 8 were the same.

The present algorithm works with an even poorer syntactic input, as summarized here.

1. INPUT: a set of finite-state approximations of sentence parts, which can be overlapping or discontinuous, with no grammatical function, subordination, or adjunct information.
2. No disjoint reference filter is used.
3. Morphological filter is used.
4. Pleonastic pronouns are recognized with finite-state patterns.
5. Reflexives simply limit the search to the current sentence, with no attempt at recognizing coarguments. No reciprocals are treated.
6. Saliency is approximated by computation based on linear order and recency. No grammatical parallelism is recognized.
7. Equivalence classes correspond to *merged entity objects* whose 'current' positions are always the most recent mentions.
8. Candidates are deterministically ordered, so no decision procedure is needed.

Given how *little* syntactic information is used in FASTUS reference resolution, the 71% accuracy in pronoun resolution is perhaps unexpectedly high. This perhaps shows that linear ordering and recency are major indicators of saliency, especially because grammatical functions correspond to constituent ordering in English. The lack of disjoint reference filter is not the most frequent source of errors, and a coarse-grained treatment of reflexives does not hurt very much, mainly because of the infrequency of reflexives.

An Example Analysis

In the IE context, the task of entity reference resolution is to recognize referential links among partially described entities within and across documents, which goes beyond third-person pronouns and identity of reference. The expression types to be handled include bare nominals, possessed nominals, and indefinites, whose referential links tend to be more 'accidental' than textually signaled, and the referential links to be recognized include subset, membership, and part-whole.

Consider one of the five articles evaluated, the one with the most number and variety of referential links, for which FASTUS's performance was the poorest. Even for the 'core' anaphoric expressions of 20 definites, 14 pronouns, and 7 names limited

Expression Type	Number of Occurrences	Correctly Resolved
Definites	25	10(40%)
Pronouns	14	10(71%)
Bare Nominals	12	3(25%)
Proper Names	7	2(29%)
Possessed Nominals	6	0(0%)
Indefinites	3	0(0%)
TOTAL	67	25(37%)

Table 3: Referential Links in the Example

to coreference, for which the code was prepared, the recall was only 51%.

Figure 1 shows this article annotated with referential indices. The same index indicates coreference. Index subscripts (such as 4a) indicate *subset*, *part*, or *membership* of another expression (e.g., indexed 4). The index number ordering, 1,...,N, has no significance. Each sentence in the TEXT region is numbered with paragraph and sentence numbers, so, for instance, 2-1 is the first sentence in the second paragraph.

Note that not all of these referential links need be recognized for each particular IE application. However, since reference resolution must consider all of the text for any particular application for the reason mentioned above, it is reasonable to assume that an ideal domain-independent reference resolution component should be able to recognize all of these. Is this a realistic goal, especially in an IE context? This question is left open for now.

Error Analysis

Table 3 shows the system's performance in recognizing the referential links in this article, grouped by referring expression types. These exclude the initial mention of each referential chain. Notable sources of errors and necessary extensions are summarized for each expression type here.

Pronouns: Of the four pronoun resolution errors, one is due to a parse error (*American* in 7-1 was incorrectly parsed as a person entity, to which *she* in 8-1 was resolved), *that* in 6-2 is a *discourse deixis* (Webber, 1988), beyond the scope of the current approach, and two errors (*it* in 3-1 and *its* in 7-1) were due to the left-right ordering of intrasentential candidates. Recognition of par-

allelism among clause conjuncts and a stricter locality preference for possessive pronouns may help here.

Definites: Of the fifteen incorrect resolutions of definites, five have *nonidentity* referential relationships, and hence were not handled. These nonidentity cases must be handled to avoid erroneous identity-only resolutions. Two errors were due to the failure to distinguish between generic and specific events. Token-referring definites (*the union* in 8-2 and *the company* in 9-1) were incorrectly resolved to recently mentioned types. Three errors were due to the failure to recognize *synonyms* between, for example, *call* (3-2) vs. *request* (3-1) and *campaign* (9-1) vs. *strategy* (9-1). Other error sources are a failure in recognizing an appositive pattern (9-1), the left-right ordering of the candidates in the previous sentence (9-2), and three 'bugs'.

Proper Names: Name alias recognition was unusually poor (2 out of 7) because *American* was parsed as a person-denoting noun. The lower-case print in *Patt gibbs* also made it difficult to link it with *Ms. Gibbs*. Such *parse errors* and *input anomalies* hurt performance.

Bare Nominals: Since bare nominals were not explicitly resolved, the only correct resolutions (3 out of 12) were due to recognition of appositive patterns. How can the other cases be treated? We need to understand the discourse semantics of bare nominals better before developing an effective algorithm.

Possessed Nominals: A systematic treatment of possessed nominals is a necessary extension. The basic algorithm will look like this—resolve the possessor entity *A* of the possessed nominal *A's B*, then if there is already an entity of type *B* associated with *A* mentioned *anywhere in the preceding text*, then this is the referent of *B*. Possessed nominal resolution also requires a 'synonymy' computation to resolve, for example, *its battle* (9-1) to *its corporate campaign* (7-1). It also needs 'inferences' that rely on multiple successful resolutions. For instance, *her members* in 8-1 must first resolve *her* to *Ms. Gibbs*, who

HEADLINE: **American Airlines₁ Calls for Mediation₁₅ In Its₁ Union₄ Talks₂**

DOCUMENT DATE: 02/09/87

SOURCE: WALL STREET JOURNAL

- 1-1 Amr corp.'s **American Airlines₁** unit said **it₁** has called for **federal mediation₁₅** in **its₁** **contract talks₂** with **unions₄** representing **its₁** **pilots₁₆** and **flight attendants₁₇**.
- 2-1 A spokesman for **the company₁** said **American₁** officials "felt **talks₂** had reached a point where **mediation₁₅** would be helpful."
- 2-2 **Negotiations_{2a}** with **the pilots₁₆** have been going on for 11 months; **talks_{2b}** with **flight attendants₁₇** began six months ago.
- 3-1 **The president₅** of the **Association of Professional Flight Attendants_{4a}**, which represents **American₁**'s **more than 10,000 flight attendants₁₇**, called **the request₃** for **mediation₁₅** "premature" and characterized **it₃** as a bargaining tactic that could lead to a lockout.
- 3-2 **Patt gibbs₅**, **president₅** of **the association_{4a}**, said **talks_{2b}** with **the company₁** seemed to be progressing well and **the call₃** for **mediation₁₅** came as a surprise.
- 4-1 The major outstanding issue in **the negotiations_{2b}** with **the flight attendants₁₇** is a two-tier wage scale, in which recent employees' salaries increase on a different scale than the salaries of employees who have worked at **american₁** for a longer time.
- 4-2 **The union_{4a}** wants to narrow the differences between the new scale and the old one.
- 5-1 **The company₁** declined to comment on **the negotiations_{2b}** or the outstanding issues.
- 5-2 Representatives for **the 5,400-member Allied Pilots Association_{4b}** didn't return phone calls.
- 6-1 Under the Federal Railway Labor Act, [if **the mediator_{15a}** fails to bring **the two sides₇** together and **the two sides₇** don't agree to binding arbitration, a 30-day cooling-off period follows]**6**.
- 6-2 After **that₆**, **the union_{7a}** can strike or **the company_{7a}** can lock **the union_{7a}** out.
- 7-1 **Ms. Gibbs₅** said that in response to **the company₁**'s move, **her₅** **union_{4a}** will be "escalating" **its_{4a}** "corporate campaign" against **American₁** over the next couple of months.
- 7-2 In a **corporate campaign₁₀**, a **union₉** tries to get a **company₈**'s financiers, investors, directors and other financial partners to pressure **the company₈** to meet **union₉** demands.
- 8-1 A **corporate campaign₁₀**, **she₅** said, appeals to **her₅** **members₁₇** because "**it₁₀** is a nice, clean way to take a job action, and **our_{4a}** **women₁₇** are hired to be nice."
- 8-2 **the union_{4a}** has decided not to strike, **she₅** said.
- 9-1 **The union_{4a}** has hired a number of professional consultants₁₄ in **its_{4a}** **battle₁₈** with **the company₁**, including **Ray Rogers_{14a}** of Corporate Campaign Inc., **the New York labor consultant_{14a}** who developed **the strategy₁₂** at **Geo. A. Hormel & Co.₁₃**'s Austin, Minn., meatpacking plant last year.
- 9-2 **That campaign₁₂**, which included a strike, faltered when **the company₁₃** hired new workers and the International Meatpacking Union wrested control of the local union from **Rogers_{14a}**' supporters.

Figure 1: Example Article Annotated with Referential Links

is *president of the association*, and this 'association' is *the Association of Professional Flight Attendants*. After it is understood as 'the members of the Association of Professional Flight Attendants,' the coreference with *the flight attendants* can be inferred. Similarly for *our women* in 8-1.

Indefinites: Some indefinites are 'coreferential' to generic types, for example, *a corporate campaign* (7-2, 8-1). Given the difficulty in distinguishing between generic and specific event descriptions, it is unclear whether it will ever be treated in a systematic way.

Conclusions

In an operational end-to-end discourse understanding system, a reference resolution component must work with input data containing parse errors, lexicon gaps, and mistakes made by earlier reference resolution. In a state-of-the-art IE system such as SRI's FASTUS, reference resolution must work with considerably impoverished syntactic analysis of the input sentences. The present reference resolution approach within an IE system is robust and efficient, and performs pronoun resolution to an almost comparable level with a high-accuracy algorithm in the literature. Desirable extensions include nonidentity referential relationships, treatments of bare nominals, possessed nominals, and indefinites, type-token distinctions, and recognition of synonyms. Another future direction is to turn this component into a corpus-based statistical approach using the relevant factors identified in the rule-based approach. The need for a large tagged corpus may be difficult to satisfy, however.

Acknowledgment

This work was in part supported by U.S. government contracts [REDACTED].

References

Appelt, Douglas, Jerry Hobbs, John Bear, David Israel, Megumi Kameyama, Andy Kehler, David Martin, Karen Myers, and Mabry Tyson. 1995. SRI International FASTUS system: MUC-6 test results and analysis. In *Proceedings of the 6th Message Understanding Conference*, pages 237-248. DARPA.

Grosz, Barbara, Aravind Joshi, and Scott Weinstein. 1995. Centering: A framework for modelling the local coherence of discourse. *Computational Linguistics*, 21(2):203-226.

Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts at Amherst.

Hobbs, Jerry. 1978. Resolving pronoun references. *Lingua*, 44:311-338. Also in B. Grosz, K. Sparck-Jones, and B. Webber, eds., *Readings in Natural Language Processing*, Morgan Kaufmann, Los Altos, CA, 1986, 339-352.

Hobbs, Jerry R., Douglas E. Appelt, John Bear, David Israel, Megumi Kameyama, Mark Stickel, and Mabry Tyson. 1996. FASTUS: A cascaded finite-state transducer for extracting information from natural-language text. In E. Roche and Y. Schabes, editors, *Finite State Devices for Natural Language Processing*. MIT Press, Cambridge, Massachusetts.

Kameyama, Megumi. in press. Intrasentential centering: A case study. In Marilyn Walker, Aravind Joshi, and Ellen Prince, editors, *Centering in Discourse*. Oxford University Press.

Kamp, Hans. 1981. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language*. Mathematical Center, Amsterdam, pages 277-322.

Kamp, Hans and Uwe Reyle. 1993. *From Discourse to Logic*. Kluwer, Dordrecht.

Karttunen, Lauri. 1976. Discourse referents. In James D. McCawley, editor, *Syntax and Semantics: Notes from the Linguistic Underground*, volume 7. Academic Press, New York, pages 363-386.

Kennedy, Christopher and Branimir Boguraev. 1996. Anaphora for everyone: Pronominal anaphora resolution without a parser. In *Proceedings of the 16th International Conference on Computational Linguistics (COLING-'96)*. Association for Computational Linguistics.

Lappin, Shalom and Herbert Leass. 1994. An algorithm for pronominal anaphora resolution. *Computational Linguistics*, 20(4):535-562.

Sidner, Candace. 1979. Towards a computational theory of definite anaphora comprehension in English discourse. Technical Report 537, MIT Artificial Intelligence Laboratory, Cambridge, MA, June.

Sundheim, Beth. 1995. Overview of results of the MUC-6 evaluation. In *Proceedings of the 6th Message Understanding Conference*, pages 13-32. DARPA.

Webber, Bonnie. 1988. Discourse deixis: Reference to discourse segments. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pages 113-122. Association for Computational Linguistics, June.

Webber, Bonnie Lynn. 1978. *A Formal Approach to Discourse Anaphora*. Ph.D. thesis, Harvard University.