

Using Decision Trees for Coreference Resolution*

Joseph F. McCarthy and Wendy G. Lehnert
Department of Computer Science
University of Massachusetts
Amherst, MA 01003-4610
Email {jmccarthy,lehnert}@cs.umass.edu

Abstract

This paper describes RESOLVE, a system that uses decision trees to learn how to classify coreferent phrases in the domain of business joint ventures. An experiment is presented in which the performance of RESOLVE is compared to the performance of a manually engineered set of rules for the same task. The results show that decision trees achieve higher performance than the rules in two of three evaluation metrics developed for the coreference task. In addition to achieving better performance than the rules, RESOLVE provides a framework that facilitates the exploration of the types of knowledge that are useful for solving the coreference problem.

1 Introduction

The goal of an *Information Extraction* (IE) system is to identify information of interest from a collection of texts. Within a particular text, objects of interest are often referenced in different places and in different ways. One of the many challenges facing an IE system is to determine which references refer to which objects. This problem can be recast as a classification problem: given two references, do they refer to the same object or different objects?

The Message Understanding Conferences (MUCs) [Sundheim, 1991, 1992, 1993] and the Tipster Project [Merchant, 1993] helped both to define the information extraction task and to push the technology of IE systems. Each of these evaluation efforts provided a corpus of news articles about a domain, a specification of the relevant information that was to be extracted from each article, the output representation of that information, and a set of key templates representing the information extracted from each article by human readers.¹ For the final evaluations, participating systems were given a set

¹This research was supported by NSF Grant no. EEC-9209623, State/Industry/University Cooperative Research on Intelligent Information Retrieval, Digital Equipment Corporation and the National Center for Automated Information Research.

²The MUC-5 evaluation actually included 4 different domains, but most participants were required to select only one

of blind texts and their output was scored against the key templates to determine how much of the relevant information they were able to extract.

The sentence analyzers used in many of these systems have shown significant improvement over the past several years. However, the discourse processing capabilities of these systems, particularly their coreference resolution components, have often been cited as weak areas [Weir and Fritzon, 1993, Moldovan *et al.*, 1992, Aberdeen *et al.*, 1992].

The IE systems developed at UMass [Lehnert *et al.*, 1991, 1992, 1993] also displayed weak coreference resolution capabilities. Each of these systems used a set of manually engineered rules to resolve some obvious types of coreference, but they tended to be very conservative, i.e., they only considered phrases to be coreferent if there was overwhelming evidence in support of that hypothesis. One of the problems with these coreference resolution components was figuring out which features of the phrases to look at when determining coreference. Another, related set of problems was determining how to combine positive and negative evidence into individual rules and then how to order the rule set. A third problem area was the accumulation of errors at that late stage of processing, e.g., from incorrectly delimited sentences, incorrect part-of-speech tags, and other sentence analysis errors.

In an effort to address these problems, a new approach to coreference resolution was begun after the MUC-5 evaluation: a system named RESOLVE was created to build decision trees that can be used to classify pairs of phrases as coreferent or not coreferent. The errors generated by the sentence analyzer were eliminated by using a special tool - the Coreference Marking Interface, or CMI - to extract a set of phrases from the MUC-5 English Joint Venture (EJV) corpus.³ In order to minimize the difficulties involved with creating and maintaining complex sets of rules, a machine learning approach was adopted, in which a decision tree determines the order and relative weight of different pieces of evidence.

The MUC-5 EJV corpus is a collection of news articles, written in English, that describe business joint ventures, i.e., associations of two or more entities (companies, governments or people) created for the purpose of owning and/or developing a project together.

RESOLVE used the C4.5 decision tree system [Quinlan, 1993] to learn how to classify coreferent phrases for the experiments reported in this paper. C4.5 was chosen primarily due to its ease of use and its widespread acceptance, however, RESOLVE can use any learning system that uses feature vectors composed of attribute-value pairs.

2 Decision Trees vs Rules

An experiment was conducted to compare the performance of the decision trees generated by RESOLVE with the performance of manually engineered rules used for coreference classification in the UMass/Hughes MUC-5 IE system. A set of references, along with the coreference links among these, were extracted from a group of texts via CMI. All possible pairings of references from each text were generated, and these pairings were used to create a set of feature vectors used by RESOLVE. The pairings that contained coreferent phrases formed positive instances, while those that contained two non-coreferent phrases formed negative instances. RESOLVE was then iteratively trained and tested on different partitions of this set of feature vectors.

The data structure used in discourse processing by the UMass/Hughes MUC-5 IE system was the memory *token*, which converted the case frame output from the CIRCUS sentence analyzer [Lehnert, 1991] into a more system-independent representation. Prior to coreference processing, each memory token contained one noun phrase, one or more lexical patterns encompassing that phrase, part-of-speech tags, semantic features, and information that was inferred from either the phrase or the context in which the phrase was found. This inferred information included the type of object referenced by the phrase, any name or location substring contained in the phrase, and some domain-specific information such as whether the phrase was a joint venture parent (one of the entities who formed a joint venture) or joint venture child (the joint venture company itself). The references marked via CMI were converted into a memory token representation in order to test the performance of the MUC-5 system's coreference module.

2.1 Data

The articles in the EJV corpus describe business joint ventures among two or more entities (companies, governments and/or people). The task definition provided for MUC-5 required IE systems to extract information about the entities involved, the relationships among these entities, the facilities associated with the joint venture, the products or services offered by the joint venture, its capitalization and revenue projections, and a variety of other related information. Since the entities involved in these joint ventures were the main focus of most of these articles, references to entities were much more numerous than references to other types of object classes, e.g., people. Therefore, entity references were selected as the focus of the experiments reported in this paper.

CMI is a graphical user interface that permits the user to mark phrases in a text, for each phrase, the user can indicate the object(s) with which the phrase is coreferent

and some additional information about the phrase that can be inferred either from the phrase itself or its local context. This additional information is parameterized and can be modified easily for use in different domains. The data used in this experiment was based on a set of phrases extracted using CMI.

As an example, consider the following sentence, from text 0970 from the MUC-5 EJV corpus

FAMILYMART CO OF
SEIBU SAISON GROUP WILL OPEN
A CONVENIENCE STORE IN TAIPEI FRIDAY
IN A JOINT VENTURE WITH
TAIWAN'S LARGEST CAR DEALER.
THE COMPANY SAID WEDNESDAY

The phrases underlined in this sentence contain relevant information that must be extracted by an IE system. The phrases in boldface refer to entity objects that are important to the MUC-5 task. As an example of the types of information collected about each phrase, consider the first phrase in the sentence

```
( string "FAMILYMART CO "
  slots (ENTITY
         name "FAMILYMART CO ")
         type COMPANY)
        (relationship JV-PARENT CHILD)))
```

Information collected about each phrase includes the string itself, the character position of the string in the source text (not shown), the index of the sentence within which the string is found (also not shown), and some slot information that can be inferred from either the string itself or its local context - the same kind of information that was contained in the memory tokens used by the MUC-5 system. In this example, the name of the entity and the fact that it is a company entity can both be inferred from the string itself. The fact that Familymart Company plans to open a store in "A JOINT VENTURE" with another entity is considered adequate evidence that the company is the parent of a joint venture (*JV-PARENT*), the fact that the sentence contains the pattern "*company name 1 OF company-name 2*" is evidence that *company-name-1*, in this case Familymart Co., is a subsidiary (*child*) of *company-name-2*, in this case Seibu Saison Group.

A second example of output from CMI can be seen below, where nationality information has been extracted from the reference to the car dealer.

```
( string "TAIWAN'S LARGEST CAR DEALER"
  slots (ENTITY
         .type COMPANY)
         (relationship JV PARENT)
         (nationality "Taiwan (COUNTRY)")))
```

³Note that the phrase "THE COMPANY" in the last clause of the sentence is not considered relevant, since it contributes no information required for the MUC-5 task - the determination of who is announcing a joint venture or when the announcement was made are not relevant pieces of information. Therefore, this phrase was not marked for use in the experiment.

In principle, much of the information gathered about a particular string could be found automatically there are numerous proper name recognizer programs, programs that extract location information, and sentence analysers that can infer relationship information - any system that exhibited good performance in MUC-5 must be good at inferring such relationships

For the purposes of our experiment, however, this information was specified by a user via CMI. The primary motivation for this was to minimise the noise in the data, coreference resolution often occurs at a late processing stage in an IE system, and earlier errors such as incorrect part-of-speech tags, incorrectly delimited sentences and semantic tagging errors can create significant noise for a coreference classifier

CMI was used to mark references to a variety of relevant object types (entity, facility, person and product-or-service) in 50 randomly selected texts*. Since references to entity objects were most numerous, this was the object class chosen for the experiment. In the 50 texts, 472 references to a total of 205 entity objects were marked using CMI

Some phrases are *multireferent*, i.e., they refer to more than one object. These multireferent phrases pose difficulties for classification, since it means that some phrases will be coreferent with other phrases in the text that have distinct referents. Thus for a set of phrase pairs which share a given phrase, more than one pair would be classified as a positive instance of coreference. Further complications are created for evaluating the performance of a coreference system when multireferent phrases are included in the data (see Section 2.4). To simplify the initial experiments reported here, multireferent phrases were excluded from the data set. The capability to handle such phrases will be incorporated in a later version of RESOLVE.

2.2 Rules used in the MUC-5 System

The coreference module of the UMAas/Hughes MUC-5 IE system was designed to minimise false positives, i.e., minimise the likelihood that two phrases that were not coreferent would be labeled coreferent. This design decision was based on the assumption that false positive errors, resulting in the merging of non-coreferent phrases in the final system output, would harm system performance more than false negative errors, which would result in coreferent phrases showing up in distinct objects in the system output. This rather conservative approach to coreference was shared by a number of MUC system developers [Appelt *et al.*, 1992, Ayuso *et al.*, 1992], though not all [Iwanska *et al.*, 1992]

Another factor influencing the coreference module was the short time allotted to developing and testing this system component. Since coreference resolution was a late stage in processing, upstream components had to be stabilised before serious development could take place on coreference. Several late-stage components were being developed in parallel, so it is difficult to assess the time

*In order to make things manageable for fMI annotator, the size of the texts was limited to 2KB, however the majority of texts in the EJV domain fall into this category

IF	both tokens come from the same <i>trigger family</i>
THEN	<u>they are not coreferent</u>
IF	each token comes from a different partition
THEN	<u>they are not coreferent</u>
IF	both tokens contain a common phrase
THEN	<u>they are coreferent</u>
IF	both tokens refer to joint ventures
THEN	<u>they are coreferent</u>
IF both tokens contain the same company name	
THEN	<u>they are coreferent</u>
IF	one token contains an alias of the other
THEN	<u>they are coreferent</u>
IF	only one token refers to a joint venture
THEN	<u>they are not coreferent</u>
7F	each token contains different company names
THEN	<u>they are not coreferent</u>

Table 1 The MUC-5 system's coreference rules

devoted exclusively to developing the coreference module, but we estimate it was two person weeks

The rules used to determine whether two phrases (represented as memory tokens) were coreferent in the MUC-5 system are shown in Table 1. Following the policy of minimising false positives, whenever none of the rules fired, the system classified the pair of tokens as not coreferent

The UMass/Hughes MUC-5 IE system used a variety of mechanisms to identify phrases referring to joint ventures (the entity formed by two or more parent entities for some particular business purpose), to identify company names within a phrase (if they exist), and to determine whether one phrase was an abas (an abbreviation or shortened form), as well as the ability to identify trigger families⁶ and partitions⁶ in the text

One of the many difficulties in developing the rule set for coreference classification was in ordering the rules. Several different orderings were tested during the development period, and the order shown above was the ordering of the rule set used for final evaluation. This difficulty in rule ordering was one of the motivations behind using a machine learning approach - we wanted to develop a system that could *learn* how to combine the positive and negative evidence

2.3 Features Used By RESOLVE

A decision tree requires data to be represented by feature vectors, i.e., vectors of attribute /value pairs. For the task of coreference classification, references were paired up, and features were extracted from the pair of references as well as from the individual references themselves. Since this experiment involved a comparison between RESOLVE and a manually engineered rule set, the

⁶A trigger family is a set of phrases all *triggered* off the same word, e.g., subject and direct object joined by the same verb phrase

A partition is a portion of the text that is focusing on the same main topic. For the MUC-5 system, distinct partitions were recognized only for texts that had bulleted items, as one might see in a news summary of the days headlines. Most texts thus had a single partition

Individual Phrases		Pair of Phrases	
Attribute	Value	Attribute	Value
NAME-1	YES	ALIAS	NO
JV-CHILD 1	NO	BOTH JV-CHILD	NO
NAME-2	YES	COMMON-NP	NO
JV CHILD 2	NO	SAME-SENTENCE	NO

Table 2 Attributes and Values for EJV entity instance

features used in this experiment were based on the antecedents of the coreference rules used in the UMass/Hughes MUC-5 IE system

For example, Table 2 shows a feature vector that represents the pairing of the phrases "FAMILYMART CO" and "TAIWAN'S LARGEST CAR DEALER". Since the two phrases are not coreferent, this represents a negative instance

Of the 8 features used in this experiment, two focus on the first reference, two focus on the second reference and four are based on the pair of references. The following is a brief description of the features that focus on individual phrases, where $i \in \{1, 2\}$

- NAME- i Does reference i contain a name? Possible values {YES, NO}
- JV CHILD i Does reference i refer to a joint venture child, i.e., a company formed as the result of a tie-up among two or more entities? Possible values {YES, NO, UNKNOWN}

The last four features focus on the pair of references

- ALIAS Does one reference contain an alias of the other, i.e., does each reference contain a name and is one name a substring of the other name? Possible values {YES, NO}
- BOTH JV CHILD Do both references refer to a joint venture child? This feature is defined as
 yes when $\forall i, JV-CHILD-i = YES$
 no when $\forall i, JV-CHILD-i = NO$
 unknown otherwise
- COMMON NP Do the references share a common noun phrase? Some references contain non-simple noun phrases, e.g., appositions and relative clauses. This feature compares the simple constituent noun phrases of each reference. Possible values {YES, NO}
- SAME SENTENCE Do the references come from the same sentence? RESOLVE does not use CIRCUS output, and thus has no notion of a trigger family as it was used in the MUC-5 system, the SAME SENTENCE feature is a very weak attempt to extract this sort of information. Possible values {YES, NO}

⁷Note that some texts contain more than one entity for which a given name might be an alias under this definition, e.g., "SUMITOMO" is a substring of both "SUMITOMO CORP" and "SUMITOMO ELECTRICAL INDUSTRIES LTD", so this feature is not always a reliable indicator of coreference

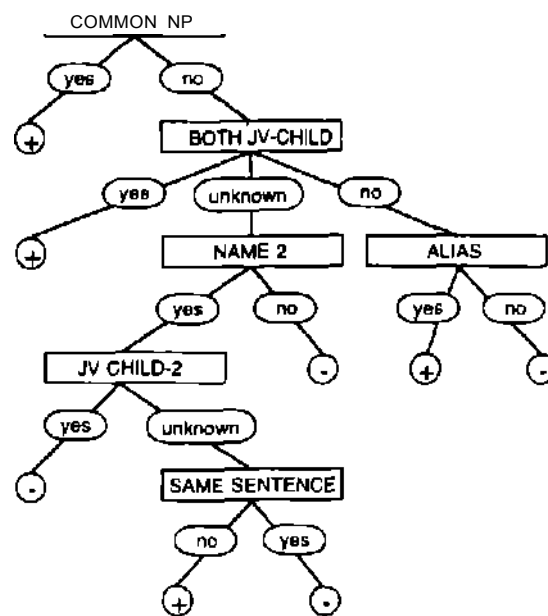


Figure 1 A pruned C4.5 decision tree

1230 feature vectors, or instances, were created from the entity references marked in the 50 texts. Of these, 322 (26%) were $\mu(“+”)$ instances - pairs of phrases that were coreferent - and 908 (74%) were $\mu(“-”)$ instances - pairs of phrases that were not coreferent. Figure 1 shows a pruned C4.5 decision tree trained on all the instances

2.4 Evaluation Methodology

Coreference is a symmetrical and transitive relation that holds among a set of two or more references, e.g., if we know that A is coreferent with B , and B is coreferent with C , then there is an implicit coreference "link" between A and C . Any coreference classification for two references has implications beyond the determination of whether that particular classification was correct or incorrect. For example, if A and B are correctly classified as coreferent, but B and C are incorrectly classified as not coreferent, a system may also incorrectly conclude that A and C are not coreferent. Thus, simply measuring the accuracy of a coreference classifier is inadequate for evaluating how well the classifier performs its task.

Two metrics that have been used to evaluate the performance of IE systems are *recall* and *precision* [Chinchor, 1991, 1992, Chinchor and Sundheim, 1993]. Recall is the percentage of information in a text that is correctly extracted by a system, precision is the percentage of information extracted by a system that is correct. For example, if a text contains four relevant items (represented

As was noted earlier, some references are *multireferent*, i.e., they have more than one referent. Thus, if B is multireferent, we cannot conclude that A is coreferent with C , for example, if $A = Sneezy$, $B = the dwarfs$ and $C = Grumpy$, we don't want to infer that $Sneezy = Grumpy$. We can ignore such complications in this paper since the experiments reported herein exclude multireferent phrases.

by {A, B, C, D} in an answer key), and a system correctly extracts the three items {A, B, C} but incorrectly extracts the two additional items {E, F} (represented by {A, B, C, E, F} in a system response), then its recall would be 75% and its precision would be 60%.

A function to combine recall and precision into a single measure of performance was incorporated into the Fourth Message Understanding Evaluation and Conference [Chinchor, 1992]. The *F-measure*, a metric used to evaluate Information Retrieval (IR) system performance [van Rijsbergen, 1979], combines recall and precision scores into a single number using the formula

$$F = \frac{(\beta^2 + 1.0) \times P \times R}{\beta^2 \times P + R}$$

where P is the precision score, R is the recall score and β is the relative weight given to recall over precision. For example, a β value of 1.0 gives equal weight to recall and precision, a value of 2.0 gives recall twice the weight of precision, a value of 0.5 gives recall half the weight of precision.

An evaluation methodology for the coreference task is being developed for the upcoming Sixth Message Understanding Evaluation and Conference (MUC-6). The metrics used for evaluating overall IE system performance are being adapted for use on this sub task (cf [Burger *et al.*, 1994]), where the answer key for each text contains a set of phrases and the coreference links among them. However, evaluation of coreference performance is complicated by the need to take into account the *implicit* coreference links among phrases. Thus, transitive closures are taken for both the answer key (the *key closure*) and the system response (the *response closure*). Recall is measured by the percentage of explicit coreference links in the key that are also found in the response closure, i.e., what fraction of correct coreference links is implied by the transitive closure of the coreference links in the system response. Precision is measured by the percentage of explicit coreference links in the response that are also found in the key closure, i.e., what fraction of coreference links in the response is implied by the transitive closure of the coreference links in the key.

2.5 Results

One experiment was run using RESOLVE. In this experiment, for each set of instances taken from the 50 texts, one set was selected for testing purposes and the remaining sets were used to train a new decision tree. This process was iterated over all 50 sets of instances. The results shown in Table 3 represent the average of these iterations: the first row shows the recall, precision and F-measure ($\beta = 1.0$) scores for unpruned decision trees, the second row shows the results for pruned decision trees⁹.

The third row in Table 3 shows the results from a second experiment, in which the rule set from the coreference module of the UMass/Hughes MUC-5 IE system

⁹Default Bettings for all C4.6 parameters were used throughout this experiment (see [Quinlan, 1993], Chapter 9, for more information about C4.5 parameters)

System	Recall	Precision	F-measure
RESOLVE (unpruned)	85.4%	87.6%	86.5%
RESOLVE (pruned)	80.1%	92.4%	85.8%
MUC-5 rule set	87.7%	94.4%	78.9%

Table 3 Results for EJV entity coreference resolution

System	$\beta = 2.0$	$\beta = 1.0$	$\beta = 0.5$
RESOLVE (unpruned)	85.8%	86.5%	87.1%
RESOLVE (pruned)	82.3%	85.6%	89.6%
MUC 5 rule set	71.8%	78.9%	87.5%

Table 4 F measures for different values of β

was applied to the memory token pairs generated from the references marked using CMI.

2.6 Discussion

When we first began applying decision trees to the coreference resolution problem, we were hoping to achieve performance that was comparable to the manually engineered rules we had used in MUC-5. We were greatly encouraged to discover that we could achieve performance that surpassed the performance of the rules from our MUC-5 system in both recall and F-measure scores.

As was noted earlier, the MUC 5 coreference rules were designed to minimize false positives. The effect of this bias can be seen in the higher precision score achieved by the rule set in comparison with both the unpruned and pruned decision trees. The difference in precision scores between the unpruned and pruned versions of the decision trees might be explained by the prevalence of negative instances (74%) in the data set, which may lead to a stronger bias to classify pairs of phrases as not coreferent in the smaller trees.

The comparative effects of false positives and false negatives in coreference classification on overall IE system performance remains an open question. However, while the precision scores achieved by the decision trees and the rule-base are rather close, especially for the pruned version of the trees, there is a large difference between their recall scores. Until we can ascertain the relative importance of high recall vs. high precision in overall IE system performance, the F-measure score that gives equal weight to recall and precision may be the best indicator of overall performance on the coreference resolution task. However, as can be seen in Table 4, when RESOLVE uses pruning, its performance surpasses that of the rule set even when the recall score is given twice the weight of precision score or when the recall score is given half the weight of precision score¹⁰.

3 Conclusions

One of the original goals of this new approach was to develop a system that achieved good performance in resolving references - performance that was at least as good as the performance achieved using manually engineered

¹⁰ The pruned decision trees yield higher F-measure scores than the MUC 5 rule set unless the recall score is given less than one third the weight of the precision score.

rules in our MUC-5 system. However, as we continue to pursue this approach, we find that there is another advantage to using decision trees: they allow us to focus on determining which features work well for resolving references.

We are encouraged by the performance of the decision trees on the coreference resolution problem. The features we have used in the experiment described above are not considered comprehensive by any means. While they have proved sufficient for attaining a certain level of performance, an examination of specific errors made by the trees shows that additional features will be needed to attain higher levels.

One area we will develop further is a set of features that incorporate syntactic knowledge. We don't have any features that identify the various syntactic constituents of a sentence, e.g., subject or direct object, nor do we have any features that identify clause boundaries (only sentence boundaries). These features will be incorporated in future experiments. Features based on focus of attention [Sidner, 1979, Grosz *et al.*, 1983], which presuppose knowledge about syntactic constituents may also prove useful. Our experiment used a feature set that was largely semantic in nature: it is interesting to see how well semantic features work as a basis for coreference resolution, and it is not surprising to see that they are also insufficient.

Ultimately, we hope to understand better which features are important for coreference classification, across different objects and different domains. Such an understanding would benefit people involved with IE system development, and should be of interest to people outside the IE community as well. We think that decision trees are an important tool in a systematic study of coreference resolution.

References

- [Aberdeen *et al.*, 1992] J. Aberdeen, J. Burger, D. Connolly, S. Roberts, and M. Vilain. MITRE-Bedford ALEM BIC MUC 4 test results and analysis. In *Proceedings of the Fourth Message Understanding Conference (MUC 4)*, pages 116-123, 1992.
- [Appelt *et al.*, 1992] D. E. Appelt, J. Bear, J. R. Hobbs, D. Israel, and M. Tyson. SRI International FASTUS system MUC 4 test results and analysis. In *Proceedings of the Fourth Message Understanding Conference (MUC 4)*, pages 143-147, 1992.
- [AVUSO *et al.*, 1992] D. Ayuso, S. Boiscn, H. Fox, H. Gish, R. Ingna, and R. Weischedel. BBN Description of the PLUM svBtem as used in MUC 4. In *Proceedings of the Fourth Message Understanding Conference (MUC 4)*, pages 189-176, 1992.
- [Burger *et al.*, 1994] J. Burger, M. Vilain, J. Aberdeen, D. Connolly, and L. Hirechman. A model theoretic coreference scoring scheme. Technical report, The MITRE Corporation, Bedford, MA, 1994.
- [Chinchor and Sundheim, 1993] N. Chinchor and B. Sundheim. MUC-5 evaluation metrics. In *Proceedings of the Fifth Message Understanding Conference (MUC 5)*, pages 22-29, 1993.
- [Chinchor, 1991] N. Chinchor. MUC-3 evaluation metrics. In *Proceeding of the Third Message Understanding Conference (MUC 3)*, pages 17-24, 1991.
- [Chinchor, 1992] N. Chinchor. MUC 4 evaluation metrics. In *Proceedings of the Fourth Message Understanding Conference (MUC-4)*, pages 22-29, 1992.
- [Grosz *et al.*, 1983] B. J. Grosz, A. K. Joshi, and S. Weinstein. Providing a unified account of definite noun phrases in discourse. In *Proceedings of the list Annual Meeting of the ACL*, pages 44-50, 1983.
- [Iwanska *et al.*, 1992] L. Iwanska, D. Appelt, D. Ayuso, K. Dahlgren, B. Glover Stalls, R. Gnselman, G. Krupka, C. Montgomery, and E. RilofT. Computational aspects of discourse in the context of MUC 3. In *Proceedings of the Third Message Understanding Conference (MUC 3)*, pages 256-282, 1992.
- [Lehnert *et al.*, 1991] W. Lehnert, C. Cardie, D. Fisher, E. RilofT, and R. Williams. University of Massachusetts. Description of the CIRCUS system as used for MUC-3. In *Proceedings of the Third Message Understanding Conference (MUC 3)*, pages 223-233, 1991.
- [Lehnert *et al.*, 1992] W. Lehnert, C. Cardie, D. Fisher, J. McCarthy, E. RilofT, and S. Soderland. University of Massachusetts. Description of the CIRCUS system as used for MUC-4. In *Proceedings of the Fourth Message Understanding Conference (MUC 4)*, pages 282-288, 1992.
- [Lehnert *et al.*, 1993] W. Lehnert, J. McCarthy, S. ooderland, E. RilofT, C. Cardie, J. Peterson, F. Feng, C. Dolan, and S. Goldman. University of Massachusetts/Hughes. Description of the CIRCUS system as used for MUC 5. In *Proceedings of the Fifth Message Understanding Conference (MUC 5)*, pages 277-290, 1993.
- [Lehnert, 1991] W. Lehnert. Symbolic/subsymbolic sentence analysis: Exploiting the best of two worlds. In J. Barnaden and J. Pollack, editors, *Advances in Connectionist and Neural Computation Theory, Vol 1*, pageB 135-164. Ablex Publishers, Norwood, NJ, 1991.
- [Merchant, 1993] R. H. Merchant. Tipster program overview. In *Proceedings of the TIPSTER Text Program (Phase I)*, pages 1-2, 1993.
- [Moldovan *et al.*, 1992] D. Moldovan, S. Cha, M. Chung, K. Hendnckson, J. Kim, and S. Kowalski. USC MUC 4 test results and analysis. In *Proceedings of the Fourth Message Understanding Conference (MUC 4)*t pages 164-166, 1992.
- [Quinlan, 1993] J. R. Quinlan. *C4 5 Programs for Machine Learning*. Morgan Kaufmann, San Mateo, CA, 1993.
- [Sidner, 1979] C. L. Sidner. Towards a computational theory of definite anaphora comprehension in English discourse. TR 537, MIT Artificial Intelligence Laboratory, 1979.
- [Sundheim, 1991] B. M. Sundheim. Overview of the third message understanding evaluation and conference. In *Proceedings of the Third Message Understanding Conference (MUC 3)*, pages 3-16, 1991.
- [Sundheim, 1992] B. M. Sundheim. Overview of the fourth message understanding evaluation and conference. In *Proceedings of the Fourth Message Understanding Conference (MUC 4)*, pages 3-21, 1992.
- [Sundheim, 1993] B. M. Sundheim. TIPSTER/MUC-5 information extraction system evaluation. In *Proceedings of the Fifth Message Understanding Conference (MUC 5)*, pages 27-44, 1993.
- [van Rijsbergen, 1979] C. J. van Rijsbergen. *Information Retrieval*. Butterworths, London, 1979.
- [Weir and Fntzson, 1993] C. War and R. Fntason. UNISYS. Description of the CBAS system used for MUC-5. In *Proceedings of the Fifth Message Understanding Conference (MUC 5)*, pages 249-261, 1993.