Anaphora Resolution: A Multi-Strategy Approach

Jaime G. Carbonell
Ralf D. Brown
Computer Science Department & Center for Machine Translation
Carnegie-Mellon University
Pittsburgh, PA 15213

15 April 1988

Abstract

Anaphora resolution has proven to be a very difficult problem; it requires the integrated application of syntactic, semantic, and pragmatic knowledge. This paper examines the hypothesis that instead of attempting to construct a monolithic method for resolving anaphora, the combination of multiple strategies, each exploiting a different knowledge source, proves more effective - theoretically and computationally. Cognitive plausibility is established in that human judgements of the optimal anaphoric referent accord with those of the strategy-based method, and human inability to determine a unique referent corresponds to the cases where different strategies offer conflicting candidates for the anaphoric referent.

1. Introduction: The Complexity of Anaphora Resolution

Anaphora is a pervasive phenomenon in natural language communication, whether it be complex multi-party human discourse or more constrained bilateral human-computer dialog. Whereas the ubiquity of the phenomenon and the necessity to cope with it in the construction of friendly natural language interfaces has long been recognized [13, 8, 9, 17, 15], no truly comprehensive computational approaches for anaphora resolution have been proposed. The RUS parser [2], the XCALIBUR system [7], and other operational natural language systems implement very rudimentary methods. And, theoretical work in linguistics is primarily concerned with certain types of syntactic intrasentential coreference, rather than pragmatic intersentential anaphora.

Anaphors typically refer back to other constituents in the same sentence, or to constituents in earlier utterances in the discourse. Syntactic information plays a central role in establishing appropriate referents for the former case, *intrasentential anaphora* [17]. But, semantic and pragmatic information is absolutely required in the latter case, *intersentential anaphora* [15, 9]. This paper addresses the problem of intersentential anaphora resolution, integrating caseframe semantics [10, 12, 5] and more global dialog coherence structures [11, 15, 14]. Empirical studies have shown that intersentential anaphora[1] is far more frequent and more crucial in designing interactive natural language interfaces[2] [8].

2. The Problem: Semantics and Pragmatics Dominate

Finding the appropriate anaphoric referent has been long recognized as a difficult problem, requiring much semantic and pragmatic knowledge. Consider, for instance, the following two sets of examples:

> *John took the cake from the table and ate it.*
> *John took the cake from the table and washed it.*
>
> *The robot pushed the box towards the conveyor belt. But, it goofed and dropped it on its way there.*

Semantic preference constraints (e.g., [18, 1]), if properly coded, suffice to resolve the first example. The preferred object of ingestation is an edible substance. It is a little more difficult to mechanize a process that excludes things such as cakes from being the object of washing. One cannot simply write a "NOT(edible)" restriction on the object case of the verb "to wash". After all, vegetables and fruits are occasionally washed prior to eating them. Perhaps a combination of typicality judgements with pragmatic knowledge extrapolating the effects of attempting to drown a cake in a sink full of water comes into play. Or, more abstract inferential constraints are appropriate, such as requiring that the object of *wash* be unchanged by immersion in water. Interestingly, Subjects given only the "...and washed it" sentence report consistently that they *didn't even consider* the cake a reasonable referent for "it".

In the robot example, there are four anaphoric referents, counting the possessive "its" and the locative "there", referring to three different antecedents. Although subjects report little difficulty ascertaining the referent for each anaphor in a consistent manner, it appears that sophisticated semantics are required. Why is the referent for "it" in "it goofed and dropped..." the robot rather than the box or the conveyor belt? One could argue that the box cannot take action, but what allows a robot to goof and not a conveyor belt? Is it something as subtle as the degree to which the former can be anthropomorphized being greater than the degree to which the latter can be anthropomorphized?

The difficulty in anaphoric referent specification in narratives has been argued convincingly by many researchers including Charniak in his work on children's story comprehension [9], where substantial pragmatic domain knowledge must be brought to bear, and by one of the authors [4], where knowledge of goals and personality traits is required to resolve difficult referents. Hence, the hypothesis that anaphor resolution in its full generality is at best a difficult problem, and at worst an almost intractable one, is well supported. Nevertheless, somewhat less ambitious endeavors can prove far more tractable, and yet be of major practical import. Hayes [13] argued for the notion of limited-domain anaphora in a natural language interface to an electronic mail system. Webber [17] demonstrated that intrasentential anaphora was more tractable than its intersentential counterpart, largely through the categorization of syntactic devices absent from larger textual or dialog segments.

This paper explores an intermediate position: addressing much larger classes of anaphors than those of Hayes [13] in a systematic manner, but stopping short of full generality, which requires unbounded pragmatic knowledge and inference. We explore the central hypothesis that anaphora resolution may be best accomplished through the combination of a set of strategies, rather than by a single monolithic method. The apparent complexities lie in the combination of these multiple strategies to produce syntactically, semantically and pragmatically sound anaphoric resolutions. In the multiple examples analyzed,[3] unambiguous resolutions reported by human subjects correspond to situations where the applicable strategies concur on the referent of an anaphor, and disagreement on the correct referent by the human subjects corresponds to situations where the applicable strategies propose different candidate referents for the anaphor in question.

3. Multiple Resolution Strategies

In this section we propose a general framework for anaphor resolution based on the integration of multiple knowledge sources: sentential syntax, case-frame semantics, dialog structure, and general world knowledge. The underlying theoretical tenet is:

Anaphor resolution is not a monolithic autonomous process; it requires access and integration of all the knowledge sources necessary for dialog and text interpretation. These linguistic knowledge sources are brought to bear as constraints or preferences encoded as multiple resolution strategies.

Each source of knowledge useful in resolving intersentential anaphora is presented below, along with corresponding examples, and a statement of the anaphoric resolution strategy.

3.1. Local Anaphor Constraints

Certain anaphors carry with them constraints (number, gender, case, etc.) which must be satisfied by the candidate referents. For instance, gender uniquely specifies the anaphor in:

John and Mary went shopping. He bought a steak. [he=John]

The strategy here is trivial:

Eliminate from consideration all candidate referents that violate the local constraints of the anaphor in question.

A variant of this strategy has been implemented in RUS and in XCALIBUR.

3.2. Case-role Semantic Constraints

Here the case-role semantics impose constraints on what can fill them. If they are filled by an anaphor (which specifies few if any semantic features), the case role constraints must be also satisfied by the referent of the anaphor, thus eliminating from consideration all candidate anaphor referents that violate constraints on the case role occupied by the anaphor. Consider our previous example, where the semantic constraints on the object case of "to eat" and "to wash" impose restrictions on the possible case fillers and prove sufficient to select a unique referent.

John took the cake from the table and ate it. [it=cake]
John took the cake from the table and washed it. [it=table]

The strategy here is also fairly simple:

Eliminate from consideration all candidate referents that violate any case-constraint imposed on the anaphor in question. Prefer those candidates that accord with typical case fillers, in the absence of hard constraints.

XCALIBUR implements this strategy directly though use of its case-frame grammar. With the I-rule mechanism, it was possible to implement an *ad-hoc* variant of this strategy in RUS as well.

3.3. Precondition/Postcondition Constraints

Using real-world knowledge and pragmatics, it is possible to say that a candidate antecedent cannot be the referent of an anaphor because some action occurring between the referent and the anaphor invalidates the assumption that they denote one and the same object or event.

John gave Tom an apple. He ate the apple. [he=Tom]

Here, "he" refers to Tom, as John no longer has the apple. The postcondition on *give* is that the actor no longer have the object being given, which conflicts with the precondition on *eat* that the actor have the item being eaten, if the actor is assumed to be John.

The strategy is simple, but requires a fairly large amount of knowledge to be useful for a broad range of cases:

Eliminate from consideration all candidate referents associated with actions whose postconditions violate the preconditions of the action containing the anaphor.

3.4. Case-role Persistence Preference

We observe a pervasive form of "linguistic inertia" that manifests as a preference to assign the referent of an anaphor to the linguistic entity in the discourse context that filled the corresponding semantic case role in an earlier utterance. This is a generalized form of case-role parallelism, which has proven crucial in ellipsis resolution [8, 7, 5], although in anaphora resolution it is demoted from the status of a categorical constraint to that of a preference.

Mary gave an apple to Susan. John also gave her an orange. [her=Susan]
Mary gave an apple to Susan. She also gave John an orange. [she=Mary]

The first anaphor refers to Susan, whereas the second anaphor refers to Mary. Clearly it is not a matter of primacy or recency, as the sentence structures are identical. Rather it is a case of structural parallelism. And, the semantic structure dominates over the syntactic one. For instance, in the first example, "Susan" is the object of the "to" prepositional phrase, whereas the coreferent anaphor is in the indirect object position: two different syntactic roles that map into the same semantic case, *recipient*. In the second example above, both syntactic and semantic structures coincide, and therefore the preference is stronger. Note, moreover, that the subject or direct-object form of the pronoun ("she" vs "her") is not the primary source of discriminant knowledge. For instance, in the example below, one has only the anaphor "it", but the same referent discrimination occurs by semantic case-role parallelism:

The robot gave the dog a bone. John also gave it some water. [it=dog]
The robot gave the dog a bone. It also gave John some water. [it=robot]

To provide more ammunition in support of semantic case role persistence, consider the following final example, with three possible referents to the anaphor "him". It is clear that "Peter" is the preferred referent, once again due to the persistence of the underlying semantic recipient case.

John carried the box of papers from Bill to Peter. He also sent him Mary's books. [he=John, him=Peter]

The semantic preference strategy can be stated as follows:

Search first for acceptable referents in the antecedent phrase (or phrases) that occur in the same semantic case role as the anaphor. If a match satisfying all constraints is found, look no further; else search the other case roles.

To our knowledge, this preference strategy has been neither proposed nor implemented prior to our work on the Universal Parser (reported below), yet it counts for a large number of anaphor resolutions in our sample set.

3.5. Semantic Alignment Preference

A form of pragmatic "Occam's razor" exists in not postulating extra roles for the same objects in different sentences in the discourse. This preference is a more general and looser form of *case role inertia*, discussed above, in that the we have inertia of the underlying action. For instance, in the example below, this preference manifests as preferring all departures to be from the park, and all arrivals to be at the club:

Mary drove from the park to the club. Peter went there too. [there=club]
Mary drove from the park to the club. Peter left there too. [there=park]

The locative anaphor "there" refers to "the club" in the first example above, but refers to "the park" in the second example, yet both sentences share the identical syntactic structure and the same basic underlying semantic case structure. However, discourse cohesion

prefers to make the sentences coreferential (pragmatically parallel) with respect to the same underlying action (leaving the park and going to the club). Therefore, the former aligns with the second (destination) part, whereas the latter aligns with the first (source) part. The strategy here is a bit more difficult to state, and certainly has not been implemented in any system to date:

> *If the clause in which the anaphor is embedded aligns with a previous clause ("aligns" means that it can represent the same underlying action, perhaps with different instantiated case fillers), or with part of a previous clause, search first for referents of the anaphor in that clause. If there are no allowable referents in the semantically aligned clause, expand the search to other antecedent clauses; else halt the search.*

3.6. Syntactic parallelism preference

Although semantic and pragmatic parallelism (case-persistence, and alignment, respectively, in the discussion above) appear to dominate over syntactic parallelism, the latter plays an important role if two clauses are directly contrasted (e.g., in a coordinate structure, or by means of explicit discourse cohesion markers [14]). Consider the following examples:

> *The girl scout leader paired Mary with Susan, but she had paired her with Nancy last time.* [she=leader, her=Mary]

> *The girl scout leader paired Mary with Susan, but she had paired Nancy with her last time.* [she=leader, her=Susan]

There is no reason to prefer different referents for the pronoun "her" in each sentence above, other than retaining as much as possible the surface syntactic order from the first coordinate clause in the second clause. The strategy here is summarized as follows:

> *In coordinated clauses, adjacent sentences or explicitly contrasted sentences, prefer the anaphoric referent that preserves the surface syntactic role from the first clause.*

3.7. Syntactic Topicalization Preference

Topicalized structures are searched first for possible anaphoric referents. Consider, for instance, the following pseudo-cleft constructions:

> *It was Mary who told Jane to go to New York. Why did she do it?* [she=Mary]
> *It was Jane who went to New York at Mary's bidding. Why did she do it?* [she=Jane]

> *It was Mary who told Peter to go to New York. Why did he do it?* [he=Peter]
> *It was Peter who went to New York at Mary's bidding. Why did he do it?* [he=Peter]

In the first set of examples, describing essentially the same underlying action, the topicalized person becomes the referent of the anaphor "she:" "Mary" in the first sentence, "Jane" in the second. And, the action associated with that person become the referent of "it." However, to stress that topicalization is a preferential rather than categorical strategy, consider the second set of examples above. The exact same semantic and syntactic structures yield "Peter" both times as the referent of "he", because localized constraints so dictate, regardless of who is topicalized. Thus, it is important to distinguish constraints from preferences in anaphora resolution. The topicalization strategy may be stated as follows:

> *Search first a syntactically topicalized part of the candidate antecedent clause (or clauses) for the referent of the anaphor. If an acceptable referent is found, search no further; else search the rest of the clause(s).*

This strategy surprisingly enough has not been exploited in any system to our knowledge, although it is easy to establish syntactic topicalization (indicated by linguistic devices such as fronting, and cleft constructions). In contrast, the much more complex phenomenon of pragmatic topicalization by dialog focus or actor focus (discussed below) was suggested by Sidner [15]. We also believe that dialog

focus can yield a useful preference for anaphoric reference selection, but lacking a computationally-adequate theory for dialog-level focus tracking (Sidner's is a partial theory), we could not yet implement such a strategy.

3.8. Intersentential Recency Preference

Thus far we have focused on the problem of selecting the best anaphoric referent among several candidates, all from a single previous sentence (or coordinated clause). When prior context contains many sentences, the question naturally arises of how far back to search for the anaphoric referent, and how to prioritize that search. At the paragraph (or dialog) level level, we advocate searching sentences in reverse chronological order, applying all the constraints and preferences to select among possible candidates within each sentence. If there are no satisfactory candidates in the previous sentence, then the one before that is considered, and so on. Although we are investigating more sophisticated techniques, these await a more comprehensive (non-linear) theory of discourse structure - and one that is precise enough to permit implementation.

4. Integrating the Strategies

In order to apply a diverse set of strategies, such as those presented in this paper, one needs to make a distinction between constraints (which cannot be violated), and preferences (which discriminate among candidates satisfying all constraints). The latter may be ranked in a partial order (as the goals trees in [4]), or may be offered a voting scheme where the stronger preferences get more votes, and where conflicting preferences of equal voting power indicate true ambiguity.

Our resolution method works by applying the *constraints first* to reduce the number of candidate referents for the anaphor in question. Then, the preferences are applied to each of the remaining candidates. If more than one preference applies, and each suggests different candidate referents for the anaphor in question, all of which have passed the constraint tests, then we consider the anaphor to have a truly ambiguous referent. Thus, when faced with conflicting knowledge sources of equal strength, we simply reduce the space of possible anaphoric referents to those that are accepted by constraints and indicated as preferred by one or more preferences. Earlier hand simulations of a slightly different method[4] on 70 examples (including those presented earlier in this paper) yielded 49 unique resolutions, 17 conflicting possibilities, and 4 anomalous cases. Human judgements correlate very well in terms of identifying the same referent as that suggested by the system in the 49 unique cases.[5] Moreover, the majority of the 17 multiple-referent cases were judged ambiguous by our subjects (the rest required complex world knowledge to establish a unique referent). Therefore, we believe that one can indeed achieve human-like performance with the multi-strategy method of determining referents to anaphors using different sources of linguistic knowledge in a semi-modular fashion.

5. A Practical Implementation

We have developed an anaphor resolver using Local Constraints, Case Role Semantic Constraints, Pre/Postcondition Constraints, Case Role Persistence, Intersentential Recency Preference, and Syntactic Topicalization Preference. The implementation occurs in the context of the Universal Parser (UP) project [6, 16] at the Center for Machine Translation at Carnegie-Mellon University. The UP uses a modified form of lexical-functional grammar [3] unifying syntactic and semantic knowledge sources to produce a complete parse of each sentence. The anaphor resolver operates *post facto* on the set of instantiated semantic case frames and syntactic trees, attempting to resolve anaphors in the parse of the newest sentence using earlier parses (semantic and syntactic) as context to mine for candidate referents. We expect the resolver to become an integral part of our multi-lingual machine translation effort.

Candidate referents are derived by extracting the noun phrases from the most-recent previous sentences that the resolver has processed. The number of sentences examined may be changed, allowing the future addition of discourse phenomena to further restrict the sentences which are examined for candidate referents.

The preferences use a voting method to determine which candidate referent is most preferred. Each preference strategy is given an individual weight, and may vote with less than its full weight for less-preferred candidates, such as case role persistence in a referent several sentences removed from the anaphor.

In addition to ruling out candidates, the case-role and local anaphor constraints may also cast votes for those allowable candidates which are most closely matched to the anaphor or correspond to typical fillers. In effect, these strategies indicate a preference in the absence of hard constraints. For example, the gender constraint would prefer a candidate reference of female gender over one of indeterminate gender when resolving an anaphor of female gender, while at the same time eliminating all candidates of male gender.

After applying the preferences, the most preferred candidate referent is unified with the reference to restrict the range of possible values as much as possible. For example, if *she* is determined to refer to *doctor*, all future anaphoric references to the doctor will be required to have female or unknown gender. However, if multiple candidates have received nearly the same number of votes, the anaphor is considered to be ambiguous.

The anaphor resolver is able to resolve partially-specified definite noun phrases with an antecedent noun phrase. To do so, along with the other local constraints, the head nouns and the remaining slots in the noun phrase are checked for agreement with the reference. The head noun of the candidate must be the same as, or an instance of, the head noun of the reference. For the remaining slots, it suffices for corresponding slots to be unifiable with each other or missing from either the definite noun phrase or the candidate referent. Unlike anaphors, which must have a suitable referent, it is not considered an error if there are no referents which pass all constraints. We believe that the ability to resolve definite noun phrases with basically the same approach as anaphors is an indication of the generality of our strategies and their implementation exploiting semantic and syntactic constraint unification methods.

The current test suite consists of ten examples, totalling 31 sentences containing 27 anaphors and three definite noun phrases with prior referents.[6] The anaphor resolver correctly resolves all but four of the anaphors, and determines the correct referent for all of the definite noun phrases. In two of the four problematic cases, the anaphor is an *it* referring to an action only indirectly mentioned, which is beyond the scope of the resolver. The remaining two anaphors are in the example

> *John carried the box from Bill to Peter. He also sent him Mary's books.*

Here, *him* remains ambiguous, and *he* also remains ambiguous between *John* and *Bill* (with the current voting scheme, John is preferred over Bill).

The following run of the anaphor resolver (edited to save space) illustrates several of the strategies. Each candidate referent is tagged with a number indicating how many votes it has received so far. The intersentential recency preference is applied at the same time that the candidates are collected and tagged because of its computational efficiency; thus, the initial list of candidates already includes the votes from intersentential recency. The case-role persistence preference is applied between pre/postcondition constraints and local constraints, because removal of eliminated candidates (in this implementation) also removes the information on which previous sentence a candidate originates from. Then, case-role constraints are applied, and if multiple candidates remain, the remaining preferences (currently only syntactic topicalization) are applied.

```
sentence 6: The doctor gave John a glass of
            water

(SENT6
    (IS-A *GIVE) (:TIME *PAST) (:AGENT *DOCTOR)
    (:OBJECT OBJECT6) (:RECIPIENT *JOHN))
(*DOCTOR
    (IS-A *PERSON))   [unknown gender]
(OBJECT6
    (IS-A *DRINKING-WATER) (:AMOUNT GLASS 1))
(*JOHN
    (IS-A *PERSON) (:GENDER M)
    (:NUMBER *SINGULAR))

frame = (:RECIPIENT *JOHN)
No referents for definite NP
frame = (:OBJECT OBJECT6)
No referents for definite NP
frame = (:AGENT *DOCTOR)
No referents for definite NP
```

[the frames are unchanged after resolution]

```
sentence 7: John drank it  [it=glass of water]

(SENT7
    (IS-A *INGEST-FOOD) (:TIME *PAST)
    (:AGENT *JOHN) (:OBJECT OBJECT7))
(*JOHN
    (IS-A *PERSON) (:GENDER M)
    (:NUMBER *SINGULAR))
(OBJECT7
    (IS-A *LIQUID) (:PRO +) (:NUMBER *SINGULAR))

frame = (:OBJECT OBJECT7)
Candidates: ((1 :AGENT *DOCTOR)
             (1 :OBJECT OBJECT6)
             (1 :RECIPIENT *JOHN))
after pre-post-cond: ((1 :AGENT *DOCTOR)
                      (1 :OBJECT OBJECT6)
                      (1 :RECIPIENT *JOHN))
after local constr: ((3 :OBJECT OBJECT6))
after case-role constr: ((3 :OBJECT OBJECT6))
referent = (:OBJECT OBJECT6)

frame = (:AGENT *JOHN)
Candidates: ((1 :AGENT *DOCTOR)
             (1 :OBJECT OBJECT61)
             (1 :RECIPIENT *JOHN))
after pre-post-cond: ((1 :OBJECT OBJECT61)
                      (1 :RECIPIENT *JOHN))
after NP agreement: ((9 :RECIPIENT *JOHN))
after local constr: ((12 :RECIPIENT *JOHN))
after case-role constr:((12 :RECIPIENT *JOHN))
referent = (:RECIPIENT *JOHN)
```
[both "John"s are coreferential]

```
(SENT7
    (IS-A *INGEST-FOOD) (:TIME *PAST)
    (:AGENT *JOHN) (:OBJECT OBJECT61))
(*JOHN
    (IS-A *PERSON) (:GENDER M)
    (:NUMBER *SINGULAR))
(OBJECT61
    (:NUMBER *SINGULAR) (:AMOUNT GLASS 1)
    (IS-A OBJECT6))

; sentence 8: He gave him an aspirin
[he=doctor, him=John]

(SENT8
    (IS-A *GIVE) (:TIME *PAST) (:AGENT *HE)
    (:OBJECT OBJECT8) (:RECIPIENT *HE))
(*HE
    (IS-A *PERSON) (:GENDER M)
    (:NUMBER *SINGULAR) (:PRO +))

(OBJECT8
    (IS-A *ASPIRIN) (QUANTITY 1))
```

99

```
frame = (:RECIPIENT *HE)
Candidates: ((1 :AGENT *JOHN)
             (1 :OBJECT OBJECT61)
             (0 :AGENT *DOCTOR)
             (0 :RECIPIENT *JOHN))
after pre-post-cond: ((1 :AGENT *JOHN)
                      (1 :OBJECT OBJECT61)
                      (0 :AGENT *DOCTOR)
                      (1.6 :RECIPIENT *JOHN))
after local constr: ((4.6 :RECIPIENT *JOHN)
                     (2 :AGENT *DOCTOR)
                     (4 :AGENT *JOHN))
after case-role constr:((4.6 :RECIPIENT *JOHN)
                        (2 :AGENT *DOCTOR)
                        (4 :AGENT *JOHN))
referent = (:RECIPIENT *JOHN)

frame = (:OBJECT OBJECT8)
Candidates: ((1 :AGENT *JOHN)
             (1 :OBJECT OBJECT61)
             (0 :AGENT *DOCTOR)
             (0 :RECIPIENT *JOHN))
after pre-post-cond: ((1 :AGENT *JOHN)
                      (1 :OBJECT OBJECT61)
                      (0 :AGENT *DOCTOR)
                      (0 :RECIPIENT *JOHN))
after NP agreement: NIL
after local constr: NIL
No referents for definite NP

frame = (:AGENT *HE)
Candidates: ((1 :AGENT *JOHN)
             (1 :OBJECT OBJECT61)
             (0 :AGENT *DOCTOR)
             (0 :RECIPIENT *JOHN))
after pre-post-cond: ((1 :AGENT *JOHN)
                      (1 :OBJECT OBJECT61)
                      (1.6 :AGENT *DOCTOR)
                      (0 :RECIPIENT *JOHN))
after local constr: ((3 :RECIPIENT *JOHN)
                     (3.6 :AGENT *DOCTOR)
                     (4 :AGENT *JOHN))
after case-role constr: ((3 :RECIPIENT *JOHN)
                         (3.6 :AGENT *DOCTOR)
                         (4 :AGENT *JOHN))
referent = (:AGENT *DOCTOR)

(SENT8
    (IS-A *GIVE)(:TIME *PAST)(:AGENT *DOCTOR4)
    (:OBJECT OBJECT8) (:RECIPIENT *JOHN))

(*DOCTOR4
    (:NUMBER *SINGULAR) (:GENDER M)
    (IS-A *DOCTOR))
        [note that the gender is now known]
(OBJECT8
    (IS-A *ASPIRIN) (QUANTITY 1))
(*JOHN
    (IS-A *PERSON) (:GENDER M)
    (:NUMBER *SINGULAR))
```

6. References

[1] Birnbaum, L. and Selfridge, M.
    Conceptual Analysis in Natural Language.
    In R. Schank and C. Riesbeck (editors), Inside Computer
        Understanding, pages 318-353. New Jersey: Erlbaum
        Assoc., 1980.

[2] Bobrow, R.
    The RUS System.
    In Research in Natural Langauge Understanding. BBN
        Report No. 3837, 1978.

[3] Bresnan, J. and Kaplan, R.
    Lexical-Functional Grammar: A Formal System for
        Grammatical Representation.
    The Mental Representation of Grammatical Relations.
    MIT Press, Cambridge, Massachusetts, 1982, pages 173-281.

[4] Carbonell, J. G.
    Towards a Process Model of Human Personality Traits.
    AI 15(1,2):49-74, November, 1980.

[5] Carbonell, J. G. and Hayes, P. J.
    Natural Language Understanding.
    In Shapiro, S. C. (editor), Encyclopedia of Artificial
        Intelligence, pages 660-677. Wiley & Sons, New York,
        NY, 1987.

[6] Carbonell, J. G., and Tomita, M.
    Knowledge-Based Machine Translation, The CMU Approach.
    In Nirenberg, S. (editor), Machine Translation: Theoretical
        and Methodological Issues. Cambridge, U. Press, 1987.

[7] Carbonell, J. G., Boggs, W. M., Mauldin. M. L. and Anick,
    P. G.
    The XCALIBUR Project, A Natural Language Interface to
        Expert Systems and Data Bases.
    In S. Andriole (editor), Applications in Artificial Intelligence.
        Petrocelli Books Inc., 1985.

[8] Carbonell, J. G.
    Discourse Pragmatics in Task-Oriented Natural Language
        Interfaces.
    In Proceedings of the 21st annual meeting of the Association
        for Computational Linguistics. ACL-83, 1983.

[9] Charniak, E.
    Towards a Model of Children's Story Comprehension.
    PhD thesis, M.I.T., 1972.

[10] Fillmore, C. J.
     The Case for Case.
     In Bach, E. and Harms, R. T. (editors), The Universals of
         Linguistic Theory, pages 1-88. Holt, Rinehart and
         Winston, New York, 1968.

[11] Grosz, B. J.
     The Representation and Use of Focus in Dialogue
         Understanding.
     PhD thesis, University of California at Berkeley, 1977.
     SRI Tech. Note 151.

[12] Hayes, P. J. and Carbonell, J. G.
     A Natural Language Processing Tutorial.
     Technical Report, Carnegie-Mellon University, Computer
         Science Department, 1983.

[13] Hayes, P. J.
     Anaphora for Limited Domain Systems.
     In Proceedings of the Seventh IJCAI, pages 416-422.
         Vancouver, BC, 1981.

[14] Hobbs, J. R.
     A Computational Approach to Discourse Analysis.
     Technical Report 76-2, Department of Computer Science, City
         College, City U. of NY, 1976.

[15] Sidner, C. L.
     Focusing for Interpretation of Pronouns.
     Journal of Computational Linguistics 7:217-231, 1981.

[16] Tomita, M. and Carbonell, J. G.
     The Universal Parser Architecture for Knowledge-Based
         Machine Translation.
     In Proceedings of IJCAI-87. Milan, Italy, 1987.

[17] Webber, B. and Reiter, R.
     Anaphora and Locial Form: On Formal Meaning
         Representations for Natural Langauge.
     In Proceedings of the Fifth IJCAI, pages 121-131.
         Cambridge, MA, 1977.

[18] Wilks, Y.
     Knowledge Structures and Language Boundaries.
     In Proceedings of the Fifth International Joint Conference on
         Artificial Intelligence, pages 151-157. IJCAI-V, 1977.

Notes

[1] Interclausal anaphora in coordinate constructions behaves much like a constrained version of intersentential anaphora, where syntactic parallelism (between the coordinated clauses) plays a more dominant role.

[2]No claims, however, are made for the relative frequency or utility of resolving intersentential vs intrasentential anaphors in processing narrative or expository texts.

[3]Although many of our anaphora instances come from actual user utterances in our experience with domain-oriented human-computer interfaces, we expect that the strategies developed here are of more general applicability. For clarity of exposition in this paper, we have selected examples not from our human-computer dialogs, but from everyday events.

[4]Using preferences to determine which candidates are tested against the constraints

[5]Often, more than one strategy suggested the same referent, increasing our confidence. Language is redundant, and it may prove useful to exploit that redundancy.

[6]The sentences in our corpus used to test the implementation are:

John gave Mary two aspirin. She took them from him.

Mary had a headache. John gave her two aspirin tablets. She took them.

The doctor gave John a glass of water. John drank it. He gave him an aspirin. He took it with another glass of water.

Mary gave an apple to Susan. John also gave her an orange.
Mary gave an apple to Susan. She also gave John an orange.

John took the cake from the table. He ate it.
John took the cake from the table. He washed it.
John took the cake from the table [ambig]. He washed it.

John carried the box from Bill to Peter. He also sent him Mary's books.

It was Mary who told Jane to go to New York. Why did she do it?
It was Jane who went to New York at Mary's bidding. Why did she do it?

John gave Peter an apple. He ate it.

Jack (age 10) went up the hill. John (age 32) went up the hill.
The boy fell down.
Jack went up the hill. The boy fell down.