Anaphor Resolution and the Scope of Syntactic Constraints

Roland Stuckardt
German National Research Center for Information Technology (GMD)
KONTEXT - Natural Language Systems
Dolivostraße 15, D-64293 Darmstadt, Germany
stuckard@darmstadt.gmd.de

Abstract

An anaphor resolution algorithm is presented which relies on a combination of strategies for narrowing down and selecting from antecedent sets for reflexive pronouns, nonreflexive pronouns, and common nouns. The work focuses on syntactic restrictions which are derived from Chomsky's Binding Theory. It is discussed how these constraints can be incorporated adequately in an anaphor resolution algorithm. Moreover, by showing that pragmatic inferences may be necessary, the limits of syntactic restrictions are elucidated.

1 Introduction

It is by now widely agreed upon that the process of resolving anaphors in natural language text is supported by a variety of strategies employing different kinds of knowledge. The process of *determining the set of possible antecedents* is governed by morphosyntactic, syntactic, semantic, and pragmatic restrictions. The same holds for preferences applied in the *antecedent selection* process: simple surface criteria are involved as well as more elaborate syntactic, semantic, or focusing heuristics. As a consequence, recent approaches to anaphor resolution apply a carefully selected blend of constraints and preferences, thus constituting multi-strategy approaches in the sense of Carbonell and Brown (Carbonell and Brown, 1988).

There are, however, implementability limitations. At discourse level, determining the set of admissible antecedents requires a representation which is ordered according to pragmatic relations (Grosz and Sidner, 1986; Webber, 1989). Although various theoretical frameworks have been suggested, the *recognition* of these relations in the case of unrestricted discourse is still beyond the state-of-the-art. Moreover, there are cases in which antecedent decisions can only be made on the grounds of domain knowledge and inferencing, and although there have been various attempts to integrate components of these kinds into anaphor resolution approaches, a satisfying solution to this problem is not available by now.

As a consequence, current anaphor resolution implementations rely on constraints and preference heuristics which employ information originating from morphosyntactic, syntactic, or *shallow* semantic analysis (cf. (Carter, 1987)). These approaches, however, perform remarkably well. An early case study revealed that a 'naive' algorithm for resolving nonreflexive pronouns, by relying merely on morphosyntactic, syntactic, and surface criteria, yields correct results for more than 80 percent of pronoun occurrences, and that the incorporation of selectional constraints results in a gain of another 3.5 percent of accuracy (Hobbs, 1978). These results have been confirmed by recent work (Lappin and Leass, 1994). The latter approach is based on a more elaborate, theory-oriented, declarative formulation of the syntactic constraints, and handles reflexive pronouns too. It proved that the incorporation of statistically measured lexical preference patterns (a dynamic, domain specific substitute for the static encoding of selectional preferences) yields a gain of only 3 per cent, and a sole application of lexical preference patterns resulted in a performance below 35 per cent. Hence, there is strong evidence that syntactic restrictions in combination with surface based and syntactic preference criteria play the central role in realistic approaches to anaphor resolution.

In this paper, an anaphor resolution algorithm is described which has been implemented as part of the KONTEXT text analysis system for the German language (Haenelt, 1994). The emphasis lies on the description of implementation techniques for syntactic constraints. Section 2 works out strategies which are applied, focusing on the theoretical background from which the syntactic constraints emerge. Section 3 describes how these strategies are coined into an algorithm for the resolution of reflexives, nonreflexive pronouns, and definite common nouns, thereby elucidating details which have to be taken into account in an adequate implementation. In section 4, a theoretical evaluation is performed, and application results are given. Section 5 points out that the structural constraints may depend on circumstances which are not a matter of syntax alone, but rather necessitate semantic and pragmatic inferencing. As a consequence, limitations concerning the implementability show up, and the scope of syntactic constraints proves to be restricted.

2 Constraints and Preferences

2.1 Morphosyntactic Agreement

A quite strict constraint requires the pronoun to agree with its antecedent in person, number, and gender. In example[1]

(1) The father visited his daughter.
She had invited him on Sunday.

the antecedents for *him* and *she* are identified uniquely as *father* and *daughter*, respectively.

2.2 Syntactic Constraints

The following data substantiate the syntactic restrictions which are to be employed:

(2a) The barber$_i$ shaves himself$_i$.
*(2b) * The client$_i$ appreciates that the barber shaves himself$_i$.*

These examples suggest that reflexive pronouns choose their antecedents in some kind of local domain. On the other hand, examples

*(3a) * The barber$_i$ shaves him$_i$.*
(3b) The client$_i$ appreciates that the barber shaves him$_i$.

indicate that the admissible structural positions of antecedents for nonreflexive pronouns are distributed complementarily, i.e. these pronouns choose their antecedents outside of their local domain. An even more stringent restriction holds for nonpronominal nouns:

*(4a) * The barber$_i$ shaves the barber$_i$.*
*(4b) * The client$_i$ appreciates that the barber shaves the client$_i$.*

But even here, configurations exist in which intrasentential antecedents are possible:

(4c) The barber who shaved the client$_i$ told the client$_i$ a story.

Chomsky provides a formal description of these observations as part of his Government and Binding (GB) Theory (Chomsky, 1981; Chomsky, 1986). Binding Theory (BT) distinguishes three types of NP, namely type A ('anaphor', comprising reflexives and reciprocals[2]), type B (nonreflexive pronouns), and type C ('referring' expressions, comprising common nouns and names). The restrictions are stated as *binding principles*:

Definition 1 (binding principles)

(A) An anaphor is bound in its binding category.

(B) A pronominal is free (i.e. not bound) in its binding category.

(C) A referring expression is free (i.e. not bound) in any domain.

where *binds* is a relation which is defined on the NP nodes of the (surface) phrase structure tree:

Definition 2 (the binding relation) *Node X* binds *Node Y if and only if X and Y are coindexed and X c-commands Y.*

where (definitions vary slightly):

Definition 3 (the c-command relation) *Node X* c-commands *node Y if and only if the next branching node which dominates X also dominates Y and neither X dominates Y, Y dominates X nor X = Y.*

The central part of the Binding Theory develops the notion of local domain to which binding principles A, B, and C refer as *binding category*:

Definition 4 (binding category) *Node X is* binding category *of node Y if and only if X is the next node which dominates Y, and which contains a subject that c-commands Y.*

Due to these definitions, the acceptability judgements for the data presented above are reproduced by binding principles A, B, and C. For each example, the subject demarcating the (local) binding category is just the ordinary subject of the subordinate clause. (One has to recall that, in phrase structure trees, the subject c-commands the content of the VP.) The notion of subject, however, is a more general one, applying also to some kinds of nominal phrase attributes, in particular certain variations of genitives and possessives:

(5) Peter listens to Sam's$_i$ story about himself$_i$.

[1]The examples are given in English. The phenomena and its implications translate directly to German.

[2]In this paper, the notion of *anaphor* is used more generally. When referring to *anaphor* in the Chomskyan sense, the notion *reflexive/reciprocal (pronoun)* is used.

2.3 Antecedent Predictability

For *cataphoric* pronominal resumptions, a constraint is applied which has been described by Kuno (Kuno, 1987). According to

(6a) The barber who shaved him_i *told the* $client_i$ *a story.*
*(6b) * The barber who shaved* him_i *told a* $client_i$ *a story.*

a *definiteness requirement* has to be fulfilled, ruling out antecedents which are not *predictable*, i.e. not already introduced in the discourse.

2.4 Case Role Inertia

In general, the constraint application will not single out a unique antecedent. Depending on the type of anaphor to be resolved, preferences are applied, comprising the rather superficial and self-explanatory criteria of recency, cataphor penalty, and subject preference. The *case role inertia criterion*, which proved to be very useful in practice, is explainable by the following example:

(7) Peter visited his brother.
He showed him his new car.

Unless given further information, there seems to be a strong tendency to choose the antecedents in a way that the syntactic and/or semantic case roles of the pronouns reproduce the corresponding roles of their antecedents. Thus, the preference rule suggests *Peter* as the antecedent for *He*, and *brother* as the antecedent for *him*. As can be demonstrated by further examples (e.g. changing from active to passive voice or vice versa), retaining the semantic case role should outvote retaining the syntactic case role. In cases in which semantic case is not available, however, promoting syntactic case parallelism serves as a good approximation.

In its effect, this preference rule approximates the often suggested heuristic of keeping rather then shifting referential focus (cf. (Sidner, 1983)).

3 Towards the Algorithm

The main question concerns the adequate implementation of Chomsky's binding principles. Some a priori remarks on theoretic subtleties and on the employed representation are in place.

3.1 Interdependency Sensitiveness

As stated by (Correa, 1988), an immediate implementation of the constraints proposed in Binding Theory is unfeasible. Chomsky states, merely as a theoretical device, a free indexing rule which randomly assigns reference indexes to surface structure NP nodes. During mapping to the semantic LF (logical form) representation, the binding principles serve as restrictions for filtering out the index distributions which are considered valid when interpreted as coreference markers. A direct implementation of this generate-and-test procedure yields an exponential time complexity.

Current approaches avoid generate-and-test by resorting to different strategies. According to the most common technique, for anaphoric NPs, a *separate* antecedent search is performed, resulting in a quadratic time complexity (e.g. (Hobbs, 1978; Strube and Hahn, 1995)). Because, however, the antecedent decisions are performed in isolation, invalid index distributions may arise. In example

(8a) The $barber_i$ *told the* $client_j$ *a story, while* he_k *shaved* him_l*.*

neither of the pronouns is confined structurally to one of the intrasentential antecedent candidates in the matrix clause. But, after a first decision, e.g.

(8b) The $barber_i$ *told the* $client_j$ *a story, while* he_i *shaved* him_l*.*

the situation changes, for one of the antecedent options of the still unresolved pronoun is no longer available. Binding principle B may be violated:

*(8c) * The* $barber_i$ *told the* $client_j$ *a story, while* he_i *shaved* him_i*.*

An interdependency between antecedent choices may arise as well when choosing between discourse antecedents, or as a consequence of relative clause attachment, which predetermines coindexing.

The approach presented below is sensitive to these decision interdependencies, while avoiding the exponential time complexity of an immediate binding constraint implementation. This is achieved by supplementing the straightforward sequential strategy with a dynamic reverification of the binding restrictions in the antecedent selection step. To avoid that desirable antecedent options are ruled out by interdependency, the choices with highest plausibility is given preference to.

3.2 Representing Surface Structure

The original statement of Binding Theory forms part of GB Theory, in which a broader set of interacting principles is formulated. Because the aim of anaphor resolution for a specific language is restricted, the representation can be simplified. Complicating details which result from the GB claim to universality may be omitted.

Besides being efficiently searchable, the simplified surface structure has to represent the structural details which are necessary for the verification of the binding restrictions. In particular, this comprises subject-object-asymmetry, the demarcation of local domains, and surface order dependent structural variations[3].

Because the KONTEXT text analysis system is based on a dependency grammar, a *mapping process* generates the required representation from a dependency tree, which is not suitable for a structural verification of the binding principles, because vital details are not structurally visible. The attempt of directly verifying BT restrictions on dependency structure, as suggested by Strube and Hahn (Strube and Hahn, 1995), does not seem adequate, because important details are ignored.

The structures which were generated for some of the above examples are as follows:[4]

(9a) (S barber (VP himself)) →*(2a)*
(9b) (S client
(VP (STHAT barber (VP him)))) →*(3b)*
(9c) (S barber (SREL who (VP client))
(VP client (VP story))) →*(4c)*

The marker nodes *STHAT* and *SREL* are delimiters of local domains, to which the binding principle verification functions are sensitive.

Special techniques are employed in representing local NP domains, which are introduced by deverbative NPs and NPs with possessive markers (saxonian genitive, genitivus possessivus, possessive pronoun, or certain attributive PPs), e.g.

(10) The barber hears his_i *story about* $himself_i$.

(S barber
(VP $story_j$
(SVATT x_story_j
(ATT his (ATT (PP himself)...)

A domain *SVATT* enforcing local reflexivation is opened. The NP *barber* and the reflexive pronoun *himself* may be coindexed only indirectly via the possessive pronoun *his*, which is of type B, and hence forced to take a nonlocal antecedent. In accordance with intuitive judgement, a local instance of the NP $story_j$ blocks the coindexing of the possessive pronoun and its dominating noun. Here again, the mechanism which copes with interdependencies is applied.[5] Technically, new NP types C' (example (10)) and B' (relative pronoun, cf. section 3.1) are introduced for which binding principles C and B are verified, respectively, but for which no antecedent search is performed.

3.3 The Algorithm

The KONTEXT anaphor resolution algorithm, as shown in figure 1, consists of three phases: constraint application, preference criteria application and plausibility sorting, and antecedent selection including reverification of constraints which may be involved in decision interdependencies.

Two binding constraint verification procedures are employed which differ in the handling of type A NPs. According to binding principle A, a reflexive pronoun requires 'constructively' a local antecedent (step 1(b)i). Example (10), however, illustrates that further nonlocal coindexings are admissible. This gives rise to a weak version of binding constraint verification, the usage of which is of vital importance to the functioning of the interdependency test step 3b.

4 Evaluation

As a proper base for comparison, the theoretical analysis is restricted to the contribution of intrasentential antecedent search. Let n be the number of NP nodes in the surface structure representation. Because the number of anaphoric NPs and intrasentential candidates is bounded by n, and the individual a priori verifications of the binding principles contribute costs proportional to the number of nodes in the surface structure tree, the worst case time complexity of step 1 is $O(n^3)$. A similar analysis, assuming a clever handling which prevents individual interdependency checks from being done more then once, reveals that the complexity of step 3 is $O(n^3)$ too. Therefore, since the scoring and sorting step 2 does not exceed this limit, the overall worst case complexity is $O(n^3)$.

In tests on architect biographies drawn from (Lampugnani, 1983), the algorithm correctly resolved approximately 90 per cent of type B pronouns (including possessives), and, as expected, all occurrences of reflexives, which occur quite scarcely in the test corpus. The set of possible antecedents tends to be reduced drastically during constraint application. Interdependency collisions did not happen too frequent. This tendency is strongly supported by the case role inertia heuristic, which promotes a complementary distribution of preferred antecedents for type B pronouns cooccurring in a domain of binding.

The strategy of considering the more plausible antecedent choices first does not eliminate interdependency collisions in general, and, moreover,

[3] This concerns certain cases of subject and object clause extraposition as well as, in particular, the object NPs contained in the VP, for which a right branching structure is generated, yielding a base for a structural determination of admissible antecedents for reflexive pronouns, which is mainly governed by subject-object asymmetry and surface order.

[4] Implementation details are ignored.

[5] This technique resembles the use of traces in Chomsky's GB theory. Because of its restricted aim, however, it is much simpler.

1. For each anaphoric NP Y, determine the set of possible antecedents X:
   (a) Verify morphosyntactic or lexical agreement with X (congruence in person, number, and gender, lexical recurrence etc, depending on the type of Y)
   (b) If the antecedent candidate X is intrasentential, check whether the binding principles of Y *and* X are satisfied: for the proposed coindexing,
      i. verify that the binding principle of Y is *satisfied constructively*,
      ii. verify that the binding principle of X is *not violated*.
   (c) If Y is a type B pronoun, antecedent candidate X is intrasentential, and, according to surface order, X *follows* Y (i.e. the resumption would be *cataphoric*), verify that X is *definite*.
2. Plausibility scoring and sorting:
   (a) For each surviving pair (Y_i, X_j) of anaphor and antecedent candidate: determine the numerical plausibility score $v(Y_i, X_j)$, which ranks X_j relatively to Y_i, based on case role inertia, recency, cataphor penalty, and subject preference, depending on the type of Y_i.
   (b) *(local sorting)* For each anaphor Y: sort their individual antecedent candidates X_j according to decreasing plausibility $v(Y, X_j)$.
   (c) *(global sorting)* Sort the anaphors Y according to decreasing plausibility of their individual best antecedent candidate.
3. Antecedent Selection: Consider anaphors Y in the order determined in step 2c. Suggest antecedent candidates $X(Y)$ in the order determined in step 2b. Select $X(Y)$ as candidate if there is no interdependency, i.e. if
   (a) the morphosyntactic features of Y and $X(Y)$ are still compatible[6],
   (b) for each NP Z whose coindexing with $X(Y)$ has been determined in the *current* invocation of the anaphor resolution algorithm: the coindexing of Y and Z which results as a side effect when chosing $X(Y)$ as antecedent for Y does *not violate* the binding principles.

   To allow for an efficient detection of interdependencies, store the selected antecedent separately from coreferent occurrences contributed by earlier invocations of the algorithm.

Figure 1: The KONTEXT Anaphor Resolution Algorithm

does not guarantee that the global maximum of plausibility is reached. Because of its practical performance, however, it proved to be a satisfactory substitute for the generate-and-test strategy.

5 Exploring the Limits

The determination of the substructure describing a local domain is not always easy. Whereas for NPs with possessive markers (cf. example (10)) the matter tends to be clear, a common source of difficulties emerges from adjectivally used participles and from *deverbative* NPs. In the latter case, e.g. a genitival attribute may instantiate, depending on the NP, either the subject (*genitivus subjectivus*) or the object (*genitivus objectivus*) (for German, cf. (Teubert, 1979)). As the following examples demonstrate, it is insufficient to know merely about the existence of a local domain. In general, it is necessary to determine the instantiation of its participants, but this, at least in certain cases, involves pragmatic inferencing.

(11a) Paul$_i$ accepts the decision for him$_i$.
*(11b) * Paul$_i$ accepts the decision for himself$_i$.*

According to acceptability judgements, *decision* introduces a local binding domain. But a change of the matrix clause verb leads to a different judgement, while the syntactic structure is preserved:

(12a) Paul$_i$ revises the decision for him$_i$.
(12b) Paul$_i$ revises the decision for himself$_i$.

The clue lies in the observation that a pragmatic restriction is governing the instantiation of the implicit local subject in examples (11), but not in examples (12). In (11a), due to the obvious conclusion that someone who *accepts* an action is not the conscious actor of it, *Paul* is pragmatically ruled out as the local subject of the *decision* domain. On the other hand, *revise* leaves open whether *Paul* or someone else is the decider. This explanation is confirmed by the following data:

(13a) Paul$_i$ revises Sam's$_j$ decision for him$_i$.
*(13b) *Paul$_i$ revises Sam's$_j$ decision for himself$_i$.*
*(13c) *Paul$_i$ revises his$_i$ decision for him$_i$.*
(13d) Paul$_i$ revises his$_i$ decision for himself$_i$.

Current approaches (Strube and Hahn, 1995; Lappin and Leass, 1994) ignore this subtlety by merely taking into account NP domains which are established by possessive determiners. As a consequence, wrong results may be obtained, e.g. in case of example (11a), as there is no possessive modifier, *Paul* will not be considered to be an antecedent candidate for *him*. With these difficulties in mind, questionable antecedent decisions may be

[6] In German, this kind of interdependency may arise, due to morphosyntactic ambiguity, in case of multiple occurrences of the pronoun *sie*.

marked as depending on particular local instantiations, by this means providing a starting point for more comprehensive considerations which take into account the relation between structural restrictions and the resolution of ellipsis.

6 Conclusion

Starting with a recapitulation of current work on anaphor resolution, it was argued for an approach which bases on syntactic restrictions.

The original formulation of Chomsky's Binding Theory proved to be unsuitable for immediate implementation. Straightforward approaches may fail in cases in which interdependencies between antecedent decisions arise. Based on this observation, an algorithm has been presented which, on the one hand, is interdependency-sensitive, but, on the other hand, avoids computational unfeasibility by following a strategy according to which the choices with the highest plausibility are considered first. For each decision, its dynamic compatibility with the earlier (more plausible) decisions is verified. The practical behaviour of the algorithm fulfilled the expectations.

There are, however, limitations to the scope of syntactic constraints. It has been demonstrated that, in general, the construction of appropriate representations for binding domains may necessitate semantic or pragmatic inferencing.

A topic which should be subject of further research is the interdependency between parse tree construction and anaphor resolution. Up to now, it has been assumed tacitly that, at the time of binding constraint application, the surface structure representation is available. The construction of this representation involves disambiguation decisions (relative clause attachment, prepositional phrase attachment, and uncertainty of syntactic function), which, due to their structure determining effects, may interfere with the antecedent options of anaphor resolution (cf. (Stuckardt, 1996)). At current, the KONTEXT text analysis system employs a processing model according to which parsing is performed prior to anaphor resolution. Because of the interdependency between parsing and anaphor resolution, however, these two problem classes should be handled at one stage of processing rather than sequentially.

References

Jaime G. Carbonell, Ralf D. Brown. 1988. *Anaphora Resolution: A Multi-Strategy Approach.* In: *Proc. 12th Int. Conf. on Computational Linguistics, COLING*, Vol. I, 96-101, Budapest.

David Carter. 1987. *Interpreting Anaphors in Natural Language Texts.* Ellis Horwood Series in Artificial Intelligence, Chichester.

Noam Chomsky. 1981. *Lectures on Government and Binding.* Foris Publications, Dordrecht.

Noam Chomsky. 1986. *Knowledge of Language.* Praeger, New York.

Nelson Correa. 1988. *A Binding Rule for Government-binding Parsing.* In: *Proc. 12th Int. Conf. on Computational Linguistics, COLING*, Vol. I, 123-129, Budapest.

Barbara J. Grosz, Candace L. Sidner. 1986. *Attention, Intention and the Structure of Discourse.* In: *Computational Linguistics*, Vol. 12, Number 3, 175-204.

Karin Haenelt. 1994. *Das Textanalysesystem KONTEXT. Konzeption und Anwendungsmöglichkeiten.* In: *Sprache und Datenverarbeitung*, 18, 1, 17-31, 1994.

Jerry Hobbs. 1978. *Resolving Pronoun References.* In: *Lingua*, 44, 311-338.

Susumu Kuno. 1987. *Anaphora and Discourse Principles.* In: Makoto Nagao (ed). *Language and Artificial Intelligence.* Elsevier, 1987.

Vittorio M. Lampugnani (ed). 1983. *Lexikon der Architektur des 20. Jahrhunderts.* Hatje, Stuttgart.

Shalom Lappin and Herbert J. Leass. 1994. *An Algorithm for Pronominal Anaphora Resolution.* In: *Computational Linguistics*, 20 (4), 535-561.

Candace L. Sidner. 1983 *Focusing in the Comprehension of Definite Anaphora.* In: *Michael Brady, Robert C. Berwick (eds). Computational Models of Discourse.* M.I.T. Press, Cambridge, MA, 267-330.

Michael Strube, Udo Hahn. 1995. *ParseTalk about Sentence- and Text-Level Anaphora.* In: *Proc. 7th Conf. of the European Chapter of the Association for Computational Linguistics, EACL*, 237-244. Dublin, March 28-31, 1995.

Roland Stuckardt. 1996. *An Interdependency-Sensitive Approach to Anaphor Resolution.* To appear in: *Proc. of DAARC96 - Discourse Anaphora and Anaphor Resolution Colloquium.* Lancaster University, 17-18th July, 1996.

Wolfgang Teubert. 1979. *Valenz des Substantivs. Attributive Ergänzungen und Angaben.* Pädagogischer Verlag Schwann.

Bonnie L. Webber. 1989. *Deictic Reference and Discourse Structure.* Department of Computer and Information Science, University of Pennsylvania, Philadelphia PA 19104-6389.