Pronoun Resolution in Unrestricted Text

Kari Fraurud

Nordic Journal of Linguistics / Volume 11 / Issue 1-2 / June 1988, pp 47 - 68
DOI: 10.1017/S0332586500001748, Published online: 22 December 2008

Pronoun Resolution in Unrestricted Text

Kari Fraurud

Fraurud, K. 1988. Pronoun Resolution in Unrestricted Text. *Nordic Journal of Linguistics* 11, 47–68.

Quantitative and qualitative studies of referential relations in unrestricted natural text are necessary both for a better theoretical understanding of referential processes, and for the development of empirically well-founded algorithms for anaphora resolution in the framework of natural language processing (NLP) systems. The aim of the study reported in this paper was to provide preliminary empirical data on anaphoric pronouns in Swedish. The relation between the pronoun and its antecedent was studied for 600 pronouns in three different types of unrestricted written Swedish text, and a simple pronoun resolution algorithm was tested on the sample.[1]

Kari Fraurud, Institute of Linguistics, University of Stockholm, S-106 91 Stockholm, Sweden.

1. INTRODUCTION

A procedure for interpreting anaphoric expressions, such as certain pronouns and definite NPs, is an essential component in any NLP system. In order to, for example, extract information about a certain entity from a machine readable written text, the system must be able to identify all noun phrases and pronouns referring to that entity. While human readers seldom have any difficulty in understanding an anaphoric expression without being aware of the great amount of linguistic knowledge and world knowledge used in the interpretation, the formulation of a computer algorithm for the automatic identification of the referential relationships within a text is a difficult task. It is an exciting theoretical challenge to try to get as far as possible in making explicit the rules and preferences underlying the human reader's comprehension of these relations. Within the fields of linguistics and artificial intelligence, a large number of mechanisms involved in the comprehension of anaphora have been suggested, including syntactic constraints on intrasentential anaphora (Reinhart 1983), selectional restrictions, formal rules for quantification (Webber 1978), discourse structure (Grosz 1977, Reichman 1981), focussing (Grosz 1977, Sidner 1979).

In the work on anaphora resolution, there are (with a rough division) on the one hand the applied systems, which make use

of a set of heuristic rules in order to achieve a reasonable degree of correctness in a program working within a strongly restricted discourse domain (e.g. Winograd 1973). On the other hand, there are the theoretical models, which assume a number of highly elaborate and yet unimplemented procedures in an effort to cover all cases of pronoun reference (e.g. Sidner 1979). However, few pronoun resolution algorithms have, to my knowledge, been tested on unrestricted natural texts. One exception is Guentner & Lehman's (1983) set of rules for the interpretation of pronouns, which is reported to work remarkably well for some 600 pronoun occurrences in a text corpus. (No figures or further specifications of text type etc. are, however, provided.) Another exception came to my knowledge when I had just finished the present study. Jerry Hobbs (1978) studied 300 pronouns in three types of written English texts. His study is very similar to mine and I will return to it shortly below.

The evaluation of NLP models in relation to unrestricted natural texts is important both from the point of view of developing applicable NLP systems and from a theoretical point of view. I will here try to illustrate a cost-benefit approach to the problem of pronoun resolution, which may simply be described as trying to do as much as possible with the simplest possible algorithm. The simple pronoun resolution algorithm evaluated in this study is based on some few principles, maximally utilizing signals in the surface text and overlooking a number of factors that may later be considered. The algorithm is tested on unrestricted natural texts in order to produce a sample of errors, which is used as a basis for further quantitative and qualitative analysis, and for a discussion of how the algorithm can be developed in order to reduce the cases where it fails to point out the correct antecedent of a given pronoun. In this discussion, quantitative considerations should play an important role, implying a careful analysis of the costs and benefits of each additional elaboration. This means, for example, that cases which are infrequent and whose resolution put high demands in terms of input text, lexicon and the knowledge involved, are given less priority. Furthermore, the development of the algorithm should be guided by other quantitative studies of referential relations in natural text.

A quantitative approach to linguistic phenomena is also theoretically motivated, as pointed out by Hakulinen and Karlsson

(1980): ". . . naive speakers have at least a rough statistical competence providing some intuitive basic knowledge about frequent and infrequent phenomena. Furthermore, there are grammatical phenomena that are inherently gradient, or at least not 'purely' qualitative". The preferences involved in the interpretation of anaphoric relations is a typical example of such phenomena, which cannot be accounted for only by categorical rules.

In section 2 of this paper, I give a short account of the forms of Swedish pronouns and describe the texts used in the study and the sample of pronoun occurrences. The following section presents the results of a quantitative study of the scope of discourse referents in relation to their animacy. In the fourth section, the simple pronoun resolution algorithm is presented and its performance is evaluated.

2. THE SAMPLE

2.1. *Animacy and Gender*

For pronouns, Swedish has two distinctions with respect to the animacy of the referent; animate/inanimate, and, within animates, masculine/feminine. In addition, there is a (grammatical) gender distinction within the inanimates; viz. t-gender/n-gender (also called 'neuter'/'uter'). In summary, there are four classes of pronouns with respect to animacy and gender: masculine, feminine, t-gender and n-gender. The subject, object and possessive forms of the main forms are the following:

Masculine: han (s), honom (o), hans (p)
Feminine: hon (s), henne (o), hennes (p)
T-gender: det (s/o), dess (p)
N-gender: den (s/o), dess (p)
The plural form is the same in all four classes:
de (s), dem (o), deras (p)

In addition to the main forms, there is a subset of the so-called demonstrative pronouns, which serve as demonstratives in parallel to the pure deictic forms 'den/det här' (this (one)), and 'den/det där' (that (one)), as well as being used anaphorically under certain circumstances. In the following, I will refer to these pronouns as the *semi-demonstratives*:

Masculine:	denne (s/o), dennes (p)
Feminine/N-gender:	denna (s/o), dennas (p)
T-gender:	detta (s/o), dettas (p)
Plural (all classes):	dessa (s/o), dessas (p)

In traditional normative grammar, a writer of Swedish is advised to use the main forms 'han' etc. only when referring to the subject of the preceding sentence, and to use the semi-demonstratives when referring to non-subjects (see e.g. Wellander 1973). The use of semi-demonstratives is, however, not obligatory (it is, for instance, seldom used in informal prose and not at all in spoken language). And, when using these forms, writers do not strictly adhere to the normative rule.

The masculine and feminine pronouns are used for humans and sometimes for animals ascribed an individuality/personality, the choice of form being determined by the sex of the referent. In some varieties of regional colloquial Swedish, these pronouns are also used for certain inanimate objects such as the sun, the clock or a boat (feminine), and the moon (masculine).

The t-gender and n-gender pronouns are used for non-human referents such as objects and animals in general. The form of pronoun is chosen according to the lexically determined gender of the corresponding noun. As a consequence, the same object may be referred to by a t-gender or n-gender pronoun, depending on which of two synonymous nouns the speaker/writer has in mind. For example:

(1) *Huset*$_i$ är mycket gammalt. *Det*$_i$ ska rivas.
(2) *Kåken*$_i$ är mycket gammal. *Den*$_i$ ska rivas.
The *house/shack*$_i$ is very old. *It*$_i$ is to be torn down.

Lexically, also nouns referring to humans have t- or n-gender. For n-gender nouns, the pronoun is *han* or *hon* depending on the sex of the referent. The same holds for most of the (considerably fewer) t-gender nouns referring to humans, but for a limited set of these nouns, such as *child* and certain nouns of occupation, a t-gender pronoun may be used instead of an animate pronoun. T-gender pronouns are also used to refer to propositions and the like.

In the excerption of pronouns, all occurrences of the pronoun *forms* enumerated above were included (excluding, for example,

reflexive and relative pronouns, which were not treated in this study). The only exception was an (in most cases) easily delimitable set of clearly non-referential uses, namely the occurrences of the t-gender form 'det' when used as a formal subject. All the remaining occurrences were regarded as potentially referential uses of pronouns.

2.2. *Texts*

Three types of text were studied. The first texts were passages from short stories, which can be seen as fairly typical representatives of narrative text, with a couple of human main actors, a number of human participants and a large set of non-human props.[2] The second type of texts were reports of the procedure of a court of justice including paragraphs of both narrative and informational character with mainly human main referents and relatively few non-human secondary referents.[3] The third text sample consisted of articles about technological inventions; common non-fiction, informational texts with almost no human referents, but a number of non-human 'main actors', or primary referents, such as *the industrial robot IRB 1000* and secondary referents, such as *the pincers* (of IRB 1000).[4] In the following, I will refer to the three text samples as the *stories*, the *reports* and the *articles*, respectively.

The syntactic complexity and the mean sentence length differed considerably in the three text types. In the articles, there were many long sentences with subordinate clauses and nominalised constructions, while the sentences in the stories were comparatively short and simple. The sentence length and complexity in the reports was something in-between that of the other two text types. The frequency of pronouns was much lower in the articles than in the other texts. With a rough estimate, the stories contained about seven times more pronouns per 100 words than the articles.

2.3. *Pronoun Occurrences*

The sample consisted of the first 200 pronoun occurrences from randomly chosen passages of each of the three text types, i.e. a total of 600. A first parameter of relevance to pronoun resolution is the type of referent that the pronoun refers to. Here I distinguish between 'human' referents (including also animals referred to by

Table 1. Types of referents. Distributions of pronouns in three text samples according to type of referent.

	Human		Object		Proposition		Total	
	N	%	N	%	N	%	N	%
Stories	186	(93.0)	11	(5.5)	3	(1.5)	200	(100.0)
Reports	157	(78.5)	34	(17.0)	9	(4.5)	200	(100.0)
Articles	12	(6.0)	149	(74.5)	39	(19.5)	200	(100.0)
Total	358	(59.7)	94	(32.3)	51	(8.5)	600	(100.0)

animate pronouns), 'object' referents (including also abstract entities introduced by a NP), and 'propositions' (see below). The distributions are shown in Table 1. Since the stories are about people and the articles about things, it is quite natural that the distribution of pronouns with human and object referents differ between the text types. However, considering the fact that the frequency of all pronouns was very low in the articles as compared to the other texts, these figures point to the more interesting observation that, in general, inanimate pronouns are much less frequent than animate pronouns (in terms of pronouns/100 words).

The pronouns here characterized as having a propositional type of referent present a serious problem in anaphora resolution. It is important to recognize that even the human interpreter's answer to the question of what such pronouns refer to may be quite vague. Sometimes it is possible to point out a clause or sentence as the antecedent of the pronoun, but more often the pronoun refers to a larger situation, a whole sequence of events or the like. For the moment, it would be a step forward if we could find a way to identify *when* a t-gender pronoun has this type of referent, so that the search for an NP antecedent would be blocked. Many of the pronouns in this sample could, for example, be identified as propositional, being subject of the following verbs: *innebära* 'imply', *medföra* 'bring about', *leda till*, 'lead to'. I have, however, found the sample in this study too small to make any further generalizations or suggestions on how propositional anaphora shall be treated in an NLP system. In the following, the 51 propositional pronouns were excluded, leaving a sample of 549 pronoun occurrences for further investigation.

3. SCOPE AND ANIMACY

In order to be accessible to pronoun reference at a particular point in a text, a discourse referent has to be, in some sense, active in the reader's mental model of the discourse at that point. One of several factors that can be assumed to affect a referent's degree of activity is its participation in recently described events and situations. In the text, this presence of a referent is often implicit, but it is also partly reflected in the explicit mentionings of the referent. In the following, I will refer to the *antecedent* of a pronoun, defined as the last mentioning of the referent, and to the *scope* of a referent, defined as the linear distance between the antecedent and the pronoun.

The task of anaphora resolution can be described as consisting of two subtasks: (i) to determine a set of possible antecedents, and (ii) to select the most likely antecedent from the candidate set. In determining the set of possible antecedents, it is important to know whether there are any limitations as to how far back in the text the antecedent of a pronoun may occur, where in the preceding text the antecedent is most likely to be found, and to which degree these constraints and preferences are influenced by other factors, such as the animacy of the referent.

The first part of the study thus concerns the scope of referents, both in general and in relation to various other properties of the referent. The linear distance between the pronoun and its antecedent was measured in terms of (graphic) sentences, with the exception of pronouns with antecedents within the same sentence, for which also a distinction was made according to whether the antecedent occurred in the same clause or not. In order to give the reader an impression of the differences between the three text types, separate accounts of the distributions within each subsample is given in Fig. 1a, and then a summary is given in Fig. 1b.

3.1. *Results*

The distributions of pronouns according to the linear distance of the antecedent differ considerably between the three text types. The fact that the proportion of pronouns with an antecedent in the same vs. the preceding sentence (A + B vs. C in Fig. 1a) is lower in the stories than in the other two text types, seem to

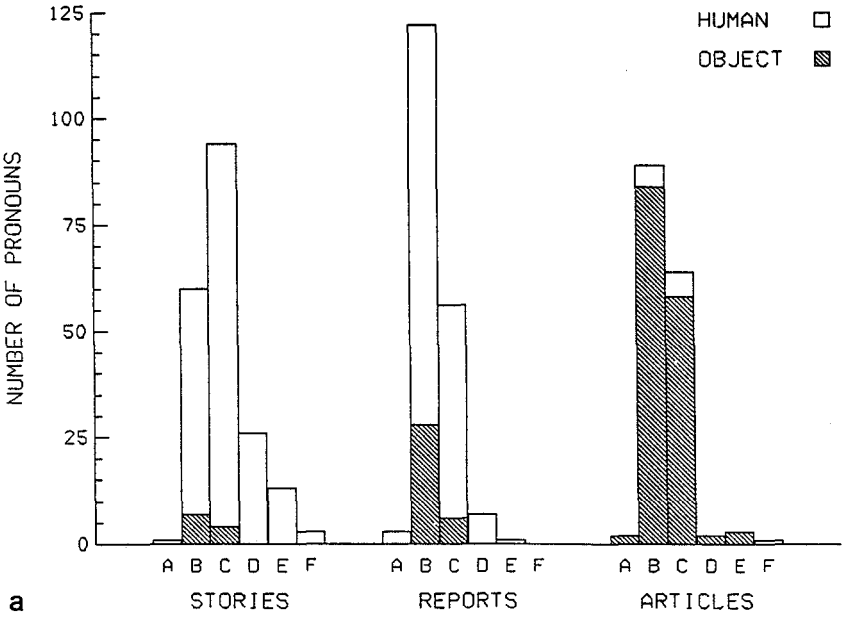

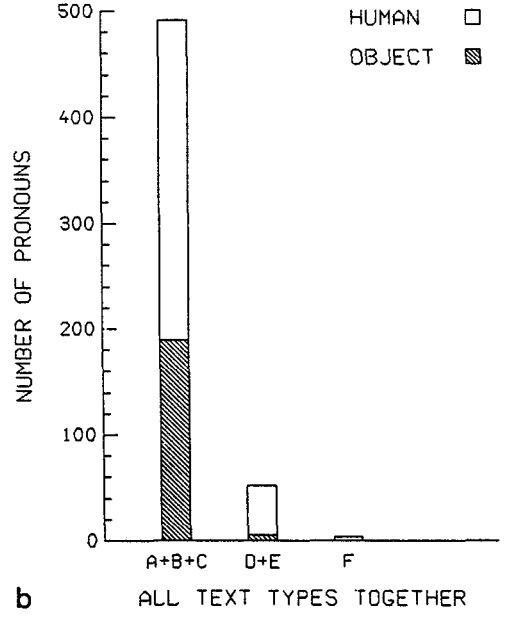

Fig. 1. Linear distance. Distributions of pronouns with human and object referents according to the linear distance between the pronoun and its antecedent in three text samples, where: the antecedent is in (A) the same simplex clause, (B) the same (graphic) sentence, (C) the preceding sentence, (D) the 2nd preceding sentence, (E) a sentence beyond the 2nd preceding sentence, and (F) there is no (explicit) antecedent.

depend largely on the differences in sentence length and complexity mentioned above. Thus, this can be regarded mainly as an effect of the crude measure of distance (i.e. in graphic sentences) used in the study, which makes it less theoretically interesting. The differences between the texts regarding the proportion of pronouns with an antecedent in a sentence beyond the preceding sentence (D + E in Fig. 1a) is more interesting, and can be attributed to their different proportions of human vs. object referents, since there is a strong tendency for human referents to have a wider scope, to which I will return shortly.

If we disregard the differences related to sentence complexity and animacy of the referent, the following general pattern emerges. For the large majority of pronouns (about 90% of the total sample), the antecedent is found within the same or the immediately preceding sentence. The number of pronouns with an antecedent beyond that distance rapidly decreases with each intervening sentence, suggesting a qualitative distinction between what I will call *short* and *long scope* referents/pronouns (A + B + C and D + E, respectively; cf. Fig. 1b). This pattern is even more pronounced among the inanimates.

The tendency for human referents to have a wider scope is especially evident in the stories, where they can have a scope over several sentences and even paragraphs. In the total sample, the maximal linear distance between a pronoun and its antecedent is 15 sentences. Three factors seem to contribute to such a long scope potential. In this case, the referent is *human* and one of the two main actors in the story, i.e. a *primary referent*, and the stretch of text between the pronoun and its antecedent contain a paragraph of direct speech, i.e. an *embedded discourse structure*. (Where there is no such embedding, the maximal distance found is 6 sentences.)

As a matter of fact, a primary referent may be referred to by means of a pronoun even in situations where another — but secondary — referent of the same gender has more recently been mentioned. And, conversely, a referent with a secondary role in the discourse is often referred to by a full lexical NP also in situations where it is the most recently mentioned referent. (Examples of this will be shown below in connection with the evaluation of the simple algorithm.)

The use of pronouns to refer to object referents is highly restricted. Only 5 of 194 inanimate pronouns have a long scope (as defined above), all five being well-established primary referents. Even primary object referents, such as, for example, the industrial robot IRB 1000 in one of the articles, are, however, regularly referred to by a lexical definite noun phrase instead of a pronoun, also in situations where the gender of the pronoun would have been a sufficiently disambiguating factor.

What is the explanation of the scope differences between human and object referents seen in the current sample? The quantitative differences could be claimed to correlate with differences in ambiguity and discourse role, i.e. the fact that there are more (competing) inanimate NPs than animate ones, and the fact that objects less frequently are primary referents. However, we would still have to account for the cases where primary object referents are referred to by a lexical definite NP even though there are no competing candidate antecedents. A low frequency of pronouns is often mentioned as a general characteristic of formal language. Thus, one possible explanation might be the formal language in the articles, where these examples are found. These texts do not contain any human primary referents for making a comparison, but the reports (partly written in a highly formal language) do, and there we find no corresponding cases of pronoun avoidance for human primary referents. The conclusion must be that the animacy of the referent is an important independent factor influencing the scope of referents. However, the question of how, exactly, it interacts with discourse role and other factors, cannot be fully answered without further studies of different types of natural text.

Within the whole sample, there are four pronouns lacking an explicit antecedent, all four being the plural 'they', which, I believe, is, in general, the most common form of pronoun used 'implicitly' in written discourse. In cases such as this, it is often hard to say whether the referent is specific but implicit, in which case the reference would have to be resolved by inference, or unspecific, in parallel to some implicit agents of passive verbs, in which case the referent can be marked as unspecified.

4. A SIMPLE PRONOUN RESOLUTION ALGORITHM

4.1. *The Algorithm*

In accordance with the cost-benefit approach described above, the following demands were put on the first version of the simple pronoun resolution algorithm. The algorithm should involve as few and simple principles as possible, and the requirements on the lexicon and on the input text should be kept down to a minimum.

The simple pronoun resolution algorithm

(1) Consider as a candidate antecedent every NP in the text that:
  (i) precedes the pronoun in the text
  (ii) agrees with the pronoun in number, animacy and gender.
  (iii) is not (co-referential with) the subject of the clause in which the pronoun occurs
  (iv) is not in a clause where the pronoun is subject

(2) Select as antecedent the NP in the set of candidate antecedents that:
  (i) is the most recent candidate in the text, i.e. the NP whose head is closer to the pronoun than the head(s) of the other NP(s),
  unless:
  (ii) there is another candidate, which is the subject of the same clause as the most recent candidate, and the pronoun is not a semi-demonstrative; in that case select the subject NP

In determining the set of candidate antecedents (1), the algorithm: (i) presupposes anaphora, rather than cataphora, (ii) checks for agreement, (iii) excludes the subject of the same clause as the pronoun, which cannot be the antecedent of a non-reflexive pronoun, and (iv) excludes all NPs in other constituents of the clause when the pronoun is the subject, in accordance with one of the consequences of the syntactic c-command constraint described in, for example, Reinhart (1983). In the selection of antecedents among the candidates (2), the basic principle is that of recency (i), sometimes overruled by the preference for subjects, which, however, does not apply to the semi-demonstratives (ii). The algorithm presupposes a lexicon with information regarding

the animacy and gender of nouns and pronouns, and an input text where NPs, subject NPs, clauses and sentences are identified.

The decision to disregard all other syntactic distinctions than that between subjects and non-subjects was based on considerations of both the cost, in terms of requirements on the input text, and benefit, in terms how much it was assumed to contribute to the efficiency of the algorithm. As shown in Källgren (1988), an almost perfect (99.5% correct) automatic identification of subject NPs in unrestricted Swedish text is possible with a very simple parser, while other syntactic categories are more difficult to identify correctly. In addition to the well-known constraints on intrasentential anaphora, there are certain preferences involved in (close) intersentential anaphora that relate to the distinction between subjects and non-subjects. When a pronoun (in the main form) has more than one possible antecedent in a preceding sentence (all other factors being equal), there appears to be a bias towards interpreting the pronoun as co-referent with the subject NP. The converse is true for semi-demonstrative pronouns, as suggested by the normative rule for the use of main vs. semi-demonstrative pronouns mentioned above.

The simple pronoun resolution algorithm was, in this first version, only designed for handling singular pronouns. Since the antecedent of a plural pronoun may be either a plural NP or two or more singular NPs in the preceding text, the resolution of plural pronouns would require additional procedures. The total number of plural pronouns was 92 (Table 2). Thus, the simple pronoun resolution algorithm was tested on a remaining sample of 457 pronouns (Table 3).

Table 2. Plural pronouns. Distributions of singular and plural pronouns with human and object referents in three text samples.

	Singular		Plural		Total	
	N	%	N	%	N	%
Stories	163	(82.7)	34	(17.3)	197	(100.0)
Reports	178	(93.2)	13	(6.8)	191	(100.0)
Articles	116	(72.0)	45	(28.0)	161	(100.0)
Total	457	(83.2)	92	(16.8)	549	(100.0)

Table 3. Results of the simple pronoun resolution algorithm. Distribution of correct and incorrect antecedent selection for singular pronouns with human and object referents in three text samples.

	Correct		Incorrect		Total	
	N	%	N	%	N	%
Stories	162	(99.3)	1	(0.7)	163	(100.0)
Reports	166	(93.3)	12	(6.7)	178	(100.0)
Articles	87	(75.0)	29	(25.0)	116	(100.0)
Total	415	(90.8)	42	(9.2)	457	(100.0)

4.2. *Results*

Evaluated on the total sample, the algorithm worked surprisingly well, providing a correct antecedent assignment for almost 91% of the pronouns. Its performance differed considerably, however, in the three text types. In the stories there was only one case of incorrect antecedent assignment, viz. the NP *God*, selected according to the principle of recency (the correct referent being the main actor of the story, mentioned in a preceding sentence). But the reports, and even more so the articles, turned out to put much harder demands on a pronoun resolution algorithm.

A possible objection against the results as presented in Table 3 is that the figures say nothing about the degree to which the algorithm actually manages ambiguity, since it might often be the case that there is only one possible candidate. There is, however, no simple solution to this evaluation problem. As long as we have not been able to formulate further restrictions on the set of possible antecedents, ambiguity would, in principle, be said to exist whenever there is one more NP anywhere in the preceding text that agrees with the pronoun.

A partial answer to the question could, however, be provided by the pronouns with an object referent. As we saw above, the scope of such discourse referents is highly restricted, 97.4% of the antecedents being in the same or the preceding sentence of the pronoun. Let us, then, consider an inanimate pronoun to be ambiguous when there is a competing NP in the same or the preceding sentence, and re-evaluate the algorithm in relation to

Table 4. Results of the simple pronoun resolution algorithm for certain pronouns. Distribution of correct antecedent selection for singular pronouns with object referents relative to the total sample and to the sample of pronouns for which there is a competing antecedent candidate in the same or preceding sentence in three text samples.

					Other candidate antecedent in same or preceding sentence			
	Correct		Total		Correct		Total	
	N	%	N	%	N	%	N	%
Stories	5	(100.0)	5	(100.0)	2	(100.0)	2	(100.0)
Reports	31	(91.2)	34	(100.0)	23	(88.5)	26	(100.0)
Articles	82	(74.5)	110	(100.0)	53	(65.4)	81	(100.0)
Total	119	(79.2)	149	(100.0)	78	(71.5)	109	(100.0)

this subsample of pronouns (Table 4). For all pronouns with an object referent, the correctness was about 10% lower than that for the total sample. When the cases where no other antecedent was considered to be present are excluded, the correctness ratio decreases to 71.5%.

These figures can be compared with the results of a pronoun resolution algorithm presented in Hobbs (1978), which is the only algorithm for which I have been able to find such statistics. Hobbs' pronoun resolution algorithm is syntactically more complex than mine, involving a search of a surface parse tree and taking into account differences in embedding and syntactic function. The algorithm works for 192 of 217 (88.5%) singular pronouns and, among these, for 55 of 71 (77.5%) of the inanimates. In the cases where, in Hobbs' words, there are "more than one plausible antecedent" (no definition is given), the algorithm selects the correct antecedent for 33 of 48 (68.8%) occurrences of the pronoun *it*. Besides the possible differences between Swedish and English, the different results for the three text types in my study should warn us not to make too much out of a comparison between two algorithms tested on two different samples. Instead, the discussion should be based on a quantitative evaluation of the contribution of each particular constraint or preference rule within an algorithm tested on the same sample.

4.3. *Some Causes of Incorrect Antecedent Selection*

For human referents, there seems to be one main factor that clearly overrules the principle of recency, viz. the status, or discourse role, of the referent. In 11 of the 12 incorrect antecedent selections in the juridical reports, the correct antecedent is a primary discourse referent. In this particular text type, the primary referent is always the person whose statement is being reported, or, among inanimates, the object of central interest (i.e. the stolen article). An example of this is seen in the following passages from two persons' statements on the same event. (In these and following examples, the pronoun and (the head of) the antecedent are marked by indices, and (the head of) the incorrectly selected antecedent is marked with an asterisk.):

(3) Magnus K har hörd över åtalet uppgivit i huvudsak: . . . (5 sentences) . . . Han viftade med handen framför Knut Js ansikte och uppmanade denne att lämna ifrån sig sin plånbok. Möjligen yttrade *han*$_i$ också något om att kalla på polis. *Knut J*$_*$ lämnade ifrån sig sin plånbok. *Han*$_i$ tog de pengar som fanns i plånboken cirka 100 kr samt två kontokort och lämnade därefter tillbaka plånboken till Knut J, vilken därefter fick lämna lägenheten.

Magnus K's statement on the indictment can be summarized as follows: . . . (5 sentences) . . . He waved his hand in Knut J's face and demanded that he (semi-demonstrative) turn over his wallet. *He*$_i$ possibly also stated something about calling the police. *Knut J*$_*$ turned over his wallet. *He*$_i$ took the money which was in the wallet, approximately 100 kronor, as well as two credit cards and then returned the wallet to Knut J, who was then allowed to leave the apartment.

(4) Knut J har uppgivit i huvudsak: . . . (4 sentences) . . . Han uppfattade situationen som hotfull. När Magnus K sade åt honom att lämna över sin plånbok gjorde *han*$_i$ det. *Magnus K*$_*$ tog de pengar omkring 100 kr som fanns i plåboken jämte två kontokort. *Han*$_i$ fick därefter tillbaka plånboken och fick lämna lägenheten.

Knut J's statement can be summarized as follows: . . . (4 sentences) . . . He experienced the situation as threatening.

When Magnus K demanded that he turn over his wallet he_i did so. *Magnus* K_* took the money about 100 kr which was in the wallet as well as two credit cards. He_i then received his wallet back and was allowed to leave the apartment.

These passages also illustrate a typical pattern, where the primary referent, with the exception of the introductory NP, is regularly referred to by a pronoun, while the secondary referent, with the exception of one semi-demonstrative pronoun, is referred to by proper name.

In the case of object referents, discourse status also appeared to play a role, interacting, however, with a number of other factors. The articles turned out to be good material for testing the pronoun resolution algorithm, offering many examples of ambiguous situations that had not shown up in the other texts. The sources of incorrect antecedent assignment were also more diverse, including some cases where even human readers may disagree about which of two or more candidate NPs would be the intended antecedent. Accordingly, the number of errors was considerably higher than in the other text types, viz. 29 (of 116). A detailed analysis of this error sample will not be presented here, partly due to lack of space and partly because the sample is too small for judgements of the quantitative effects of all possible elaborations of the current formulation of the algorithm. Let me, thus, confine myself to a few thoughts about conceivable additional constraints and preference rules.

The formulation of the correct preferential principles actually calls for a considerably larger sample than the one analysed here. The principle of recency is obviously too rough, although I still believe it correctly reflects the default case. The principle of preference for subject NP's performs well in most cases, but will of course fail when a new discourse referent is introduced and becomes the topic of the following sentence (ex. 5), and also sometimes intrasententially (ex. 6):

(5) *Lastslussen*$_*$ är utrustad med en eldriven *vädertätningsdel*$_i$. Den$_i$ består av s k dragspelstätningar som ansluter mot lastbilens kapell.

The loading *lock*$_*$ is equipped with an electrically driven weather caulking *section*$_i$. *It*$_i$ consists of so-called accordion caulkings which connect with the vehicle's cover.

(6) När *datorn*$_*$ anropar *bilen*$_i$ svarar *den*$_i$ automatiskt med status och inställt område.

When the *computer*$_*$ calls up the *car*$_i$ *it*$_i$ responds automatically with the status and previously tuned in area.

It is often suggested that pronoun interpretation involves preferences related to the degree of embedding of NPs. So does, for example, the set of pronoun interpretation rules proposed by Guentner & Lehman (1983), which include principles by which NPs in a matrix clause or phrase are preferred over NPs in embedded clauses or phrases, head nouns are preferred over noun complements and accusative object NPs are preferred over other non-subject NPs. Interestingly enough, these assumptions were not confirmed in the present sample of Swedish texts. Although the addition of such preference rules might seem to be motivated in some of the cases of incorrect antecedent assignment, it would not improve the overall outcome of the algorithm, since there are so many exceptions to the rules.

The problem with these rules is that they are too general. An interesting case in point is the genitive complement, which appears to be preferred over the head (or whole) NP (see also example 8 below):

(7) Efter det att *ringens*$_i$ felaktiga *form*$_*$ blivit korrigerad i backarna, flyttas *den*$_i$ tillbaka till mäthuvudet för inspektion.

Following the correction of the *ring's*$_i$ deficient *construction*$_*$ in the jaws, *it*$_i$ is replaced at the measuring head for inspection.

This is in accordance with the observation that the genitive construction often signals a hierarchical relation between two referents (cf. Brodda 1975 and Fraurud 1986, 1988).

Although the difference between primary and secondary referents has proved to play an important role, I am not prepared to formulate a preference rule involving this feature. So far, all we know is that, in general, primary referents are preferred over secondary referents, and also that a secondary referent *may* be accessible to pronoun reference within a limited scope. Furthermore, we have yet no formal definition of a primary referent, and it may be expected that the discourse role is a much more gradient phenomenon than implied in the discussion above. The

definition problem is highly related to the problem of identifying primary referents at a particular point of a particular text. In one of the text types, the reports, both the primary referent and its scope could easily be identified, but in most texts this presents a serious problem (concerning other possible ways in which the discourse role of a referent is expressed in the text, see Hellberg (1984) and Fraurud (1986)).

The addition of a principle of preference for concrete over abstract referents would slightly improve the algorithm's selection between inanimates (ex. 7 above). Some hypotheses to be tested on a larger sample can be summarized in the following hierarchy for pronoun accessibility:

human < animate < concrete < abstract

The following two additional constraints would reduce the number of errors in the current sample by some percent.

(i) Coordinate NP constraint: Exclude from the set of candidate antecedents NPs that are coordinated with other NPs, unless the referents of the coordinate NPs are of different sex.

This constraint applies to examples like (8), although, in this case, it is only one of several factors:

(8) *Pumpens*$_i$ lilla format och låga *vikt*$_*$ gör att *den*$_i$ kan placeras närmare sugstället än konventionella vakuum-pumpar.

The *pump's*$_i$ compact size and low *weight*$_*$ enables *it*$_i$ to (litt.: makes that it can) be placed closer to the suction area than (is possible with) conventional vacuum pumps.

In this example, *pumpen* is n-gender, *format* is t-gender and *vikt* is n-gender. Notice that (grammatical) gender contrast, as opposed to the semantic contrast between masculine/feminine, is not sufficient to allow for pronoun reference into a coordinate NP (cf. Tingbjörn 1979). This should be regarded as a special case of a more general phenomenon, as exemplified by the following (constructed) sentences (where: *lejon* is t-gender and *hjort* is n-gender).

(9) *Pojken*$_i$ jagade *flickan*$_j$, men *hon*$_j$ sprang snabbare än *han*$_i$.

The boy$_i$ chased *the girl*$_j$, but *she*$_j$ ran faster than *he*$_i$ (did)

(10) ?*Lejonet*$_i$ jagade *hjorten*$_j$, men *den*$_j$ sprang snabbare än *det*$_i$.

The lion$_i$ chased *the deer*$_j$, but *it*$_j$ ran faster than *it*$_i$ (did).

When asked about the acceptability of sentence (10) and similar constructions, native Swedish speakers find them 'odd' and would prefer some paraphrase where the pronoun reference is avoided. Again, we see an example of the asymmetry between animate and inanimate pronouns, which, in part, may have to do with the fact that the gender contrast in inanimate pronouns is devoid of semantic content.

The second constraint that can be formulated is the following example of an intrasentential constraint involving both syntactic and lexical information:

(ii) Constraint on subjects of certain verbs with *att*-clause complement: If the pronoun is the subject of an *att*-clause being the complement of one of the verbs *göra* 'make/cause', *innebära*, 'imply' or *medföra* 'bring about', then exclude from the set of candidate antecedents the subject NP of the matrix clause.

(Note that these verbs are partly the same as those mentioned in connection with the propositional pronouns.)

To my surprise, the algorithm proved to be too restrictive in one respect, viz. the agreement requirement. In three cases the algorithm failed because what would be considered the antecedent did not agree with the pronoun in gender or number. In the following example, the incorrectly selected (t-gender) antecedent is the semantically salient attribute NP in what, syntactically, would count as the correct antecedent, viz. the whole NP (in n-gender) (although this might be discussed):

(11) NYTT VÅGSYSTEM

ASEA har utvecklat en ny *generation*$_i$ av Pressductor *vågsystem*$_*$. *Det*$_i$ är direkt anpassat för ASEA Master datorbaserade styrsystem, men kan även anslutas till andra datorer.

NEW WAVE SYSTEM

ASEA has developed a new *generation*$_i$ of Pressductor wave *system*$_*$. *It*$_i$ is directly designed for use with the ASEA

> Master computer-based steering system, but can also be connected with other computers.

However, the major shortcomings of the current algorithm as regards the determination of the set of possible antecedents is that it allows for too many possible antecedents, including those that either should not count as proper introductions of a discourse referent (see Fraurud 1986), or which no longer are accessible for pronoun reference (such as, for example, the short scope inanimates discussed above). This property of the algorithm did not produce errors, but makes it theoretically less adequate and more difficult to implement.

An even more serious practical problem is that the algorithm gives no indication of *when* it fails, which, in turn, can lead to a domino-effect, where subsequent pronouns are misinterpreted due to the first failure. This problem can probably not be completely solved, but it is possible to conceive of a procedure that measures the degree of ambiguity involved in the selection of an antecedent and, thereby, the certainty of the antecedent assignment.

5. CONCLUSION

In this paper, I have presented preliminary empirical data for an analysis of the relative influence of different mechanisms involved in pronoun interpretation. Like many other linguistic phenomena, these mechanisms are often inherently non-categorical. The analysis of quantitative data may contribute to our understanding of such phenomena by revealing default values and preferences and by suggesting qualitative distinctions in situations where, at a particular point of the scale of a gradient phenomena, quantity turns into quality. In the present study, one case in point was the distinction between long and short scope (here roughly described in terms of graphic sentences), derived from the observation that non-human referents (with a few exceptions) and secondary human referents were accessible to pronoun reference only when the antecedent appeared in the same or preceding sentence.

The cost-benefit approach to anaphora resolution outlined here seems to have a reasonably good chance to prove its practical usefulness, considering the fairly good results of the very simple

algorithm and the suggested possibilities for improving its performance by integrating knowledge derived from error analysis and further quantitative studies. Obviously, 100% correctness is not a realistic goal for pronoun resolution in unrestricted natural text. However, also partial solutions, resolving most of the pronouns in a text, will considerably improve the practical applicability of NLP systems to, for example, information retrieval.

I want to conclude with a methodological remark of relevance to further work along the lines sketched out in this paper. A drawback in the present study was that I had to use a manual method for testing the algorithm. In order to remove possible hidden implicitness, which often occurs in algorithms that are only verbally formulated, the algorithm was implemented in the form of a computer program written in Lisp. The program, however, was only tested on a short constructed text, since currently available corpora of unrestricted Swedish text are either too limited in size and selection of text, or do not meet the (relatively low) input demands of the algorithm. This is an important obstacle in quantitatively oriented research on discourse, since manual work is slow and inefficient, limiting the amount of text that can be analysed within a reasonable amount of time, and decreasing the reliability of the results. In the light of some promising results of a parser for unrestricted Swedish text reported in Källgren (1984, 1988), I see, however, hope for better future possibilities for quantitative analysis of referential relations in natural text as well as other discourse phenomena.

ACKNOWLEDGEMENT

I wish to thank Östen Dahl, Gunnel Källgren and Christina Hellman for their comments on a draft of this paper.

NOTES

1 The work was carried out within the project Discourse Semantics and Discourse Strategies, financed by the Swedish Council for the Humanities and Social Sciences.

2 From Lars Ahlin, *Inga ögon väntar mig*, Tidens förlag, Stockholm, 1948.

3 From a machine readable corpus of reports of the procedures of the Supreme Court 1981.

4 From Framsteg inom forskning och teknik, IVA, Ingenjörsvetenskapsakademins Meddelande 240/1982, 243/1983, 247/1984, Stockholm.

REFERENCES

Brodda, B. 1975. Har genitiven någon grundbetydelse? In *Nionde sammankomsten för svenskans beskrivning*, Göteborg, 18–19 October 1975.

Fraurud, K. 1986. The Introduction and Maintenance of Discourse Referents. In Dahl (ed.), *Papers from the Ninth Scandinavian Conference of Linguistics*, Stockholm, January 9–11, 1986, pp. 111–122.

Fraurud, K. 1988. Hierarchical Relations between Discourse Referents and the Interpretation of Antecedentless Definite Noun Phrases. In Dyvik (ed.), *Papers from the Tenth Scandinavian Conference of Linguistics*, Bergen, June 11–13, 1987 (forthcoming).

Guentner, F. & Lehman, H. 1983. Rules for Pronominalization. In *COLING-83 Proceedings from the 9th International Conference on Computational Linguistics*, pp. 144–151.

Grosz, B. J. 1977. The Representation and Use of Focus in a System for Understanding Dialogues. In *Proceedings from the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, Mass., pp. 67–76, Los Altos: William Kaufmann, Inc., 1977. Reprinted in Grosz, B. J., Sparck Jones, K. & Webber, B. L. (eds.) (1986), pp. 353–363.

Grosz, B. J., Sparck Jones, K. & Webber, B. L. (eds.) 1986. *Readings in Natural Language Understanding*. Los Altos: Morgan Kaufmann Publishers, Inc.

Hakulinen, A. & Karlsson, F. 1980. Finnish Syntax in Text: Methodology and Some Results of a Quantitative Study. *Nordic Journal of Linguistics* 3, 93–129.

Hellberg, S. 1984. Satsens subjekt och textens. *Nysvenska studier* 64, 107–160.

Hobbs, J. 1978. Resolving Pronoun References. In *Lingua* 44 (1978). Reprinted in Grosz, B. J., Sparck Jones, K. & Webber, B. L. (eds.) (1986), pp. 339–353.

Källgren, G. 1984. *Automatisk excerpering av substantiv ur löpande text — Ett möjligt hjälpmedel vid datoriserad indexering?* Report 1984:1. IRI, University of Stockholm.

Källgren, G. 1988. What Good is Syntactic Information in the Lexicon of a Syntactic Parser? In *De Nordiska Datalingvistikdagarna 1987* (forthcoming).

Reichman, R. 1981. Plain Speaking: *A Theory and Grammar of Spontaneous Discourse*. Technical Report No. 4681, Bolt, Beranek and Newman, Inc.

Reinhart, T. 1983. *Anaphora and Semantic Interpretation*. London: Croom Helm.

Sidner, C. L. 1979. Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse. Ph.D. thesis. Dept of Electrical Engineering and Computer Science. Mass.: MIT.

Tingbjörn, G. 1979. Kontrastiv minigrammatik. In Hyltenstam, K. (ed.), *Svenska i invandrarperspektiv*. Lund: Liber.

Webber, B. L. 1978. *A Formal Approach to Discourse Anaphora*. New York: Garland Publ.

Wellander, E. 1973. *Riktig Svenska*. Stockholm: Esselte Studium.

Winograd, T. 1973. A Procedural Model of Language Understanding. In Schank, R. & Colby, K. (eds.) 1978. *Computer Models of Thought and Language*. New York: W. H. Freeman. Reprinted in Grosz, B. J., Sparck Jones, K. & Webber, B. L. (eds.) (1986).