Anaphora and its Resolution

1. Anaphora

A pervasive property of natural language is the use of **referring expressions**, or **anaphors** These are expressions that denote entities previously introduced into a discourse, and include definite noun phrases and pronouns, e.g.

I saw a car with wide-wall tyres. The car/It was a Ferrari.

A natural language understanding system must be able to determine which discourse entity a referring expression refers to. This task is hard because it involves a variety of types of knowledge, viz. syntactic, semantic and discourse.

In many cases, the problem of determining the referent of a referring expression reduces to finding its antecedent, i.e. some previous portion of the text that introduces the same entity. In a language which has a gender system, it is generally the gender of the antecedent that determines the gender of the anaphor, rather than the semantic nature of the entity denoted by the antecedent. So, for instance, the ambiguity of the pronouns in a text like:

I took the sugar off the shelf and put it into my basket. When I got to the counter I paid for it.

will not arise in German, because 'sugar' will be the only one of the possible antecedents with the correct gender. However, this superficial strategy will not always be appropriate. For instance, in German again, where grammatical gender and natural gender diverge, one gets the following effect.

Das Maedchen .... Es .... Sie ....

and 'Sie' a feminine one. Immediately following the full noun phrase, grammatical gender determines the pronoun to be used, but as the textual distance between the two increases, it becomes more likely that natural gender will be used.

Other examples where pronoun resolution does not reduce to finding a textual antecedent may be illustrated by the following:

My uncle doesn't have a spouse, but my aunt does, and he is lying on the floor.

Here there is no antecedent, since it is part of an understood verb phrase that has undergone ellipsis.

My neighbours have a monster Harley 1200. They are really huge bikes.

Here the pronoun is plural but the only plural noun phrase in the preceding discourse is not the antecedent.

When the anaphor is a definite noun phrase, there is an even greater possibility that there will be no textual antecedent at all, viz:

A bus came round the corner. The driver was wearing a yellow tie.

Note also that referring expressions may refer forward, a phenomenon called cataphora.

For instance, contrast the possibility of understanding the pronoun in these sentences as referring to (the same as) Tweetie Pie:

Creeping up on him, Tweetie Pie saw Sylvester.
Creeping up on Tweetie Pie, he saw Sylvester.

2. Anaphora and Syntax

The obvious value of syntax to the question of anaphor resolution is that it rules out certain otherwise plausible candidates for antecedents. As well as the examples of cataphora given above, consider the following:

I saw a driver in his truck
I saw a driver of his truck

In the model of phrase structure called X-bar syntax that we saw in HO4, the essential difference between these can be captured as follows:

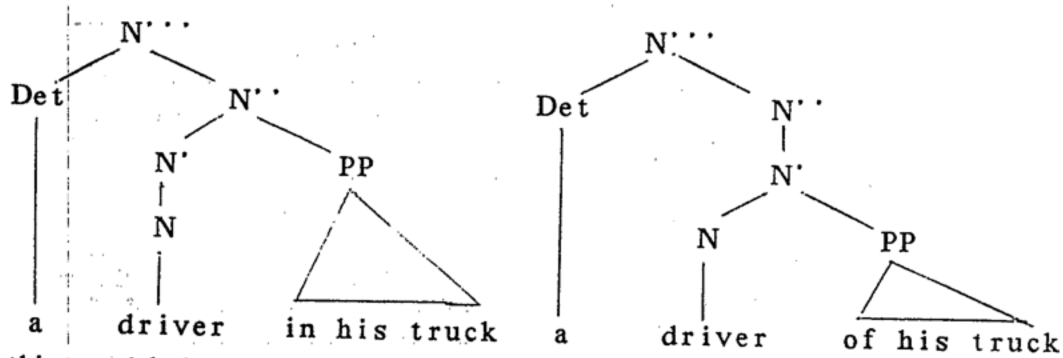

Hobbs uses this model in defining his naive algorithm for determining pronoun antecedents. Naive means that syntax and selectional restrictions only are used. The algorithm requires that the path between a pronoun and the NP (N''') node that dominates the potential antecedent should not pass through the N' that dominates it. This is just a way of saying that a noun cannot be coreferential with a pronoun in one of its subcategorised complements, which are the daughters of N' (as opposed to N'').

Hobbs algorithm essentially defines three positions with respect to a pronoun where an antecedent might be find, in what seems to be decreasing order of likelihood. The positions are defined with respect to nodes of two syntactic categories, NP and S, which I will call X-nodes.

1. Under the first X-node that dominates the pronoun (not including the NP that the pronoun is). Here antecedents may be found only to the left, so

*(1) * He_i saw $John_i$*
*(2) * He_i saw $John_i$'s friend*

* - non gramatical sentence
i= same entity.

and only under another X-node:

*(3) * $John_i$ frightens him_i*
(4) $John_i$'s poodle frigtens him_i

(5) that John$_i$ might have aids frightens him$_i$

It would seem that the reason why it fails in (3) is that this is precisely when we would use a reflexive pronoun to indicate coreference.

2. Under a higher X-node, including the X-node itself. The latter possibility is that illustrated above with *driver in his truck*, and is subject to this N' constraint.

Below and to the left of this node, Hobbs algorithm allows any NP to be an antecedent. Actually, this is a little generous, as he points out. But the problem is distinguishing cases like (6) and (7):

*(6) * John$_i$ likes the picture of him$_i$* ← - dubious. (correct?!)
(7) John$_i$ likes the picture of him$_i$ that hangs in the Tate

There seems to be a similar possibility of using a reflexive as in position 1. But this gets weaker as one goes to higher nodes above the pronoun, as in (8)

(8) John$_i$ believes that the picture of him$_i$ is a good likeness

Below and to the right of these higher X-nodes, an 'antecedent' may be found only above any other X-node. Contrast (9) and (10):

(9) When he$_i$ got home, John$_i$ had tea
*(10) * When he$_i$ got home, John$_i$'s brother had tea*

3. In a previous sentence. Here any noun phrase may be taken as antecedent.

In all those cases where the algorithm says that any node in a sub-tree may be an antecedent, the tree is traversed in left-to-right breadth first order. This will mean, for instance, that subjects will be proposed before objects, and these before more deeply embedded NPs.

Hobbs claims that once simple selectional restrictions are incorporated, this algorithm finds the correct antecedent in over 90% of cases. Selectional restrictions are just statements of the semantic type restrictions that a predicate imposes on its argument. If the pronoun occurs in a predication, we will be able to say what type of entity it must be, and this can eliminate possible antecedents that are not of this type. Hobbs exemplifies with the sentence:

The castle in Camelot remained the residence of the king until 536 when he moved it to London

Here, several of the antecedents proposed by the naive algorithm are eliminated from consideration by the selectional constraints:

Dates can't move
Places can't move

Large fixed objects can't move

Of course, such selectional constraints can never help with pronouns like 'he' and 'him'.

The algorithm takes no account of the possibility of sentential antecedents. If it did, it would do a lot worse than it does. But of course, it gets things wrong, like:

The newspaper reported that Lawson had claimed the economy was improving, bit I didn't believe it.

Even if it did include S nodes as possible antecedents, it would still get this wrong, since it would propose possibilities in the order:

the newspaper
Lawson had claimed the economy was improving
the econonmy was improving

which seems the inverse of the intuitive order.

Even if a pronoun-resolving algorithm was made considerably more sophisticated than this one, by incorporation of large amounts of real-world knowledge and a sophisticated model of the ongoing discourse, it would still need to incorporate the basic syntactic constraints on coreference that Hobbs has tried todefine in his algorithm. Nevertheless, these syntactic constraints are obviously insufficient, so Hobbs also proposed a resolving mechanism that used real-world knowledge for resolution.

3. Anaphora and Semantics

Hobbs supposed that the sort of information that a lexicon for a sophisticated natural language understanding system could be expected to contain would be adequate to resolve even those problems that had always been claimed to require arbitrary amounts of real-world knowledge. The classic example is Winograd's:

"*The councillors refused the woman a permit to demonstrate because they advocated violence*"
"*The councillors refused the woman a permit to demonstrate because they feared violence*"

The pronoun 'they' would have to be resolved in an MT system from English to French, say, since it would presumably translate as 'elles' in the first sentence and 'ils' in the second. Actually, Hobbs uses a sex-neutral version of the sentences, which is more idealogically acceptable, but doesn't illustrate the point about translation.

The basic idea of the semantic resolution approach is to use axioms associated with word meanings to infer further properties of entities introduced in the sentence. Then a sort of default inference rule called 'knitting' applies. This says that if at some point we have inferred $p(X_1)$ where X_1 is the variable corresponding to the pronoun E_p, and we have inferred $p(X_2)$ for a variable X_2 introduced by some other expression E_l, then in the light of any evidence to the contrary, we assume $X_1 = X_2$, that is, E_l is the antecedent of E_p.

What determines the inferences that we make from the axiomatic lexicon. There are two major guiding principles. One, we try to establish relations between sentences as being one of a small set. Two, we try to interpret general (i.e. ambiguous or vague) predicates according to their immediate context.

3.1. Relations between sentences

Common relations that obtain between sentences include Contrast, Violated Expectation, and Cause. Sentence patterns like $S_1 but S_2$, or $S_1. However, S_2$ tend to indicate that S_1 and S_2 stand in a Contrast relation, while $Since S_1, S_2$ and $S_1. Therefore S_2$ imply a Cause relation. We try to draw inferences that will confirm these patterns. For example, if we think two sentences stand in a Contrast relation we try to infer from them two propositions P_1 and P_2 such that:

(a) the predicates of P_1 and P_2 are contradictory or lie at opposite ends of a scale;

(b) one pair of corresponding arguments of P_1 and P_2 are identical;

(c) the other pair of corresponding arguments are similar but different

If we think two sentences S_1 and S_2 stand in a Cause relation, then we try to infer some P_1 from the first and some P_2 from the second and to establish some causal chain of inferences from P_1 to P_2.

3.2. Predicate interpretation

Hobbs subsumes under this header such inferences as the following:

(a) from into(Y,X), if Y is a motion and X a region, infer that the endpoint of Y is inside X

(b) infer the part from the whole, as in

He landed on (the roof of) the building

(c) infer missing quantity words, as in

(The price of) Coffee is higher this month

Hobbs claims that one of the above two inference generating mechanisms in combination with knitting is adequate for resolving most pronouns.