

Chapter 6

CONSTRAINTS AND DEFAULTS IN ANAPHOR RESOLUTION

There is, of course, no firm dividing line between the act of deciding what the candidates for an anaphor's antecedent are and the act of deciding among them; it all depends on how much information there is to limit the possibilities during the search. We can imagine at one extreme a two-pass system which computes when necessary, or always maintains, a focus as we have discussed above, and then chooses among them when necessary, and at the other extreme a one-pass system which applies both focus and anaphor-specific constraints to each entity when looking for a particular referent. Combination approaches are also possible. I know of no evidence favouring one of these approaches over the others on theoretical grounds, nor is it clear when each is the most computationally efficient.

So far in this thesis, I have tacitly assumed that in determining the candidates – the focus – we have no information about a particular anaphor occurrence, but are rather generating the maximal set of entities that some anaphor could refer to at the present point in the text. In this chapter, now, I consider the additional constraints imposed by having information on a particular anaphor that needs resolving, and the problem of default referents. It is unimportant to the present discussion at what point anaphor-specific information is used.

Many anaphor-specific factors have been discussed earlier in this thesis; in these cases, the reader is referred back to the appropriate sections.

6.1. Gender and number

While gender and number are strong constraints on reference, we saw in section 2.3.1 that they are not absolute: a plural anaphor can have a singular antecedent, a feminine one a masculine antecedent, and so forth.

6.2. Syntactic constraints

Linguists have discovered many syntactic constraints on anaphoric reference; see section 3.2.2.

6.3. Inference and world knowledge

In sections 2.4.2 and 3.2.6, we saw how world knowledge and inference may need to be applied.

6.4. Parallelism

Consider the following texts:

- (6-1) Ross likes his⁽¹⁾ beer and Daryel his⁽²⁾ carrot juice, but Bruce swears by his⁽³⁾ Samoa Fogcutter (two parts gin, one part red wine).
- (6-2) Roger makes some great drinks at home. Ross likes his⁽¹⁾ beer and Daryel his⁽²⁾ carrot juice, but Bruce swears by his⁽³⁾ Samoa Fogcutter.

In each *his* refers to the immediately preceding name, and in the additional context of (6-2), each refers to Roger. That each *his* is dealt with in the same way, in a certain sense, is the not uncommon linguistic phenomenon PARALLELISM. Parallelism can operate at both the syntactic and semantic levels. Its effects are quite strong: there is, I conjecture, NO context in which can be embedded such that the *hiss* aren't dealt with in a parallel manner (in which *his*⁽¹⁾ is someone in a previous sentence, *his*⁽²⁾ is Daryel, and *his*⁽³⁾ is Ross, for example).

Clearly, an anaphor resolver needs a knowledge of parallelism, although I am not aware of any attempt to formalize the phenomenon, let alone implement it. Note that parallelism is particularly important in resolving surface count anaphora (see section 2.3.2).

6.5. The preferred antecedent and plausibility

In section 2.6, when discussing the problems of ambiguous text, I introduced the notion of a PREFERRED or DEFAULT ANTECEDENT. The preferred antecedent rule says "if you cannot decide on a single 'right' antecedent for the reference, choose from the uneliminated candidates the one that has quality *X* in the greatest proportion; if no candidate has significantly more of quality *X* than the others, treat the sentence as genuinely ambiguous". In this section, I will look at the nature of quality *X*, and will start by immediately prejudicing the discussion by giving *X* the name PLAUSIBILITY.

Let us first recall two potentially ambiguous examples from section 2.6:

6.5 The preferred antecedent and plausibility

- (6-3) Daryel told Ross he⁽¹⁾ was the ugliest person he⁽²⁾ knew of.
 (6-4) The FBI's role is to ensure our country's freedom and be ever watchful of those who threaten it.

The default interpretation of (6-3) is that Daryel is insulting Ross ($he^{(1)} = Ross$, $he^{(2)} = Daryel$), rather than being self-critical ($he^{(1)} = he^{(2)} = Daryel$). This may be simply because insulting behaviour is more common than openly self-critical behaviour with respect to personal appearance in western English-speaking cultures. That is, an insult is the most plausible interpretation of (6-3), and the corresponding antecedents are chosen accordingly. Similarly, in (6-4), *it* is more plausibly *our country* or *our country's freedom* than *the FBI* or *the FBI's role*.

Moreover, Kirby (1977, 1979) has shown in psycholinguistic experiments that plausibility of meaning is a factor in the time taken to understand a structurally ambiguous sentence – ambiguous sentences lacking a single, obviously most plausible interpretation take longer. This suggests that plausibility could also be relevant to ambiguous anaphors.¹

Plausibility differs from other constraints mostly in its weakness. For example, the gender constraints that make (6-5) so bad:

- (6-5) *Sue found himself pregnant.

can be broken in certain cases (see 2.3.1), but in most contexts are very strong and not really a matter of degree. Plausibility, on the other hand, IS a matter of degree, and always requires evaluation relative to the other possibilities.

Is plausibility the only factor (other than theme, of course) in assigning the default antecedent? Or conversely, is there a well-formed anaphorically ambiguous text in which a preferred antecedent exists but is neither the theme nor the candidate that gives the text its most plausible reading? I have not been able to construct such an example, but am not willing to assert that none exist. If they do exist, they are probably rare enough for an NLU to ignore with reasonable impunity.

The computational problem of deciding how plausible an alternative is, is extremely difficult. While it relies on knowledge of real-world norms, inference plays a part too. For example, one is unlikely to find explicitly in a knowledge-base grounds on which (6-4) can be resolved, namely:

- (6-6) If *X* guards *Y*, then it makes more sense for *X* to keep under surveillance all who threaten *Y* rather than just those who threaten *X*.

Working out what "makes most sense" can involve an extremely complex and time-consuming process of generating and evaluating consequences.

However, there is at least one form in which plausibility becomes computationally simple, and we shall examine this in the next section.

¹It remains for someone to perform a properly controlled experiment to test this hypothesis. But see also the next section, on causal valence.

6.5 *The preferred antecedent and plausibility*

6.6. Implicit verb causality

One guise in which plausibility turns up is IMPLICIT VERB CAUSALITY or CAUSAL VALENCE. In a series of experiments (Garvey and Caramazza 1974; Garvey, Caramazza and Yates 1975; Caramazza, Grober, Garvey and Yates 1977), it was shown by Catherine Garvey and her colleagues that the causal valence of a verb can affect the antecedents assigned to nearby anaphors. For example, consider these texts (from Caramazza et al 1977):

(6-7) Muriel won the money from Helen because she was a skillful player.

(6-8) Ronald scolded Joe because he was annoying.

People tend to interpret *she* in (6-7) as *Muriel*, the first NP of the sentence, and *he* in (6-8) as *Joe*, the second NP. In general, with sentences of the form:

(6-9) *NP1 VERB*ed *NP2* because {he | she} . . .

(where both *NP1* and *NP2* are of the same gender as the pronoun) there is a distinct tendency for people to construct and interpret the sentence such that the pronoun refers to *NP1* in the case of some verbs, and *NP2* in the case of some others. (Some verbs are neutral.) The strength of this tendency is the verb's causal valence.

Garvey et al (1975) determined the causal valence of a number of verbs by asking subjects to complete sentence fragments in the form of (6-9) with a suitable reason for the action described therein; to distract them from the potential ambiguity, subjects were told that the experiment was about people's motivations, and apparently the subjects performed the task unaware of the ambiguity. For each verb, the proportion of responses favouring *NP2* as the referent was defined to be its causal valence. In a subsequent experiment (Caramazza et al 1977) it was found that subjects took longer to comprehend sentences such as this:

(6-10) Patricia won the money from Janet because she was a careless player.

where semantics force an interpretation contrary to the usual causal valence of the verb.

We can see that if an NLU system had the implicit causality of each verb marked in its lexicon, this information could be used to help find the preferred antecedent in potentially ambiguous cases.²

The phenomenon of causal valence may be explained as simply being a special effect of plausibility. The causal valence data in Garvey et al (1975), Caramazza et al (1977) and Grober, Beardsley and Caramazza (1978) suggest that verbs with an *NP2* bias are exactly those describing an action normally

²The similar constraints which verbs of introspective experience place on anaphors could also be included; see Springston (1976) and Caramazza et al (1977).

performed in response to an external cause, while *NP1*-biased verbs describe an initiating action.

So, for example, in (6-11), where the verb is *NP1*-biased:

(6-11) Ross apologized to Daryel because he . . .

it is most likely that Ross has initiated the action – the cause lies with him – and so he is the actor in the subordinate clause, and hence in turn probably the referent of its surface subject. On the other hand, in the case of (6-12) with an *NP2*-biased verb:

(6-12) Ross scolded Daryel because he . . .

it is most likely that Ross is responding to something Daryel has done, and hence the cause lies with Daryel. It follows that a text like (6-13), in which it is hard to determine the initiator with any confidence, is more ambiguous than one in which there is an actor who is clearly the initiator:

(6-13) Ross telephoned Daryel because he wanted an apology.

Unfortunately, the nice computability of implicit causality does not seem to generalize; with the exception of interrogativization (Garvey et al 1975) and certain strong modal verbs (Grober et al 1978), most linguistic variations on the ‘‘pure’’ form of (6-9), such as negation, passivization or the use of *but* instead of *because*, tend to attenuate the effect of *NP2*-biased verbs, moving them towards *NP1*. It is possible that analogous measures may be found that apply in different contexts from (6-9). However, unless these contexts are rather general, such measures are of little use; indeed, one wonders if enough sentences of the form of (6-9) are ever encountered to make the inclusion of implicit causality in an NLU system a worthwhile endeavour.

6.7. Semantic distance

To check for the possibility of an antecedent being non-identically related to a referent (see section 2.4.2), the SEMANTIC DISTANCE between the referent and its candidate antecedents needs to be considered. The semantic distance between two concepts or entities is simply a metric of how ‘‘similar’’ they are. If a candidate is within a certain threshold semantic distance of the referent, then the possibility that it is an antecedent must be considered.

How to compute a semantic distance and set a threshold are major research problems that underlie much of the research in anaphora understanding. In sections 5.2, 5.3 and 5.5 we saw approaches in which a knowledge representation was used to provide a measure of semantic distance. However, as we saw in 2.4.2, computing the semantic distance relationship may involve complex inference, and no-one has yet attempted a general solution.

6.7 Semantic distance