

The Use of Focus as a Tool for
Disambiguation of Definite Noun Phrases

by

Candace L. Sidner
MIT Artificial Intelligence Laboratory,
Cambridge, Massachusetts

1. Introduction to the Problem

When speakers¹ utter or write sentences, they use certain words in the sentence to refer to people, places, object, times, events and ideas which exist in the real world. When sentences are formed into units of two or more sentences, certain words refer back to other referring expressions in the previous sentences. Among the words which can be used to refer to the real world as well as to refer back (which is called the process of co-reference) are noun phrases containing a definite article, such as *the* (called defnps, hereafter). Several aspects of comprehension of defnps are open problems:

- 1.) What is the definition of the reference of a defnp? That is, what do we mean by *reference* in computational linguistics?
- 2.) How are defnps which are used to co-refer into a discourse distinguished from those which refer to real world objects outside the discourse?
- 3.) What methods of search will distinguish the referent of a defnp which refers to an entity outside the discourse context?
- 4.) What different ways can defnps be used to co-refer to other entities in the discourse?
- 5.) How can co-reference of defnps be detected?
- 6.) What inferences and data structures will be needed for the detection process? The work of Winograd [1971], Charniak [1972] and Rieger [1973] suggest that inferencing is crucial to the interpretation of co-reference.

This paper presents a viewpoint from which to answer these questions based on the concept of focus, as developed by Grosz [1977] and the author (Bullwinkle [1977]). This paper extends Grosz' work by establishing a framework for communication and a set of rules for use of focus in discourse. The approach taken here represents an alternative to the inference driven schemes of Charniak and Rieger.

1. I will use the term *speaker* to refer to the producer of a spoken or written discourse and *hearer* to refer to the receiver of the discourse.

2. The Communication Process and Focus

The description of the communication process given here contains four simple assumptions which are generally true and will be taken as true in this work. First, the speaker is assumed to be communicating about something. This assumption implies that the speaker is not speaking gibberish, that there are referring expressions and either requests, questions, assertions or acknowledgements being made. The something which the communication is about will be called the focus of the discourse.² Second, the hearer is assumed to be able to identify what the focus of the discourse is. The speaker wants to communicate about something, and for the communication to occur, the hearer must be able to distinguish what the speaker is communicating about. Third, the speaker is not trying to confuse or deceive the hearer. The speaker uses referring expressions with the intention of referring to someone or something, or with the intention of describing something or some event. In Gricean³ terms, the byword is "Be perspicuous." The final assumption claims that the speaker assumes the hearer knows certain knowledge about the real-world which can be referred to during the communication process. Recent research (Cohen [1978]), as well as the well known work of Searle [1969] and Austin [1962], describe models of the speaker's knowledge of what the hearer believes. In this chapter, the weakest form of such a model is assumed: the speaker assumes the hearer has enough real-world knowledge in common with the speaker to know the entities which the speaker refers to, and that knowledge is what the speaker draws on in constructing a message for a hearer. These four assumptions will play an important part in the discussion of co-reference interpretation which follows.

2. I don't want to suggest that only one thing can be communicated in a discourse, for speakers do direct their attention from one thing to another. For the moment, I will speak of the focus as the first center of attention in a speaker's discourse.

3. Grice, H.P. "Logic and Conversation" etc.

This paper makes the claim that the focus acts as an index function for referring expressions. For those referring expressions which are anaphoric, the focus indicates where to look for an antecedent. For those referring expressions which are names or descriptions of things in the world, the focus acts as a generation center for a process that chooses a representation of a real world entity which fits the name or description. However, the focus of a discourse alone is not sufficient to produce the indexing behavior. The focus must be used in conjunction with a hierarchical semantic network of associations. The network will indicate what other concepts are related to the focus. It is a codification of some of the general knowledge speakers and hearers have about the real world. The network is a dynamic structure because the hearer adds to his/her general knowledge in the process of interpreting a piece of discourse. Focus must also act with a third piece of computational machinery, an inferencing mechanism. It is used to infer from general knowledge and some suppositions that a certain proposition is true.

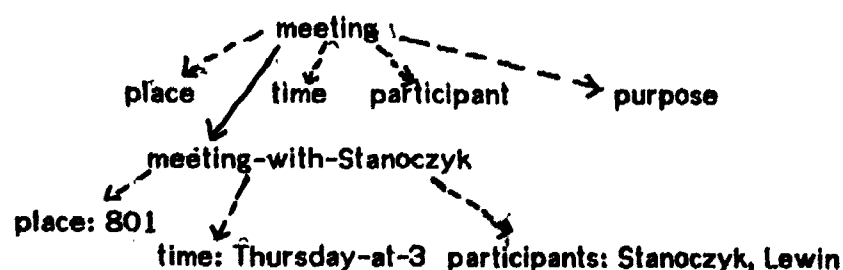
An example will be helpful here. In the discourse below, the focus of discussion is the meeting of D0-1.

- D0-1 I want to schedule a meeting with Ira.
 2 It should be at 3 p.m.
 3 We can get together in his office.
 4 Invite John to come, too.

All four sentences give information about the focussed entity. Thus in D0 both sentences 3 and 4 make no direct reference to the meeting of D0-1. As human hearers, we know that these sentences are related to the rest of D0 because they give information about the focus *meeting*. In D0-3 there are three clues which connect this sentence and the rest of the discourse: the use of *get together*, the co-reference of *we* to the participants of the meeting, and *his office* establishing a place for a meeting. D0-4 introduces an additional participant which can be surmised from the use of *invite*, and the fact that the ellipsis of the event that John is invited to is the focus.

A piece of the hierarchical net needed for D0 is given below in figure 1. A prototype meetings has associated places, times, participants, and purposes. The relation between meeting and place is one of occurrence while the relation between meeting and purposes is one of causality. When D0-1 is encountered, the hierarchical net grows a new member: an *instance* of meeting from D0-1. It inherits the associated entities of meetings and some specific values for the participant entity. D0-2 indicates that something (called *it*) will occur at a particular time. The focus of D0-1 is meeting, so the focus, meeting, is proposed as the antecedent of *it*. To

Fig. 1. Instances of a General Meeting Concept



confirm the proposal, the inference mechanism checks to see if meetings occur at times. Indeed they do, so the proposal of meeting as antecedent of *it* is accepted.

The explanation about the use of focus is not really so simple because the focus of a discourse changes. The interpretation of focus requires a means of recognizing that the focus has changed to some other entity. In D0 the focus begins on meeting, but the *it* in D0-3 has *my office* as its co-referent, not the meeting. Detecting this co-reference requires a means of noticing a shift of focus and using the inferencing mechanism to confirm the choice of co-referent. Focus shift detection will not be discussed here; the reader is referred to Bullwinkle [1977] for a discussion of focus shift where the term "sub-topic shift" is used.

- D0-1 I want to schedule a meeting with George, Jim, Steve and Mike.
 2 We can meet in my office.
 3 It's kind of small, but the meeting won't last very long anyway.

3. Reference in Computational Terms

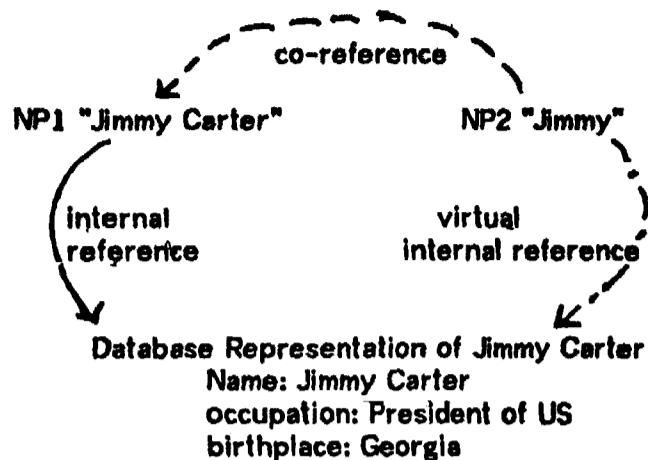
The theory presented here distinguishes two kinds of referring. The first is an internal reference between a noun phrase and some pre-existing database object. That database object represents a real world entity. In Figure 2 below internal reference links the noun phrase NP1 "Jimmy Carter" to a representation of Jimmy Carter (who is described as president of the US, etc.). How the noun phrase and the database object refer to the real world is the classical semantic problem of reference (cf. Kripke [1972] among others) and is beyond the scope of this work. The other kind of referring is co-reference. Co-reference links a noun phrase to another noun phrase. The two noun phrases are said to co-refer, and both internally refer to the same database object, both refer to the real world object. In Figure 1, the dashed link from NP2 "Jimmy" to NP1 is a co-reference link. The dot-dash link from NP2 to the database object is a virtual internal reference link which results from the co-reference link from NP2 to NP1 and from the internal reference link from NP1 to the database object. Internal reference and co-reference links are distinguished because co-reference links can be established more easily using discourse context, which will be discussed in detail later in this paper. In the remainder of this paper when I speak of internal reference, I will drop the phrase "internal" and use only "reference."

A computational theory of co-reference comprehension must answer the following questions about the use of referential terms in natural language:

- (1) Does the expression refer to someone or something?
- (2) What conceptual entity in the memory or the database of the hearer's knowledge, if any, is denoted by the referring expression?
- (3) When does a given expression refer to the same entity as another referring expression?

The expression *Julius Caesar* is used to refer, and can refer to the person represented in the hearer's knowledge as Julius

Fig. 2. Reference Links Between Noun Phrases



Caesar. To answer the first question above, the hearer must decide that names are referring expressions. To answer the second question, the hearer must decide 1) whether *Julius Caesar* refers uniquely and 2) what conceptual entity in the hearer's memory represents the hearer's real-world referent. These two decisions together with the initial assumptions appear to make necessary and sufficient conditions for comprehension since by deciding that *Julius Caesar* refers uniquely and choosing a conceptual entity, the hearer has decoded what entity the speaker was referring to.

There are, however, situations where the hearer's choices to the above decisions and the speaker's intended referent do not coincide. Suppose the hearer decides that *Julius Caesar* refers uniquely and refers to Julius Caesar, who was a Roman emperor. The speaker may also have intended it to refer uniquely, but to the author's deceased cat, whose name was Julius Caesar. Now there are three possibilities: either the hearer knew about Julius Caesar the cat, but decided the expression referred to Julius the emperor, or the hearer only knew about the emperor, or the hearer didn't know of either. In the last case, the hearer "found" a referent by a chance from randomly linking up the name and some memory representation. The last possibility does not fit a description of reference comprehension of any kind. Randomly hooking up information from one's memory to what appears to be a referring expression may be a cognitive act, but intuitively no one would call it reference comprehension.

In the case where the hearer only knew about the emperor, it seems safe to conclude that the reference may have been comprehended, but *incompletely*. As we shall see, there are many other clues in communication about the referent of terms than those given by referring expressions in isolation. Without these, reference comprehension is incomplete because the hearer has no means of knowing whether s/he may have the wrong referent. Even with the best set of clues, the hearer may still choose Julius the emperor. Here we will say that comprehension has taken place, completely but incorrectly, because the hearer has used all the relevant communication knowledge to decode the speaker's message. What can be concluded is that the speaker's rules for reference generation and/or the speaker's knowledge of the hearer is faulty (thereby contradicting the

speaker's assumption above).

In the case where the hearer knows of both possibilities and chooses the incorrect one, the hearer may have erred due to failure to follow other communication clues or again because the speaker's rules and knowledge were lacking. In conclusion, a referring expression is *comprehended as intended*, if and only if the same referent as that intended by the speaker is chosen from the entities in memory. The expression is otherwise just *comprehended* when the hearer chooses an entity from memory which is denoted by the referring expression using all the available communication clues but does not choose the same entity as intended by the speaker. An expression will be considered *incompletely comprehended* if the hearer fails to use all the communication cues available at the time the communication occurs.

So far I have not considered the possibility of error on the part of the hearer because of the hearer's beliefs. Suppose, for example, that the hearer believes the speaker hates to even speak of cats. Then the hearer may conclude that *Julius Caesar* is most likely a reference to the emperor of Rome. I am not going to consider this possibility in the forthcoming discussion; instead I will restrict the discussion to cues from the communication process. Hearer beliefs raise a separate set of philosophical as well as computational problems and extends the scope of this study too broadly. However, the issues are significant in the total picture of reference and co-reference comprehension.

In the remainder of this paper I will consider co-reference comprehension just from the hearer's point of view. Thus in discussing referential and co-referential expressions, I will be concerned with a model of how the hearer disambiguates these expressions used in discourse. By symmetry, one might suppose that the generation of referential expressions by a speaker could make use of a similar model. Such a supposition will remain untested in this paper and is to be verified by later work. Furthermore, I will not be concerned with comprehension as intended since this process requires the additional information of what the hearer believes that the speaker knows about. Instead I will point out at various times how the theory under discussion would need modification if hearer's beliefs were included.

4. Problems with Definite Noun Phrases

Definite noun phrases can be used to refer to entities in the real world. Russell [1905] says of the expression *the author of Waverly* that it denotes Sir Walter Scott, and that when it is strictly used, a defnp denotes uniquely.⁴ Thus by using a definite article, a speaker is saying in effect "there is one object in the world denoted by the phrase that follows, and I mean that one." Of course a defnp may be used to denote someone without *actually denoting* anyone, as is the case with *the woman who wrote Waverly*. This defnp is used to refer to someone, but there is no conceptual representation in the hearer's (or for that matter, the speaker's) memory which corresponds to a real world

⁴4. By *strictly used* Russell means used without ambiguity.

entity assuming the normal case.⁵ Nothing in the syntactic or semantic form of the expression itself suggests that the expression has no denotation. How can the hearer determine whether the defnp refers to someone or not? Of course, if there exists a memory entity the author of *Waverly*, which is attributed as male, the hearer can decide that the expression does not refer to anyone on the basis of a contradiction. But if no memory entity exists, the hearer cannot decide whether *the woman who wrote Waverly* refers to anyone. This way of looking at defnps, however, fails to account for all the phenomena of defnps because it involves an assumption which is not true.

The Russellian analysis has difficulties because defnps are not always used to refer. The problem is not only whether a particular defnp actually denotes a real world object; it is also a question of whether the defnp is intended to refer at all. Even more surprising, a defnp may be used to refer, but the speaker may not intend for the hearer to know the referent of the defnp; the defnp form is used to indicate that the referent is knowable, but possibly not significant for the communication at hand. Donnellan [1977] points out that some defnps are used attributively. If we happen upon Smith who lies dead with foul wounds, one can say "Smith's murderer is insane." Used attributively, *Smith's murderer* does not refer to anyone, and the phrase does not describe a particular person. It is as if to say, Smith was murdered and the murderer, whoever that may be, is insane. Thus the speaker using an attributive defnp does not assume that someone fits the description, whereas with a referential defnp the speaker expects the hearer to realize who is being pointed to.

The other distinction a speaker can make is to use a defnp to indicate that the referent is knowable. Thus if one says:

- (1) Larry read a lot of linguistics in the hospital.
- (2) Larry read a lot of linguistics in a hospital.

the (2) usage is not the same as the (1). While the hearer does not know which hospital *the hospital* refers to, it is clear it refers to some particular one. Comprehension of the referential term does not involve finding a memory entity which represents the real world entity that the expression refers to. For reference comprehension, this concern is considerable since *the Chinese government* in (3) does not demand reference disambiguation, while (4) does; the disambiguation is difficult because the expression can refer to more than one thing.

- (3) John got help from the Chinese government in adopting an Oriental child.
- (4) Get a visa for your trip from the Chinese government.

Another difficulty with defnps is that sometimes they are used not to refer to or to describe specific individuals or objects, but to characterize a class of entities with the

5. Possible world semantics will not be discussed here. Issues of transworld identity and designation by definite descriptions may require more machinery than is considered here.

properties of the head noun phrase and any of its modifiers. Thus (e1) used in this way does not refer to an individual. It characterizes a member of the class of individuals who are men and book writers. (e1) is similar to attribution except that the description applies to a class.

(e1) the man who writes books

So far, then, the following classes of defnps can be stated. A defnp that is used to refer uniquely to one entity, whether or not such an entity exists in the real world, is a specific defnp. A defnp that characterizes a class of entities by means of an individual whose properties are delineated by the properties of the head noun phrase and its modifiers is a generic defnp. A defnp is attributive if it describes an entity without referring. A defnp can be ambiguous in use (u-ambiguous hereafter) if its use as a specific, attributive or generic is not identifiable, while a defnp is ambiguous in reference (r-ambiguous hereafter) if it is used specifically and there is more than one entity fitting the description of the defnp.

5. General Role of Context in Disambiguation

Little of language, if any at all, is said without some surrounding contexts of information. For example, most conversations happen in a location where there are other objects present. Most stories have at least the context of there being a story teller, a hearer and the story being told. There are contexts with more presumed common knowledge, such as what the hearer knows of the speaker's own identity or some shared additional information between them (e.g. they have children or parents in common).

Contexts are needed to determine what a defnp refers to. If I say (5), when I am standing in my kitchen with a friend, the defnp, if specific, must refer to some unique object in the world.

(5) Get me the hot dish holder.

There may be lots of hot dish holders denoted by conceptual entities in my friend's mind, but I am referring to a specific one. Since nothing in (5) distinguishes the one I mean from the whole collection, either I have misused the language, or there is a context which contains only one such hot dish holder, and my friend is aware of that context at the time of my saying (5). In this case, the necessary context is the kitchen, and the referent is probably an item in the kitchen. Reference made to an object external to the conversation is called extra-sentential reference. It is discussed here to exemplify the role of one class of contexts used in reference determination. I call contexts of reference which exist in addition to the one created by the discourse implicit contexts. In this paper I will show how use of implicit contexts can avoid the problem of searching a general database for the entities denoted by defnps.

Other defnps make use of different implicit contexts. Instead of an implicit context consisting of objects near the speaker, the implicit context may be events that the speaker believes are common to the hearer. The speaker who opens a dialogue with (6) below is assuming some previous context (a discussion with the hearer or some other situation) where the reference of *the A.I. Lab Language Group* was first established. In (7), the speaker is again assuming a pre-established referent, but since the hearer may know of several different

dogs, some specific context must be chosen that will distinguish a single dog. Later in this paper some heuristics for choosing a context will be discussed.

- (6) The A.I. Lab Language Group wants to meet next week.
 (7) The dog is sick again.

Contextual information of yet another kind appears in story telling. At the beginning of a story, the hearer expects characters to be introduced. Sometimes this is done with indefinite noun phrases, which are a way all discourses introduce new items, but often a story-teller uses names or defnps as (8) below shows.

- (8) The heiress lived the life of a recluse. She died under mysterious circumstances, but the murderer was never found.

(8) is not a case of cataphoric referencing (referring forward in a text) since the phrase *the heiress* can fully specify an object itself. However, hearers of (8) do not have to search their memories for a referent to *the heiress* in (8). They use the context of story beginning to guide them in reference disambiguation.

6. Distinguishing Generics

Defnps must be disambiguated as generic or non-generic. As will be shown later, generics in the midst of a discourse can be easily disambiguated, but in an initial sentence only implicit contexts may exist in which a co-referent can be found to use in disambiguation. Implicit contexts may be helpful in some cases, but in general they are not sufficient to indicate the interpretation. However certain rules can be postulated based on observable sentence data. This data indicates that there are several levels of sentential and phrasal information used for disambiguation. The rules which will be summarized⁶ here give preferences for generic and non-generic readings.

Some rules govern whether the defnp itself is preferred as a generic or non-generic reading. A small collection of phrases like *the sun*, *the moon* and *the president* default to specific, well defined entities. Which entities are defaulted to depends on the presence or absence of an implicit context to which the phrase may co-refer. For other defnps, a "yes" answer to (1) of the following questions indicates a specific reading preference, while a "yes" to (2) and (3) indicates a generic.

- 1.) Is there a specific individual so described?
- 2.) Is there a class use acceptable for this pair?
- 3.) Are there many individuals described but none outstanding?

"No" answers to all causes a preferred specific reading. Question (2) is necessary because generic readings are difficult to obtain for certain classes of entities. Thus color words like *black*, *red*, *yellow* and *white* applied to *man* describe a class specification while the other color words indicate a non-generic description.

6. See Sidner [forthcoming] for a full specification of the interpretation rules.

In addition to phrasal preferences, predicate argument relations for certain verbs may indicate a preferred non-generic reading. In these cases, a u-ambiguous defnp will be taken as a non-generic, and a generic defnp will cause an odd sentence.

- (9) The black man was moving towards the window.
 (10) The woman who reads *Total Woman* is coming to dinner on Saturday.

Some classes of speech acts are also distinctly generic. The *is-a* sentence below is always generic; this reading may result from the use of *is-a* to indicate further characteristics of the subject. The announcement speech act in (12) is generic as long as an implicit context does not exist which contains an acceptable co-reference for the defnp.

- (11) The elephant is a large mammal.
 (12) I want to tell you about the orangutang.

Speaker-hearer assumptions about perspicuity can force a reading to be generic or non-generic. The defnp in (13) is forced to be read generically because a specific reading would be r-ambiguous and therefore not perspicuous. On the other hand, (14) is odd since *invite* requires a non-generic object. However, because the speaker is assumed to be maintaining perspicuity, the hearer may attempt to read the defnp in (14) as a non-generic.

- (13) Bill considers the black man to be the source of Boston's social unrest.
 (14) Invite the man who reads *The Grapes of Wrath* to dinner.

I want to emphasize that the noun phrase, verb phrase and sentence level rules are only preferences for readings. In the worst case, as (15) shows, an initial sentence may contain a u-ambiguous phrase which, while preferred as generic, can be used either way.

- (15) The robot is replacing the car.

7. The Explicit Backwards Co-reference Rule for Defnp Disambiguation

Using the concept of focus, rules governing the co-referential use of defnps in discourse can be stated. The rules for defnp co-reference which follow depend upon the ability of the hearer to identify focus. This process is a complex one and will not be discussed here. The reader is referred to Sidner [forthcoming] for full details. In brief, the focus of a sentence depends upon predicate argument relations and in some cases, special syntactic forms, such as clefts and pseudoclefts.

In the simplest formulation, the rules for defnp co-reference states: the discourse focus provides a reference point for the co-references of defnps. As I will show below, the rule contains several subparts which must be stated separately. In this paper I will refer to cases of a defnp used anaphorically as explicit backwards co-referencing (EBC). The EBC rule states that a defnp with the same noun phrase head as the focus, and which appears in a sentence following the sentence with the focus, is co-referential to the focus. The more common forms of explicit backwards co-referencing are found in D1 and D2 below:

- D1-1 I want to have a big party, with lots of guests.
 2 The party ought to be on Saturday so everyone can come.

- D2-1 I'm going to tell you about the elephant.
 2 The elephant is the largest of the jungle mammals.
 3 He weighs over 3000 pounds.
 4 At one point in its existence, the elephant had to protect itself from the lion,
 5 but now its herds are so large, that most lions won't even venture near.

What the reader will notice about D2 is not only the co-referentiality of focus for the second and third uses of *the elephant*, but also the co-referentiality implies that these uses are generic. Where defnps in isolation are often ambiguous on the generic-specific classification, in discourse context, this rarely occurs since the focus provides the class type for the defnp. As stated, the EBC rule makes a true prediction about u-ambiguous defnps which occur in sentences following the focus: they are co-referential with the focus, and hence disambiguated as non-generic. D3-2 below contains a defnp which is u-ambiguous in isolation, but in the discourse context, it refers to George's elephant, the reference of the focus.

- D3-1 I sent George an elephant last year for a birthday present.
 2 The elephant likes potatoes for breakfast.

The EBC rule is inaccurate when applied to strictly generic defnps, and where it fails, the role of phrasal and sentential level processing in co-reference comprehension is indicated. D4 is an indication of the problem. D4-2 is generic in isolation. Even in the context of D4, where the focus is Mary's ferret, hearers interpret the underlined defnp as generic.

- D4-1 Mary got a ferret for Christmas last year.
 2 The ferret is a very rare animal.

The context cues of discourse are not strong enough to reverse a strongly generic reading of a defnp. In order for this to be so, sentential level processing must have occurred without consideration of the demands of the context. Since the EBC rule as stated predicts co-reference in cases like D4, it must be revised: specific and u-ambiguous defnps which contain the same noun phrase head as the focus, and which follow the focus in the discourse, co-refer with the focus.

A further refinement on the EBC rule is needed. Consider *the fairy tale book* in D5-2. The EBC rule predicts it will be co-referential with the focus of book in D5-1. In fact, English speakers find D5-2 an odd sentence in the discourse because it is not clear what *the fairy tale book* has to do with the rest of D5.

- D5-1 I bought a book today.
 2 The fairy tale book is by the Brothers Grimm.
 3 It is really well illustrated.

It seems that defnps which co-refer with the focus cannot contain anymore information than is known about the focus. Thus one could say following D5-1. "The book I bought is a fairy tale book by the Brothers Grimm" (since D5-1 states that

the speaker bought the book), but one cannot say D5-2. Why can't a defnp that contains more information than the focus co-refer to the focus? Returning to the discussion of focus-shift earlier, a referring expression following the focus is either co-referential to the focus or introduces an entity which is the potential new focus of the discourse. The difficulty with phrases like *the fairy tale book* is that one cannot tell if it is intended to co-refer, or because it is somewhat different from the focus, intended to be used as a potential new focus. The EBC rule must be revised to state: specific and u-ambiguous defnps which contain the same noun phrase head as the focus, which follow the focus in the discourse, and which do not contain more information than is known about the focus co-refer with the focus.

The EBC rule explains why a negative existential cannot be referred to using a defnp. A sample case, from Karttunen [1968], is given in D6. D6-2 is generally regarded as an unacceptable sentence following D6-1. The sentence is certainly grammatical, so the assumption by Karttunen is that the referential term *the car* is being used in some inappropriate manner.

- D6-1 I don't have a car.
 2 * The car is black.

The EBC rule predicts that *the car* co-refers with the focus in D6-1. But *a car* in that sentence does not have a referent (because the speaker has just said so). Thus the use of the defnp in D6-2 causes the hearer to expect a reference when in fact there is no referential entity.

A similar case, (16), also from Karttunen, does not involve negative existentials, but entities within modal contexts:

- (16) * Mary expected a present from John although the present was expensive.
 (17) Mary expected a present from John although the present wasn't the thing that worried her.

The defnp in (16) according to the EBC rule must co-refer with the focus. What is significant is that the co-reference is acceptable, as (17) shows. What is odd about the second clause of (16) is the predication. This paper cannot give an account of such semantics, but intuitively, it seems odd to predicate the property of being expensive to something one expects. Thus as long as there is a co-referential entity specified by the focus, a defnp may be used, but the predication about the defnp must be semantically meaningful.

Another form of explicit backwards co-referencing is slightly different than the previous examples. It involves the use of lexical generalization of the focus. Grosz [1977] first categorized the relation of focus and defnps with a more general noun in the noun phrase head. In D7, *the poor old beast* is a lexical generalization of *the dog*, that is, its head noun is a term which is a class generalization of the focus.⁷

7. This term comes from the observation of Halliday and Hasan [1976] that lexical cohesion includes the use of reiteration of four types: same word, synonym, superordinate, and general word.

Determining the class generalization of the focus is possible when the focus is represented in the way that is assumed in this paper: as an association network with an is-a hierarchical structure. Using that hierarchy, it is possible to determine whether a phrase like *beast* is hierarchically related to *Salamut*.

- D7-1 Harold took his Salamut to the vet yesterday.
2 The poor old beast was quite lame.

One might expect that some constraint on the amount of information in the lexical generalization of the focus is needed. This is the case, since the underlined defnp in D7-2' is unacceptable following D7-1 as a co-referent with the focus.

D7-2' The beast who is old was quite lame.

-2" The mangy, snarling, unfriendly beast was quite lame.

It appears from all the cases I can find that any post-nominal modifiers on a noun which is a lexical generalization of the focus force the defnp to be non-coreferential with the focus, while pre-nominal modifiers, no matter how complex, preserve co-referentiality. It is unclear why pre-nominal modifiers and post-nominal modifiers have these different behaviors.

8. Implicit Backwards Co-reference

Many definite noun phrases which occur in discourse are not cases of backwards co-reference to the focus. Grosz [1977] suggested that the focus brings other items implicitly into focus as well, by means of association. Such defnps are related to the focus in one of several ways. Since the focus is well specified, these relationships can be easily determined. The focus acts as an anchor point for finding references for such defnps. In D8, the defnp *the time* refers to the time of the discourse focus, the meeting. This defnp use I will call implicit backwards co-reference. Such cases are to be distinguished from explicit backwards co-referencing because the defnp is co-referential with an entity that is closely associated with the focus rather than to the focus itself. The phenomenon of association between two noun phrases has been cited by Norman and Rumelhart [1975].

- D8-1 The pa group wants to have a meeting.
2 The time will be 3 p.m. on Tuesday.

Implicit backwards co-referencing is constrained by the association network surrounding the focus. Any entity closely associated with the entity which represents the focus can be mentioned using a simple defnp. Thus in D9, sentences with acceptable defnps as well as ones with unacceptable defnps are given.

- D9-1 I went to a new restaurant with Sam.
2 The waitress was nasty.
3 The food was great.
4 The soup was sa ty, but the wine was good.
5 * The rug was ugly.

Non-simple defnps have infinitely more variety because the modifiers can specify the relation of the defnp to the focus at hand as in D9-6. Non-simple defnps which do not suggest some connection are less acceptable, but hearers, in reliance on the perspicuity maxim, may attempt some connection. Thus

if D9 included D9-7 below, some hearers might attempt to connect the defnp with focus.

D9-6 I like the band that plays there.

-7 The elephant with the green tutu danced an impressive jig.

Another use of focus is as an inference point for inferred co-references. Inferred co-references, like *the murderer* in (8), presented here as D10, are not mentioned explicitly in the previous discourse nor can they be considered closely associated to the focus on general principles. Their use reflects an inference about the focus on the part of the speaker.

D10-1 The heiress lived the life of a recluse.

2 She died under mysterious circumstances, but the murderer was never found.

In D6, *the murderer* represents an inference that the heiress' death was due to a specific type of circumstance, a murder. Such an inference is possible given a Fahlman [1977] type net with two inference points like heiress and murderer (and the information associated to heiress from the context thus far); from the net, the relation of the murderer to the heiress can be inferred. Such an inference does not produce a real-world entity to which *the murderer* refers. Instead the inferred relation of murderer and heiress provides sufficient information to produce the entity if it exists in the database. When a denotation does not exist in the database, the inference between the murderer and the heiress suggests that the speaker is attributing of some individual that s/he is a murderer.

A concrete example will illustrate my point. Suppose the hearer knows that the heiress was killed by Jones. Then on hearing D10, the hearer not only concludes that *the murderer* refers to the murderer of the heiress, but also based on that conclusion, the hearer decides that Jones is denoted by the referring expression. However, another hearer upon hearing D10 and not knowing what the first hearer knew, could only conclude that *murderer* is attributed of a person who is assumed to have murdered the heiress. The referent is not known to the second hearer, but if someone were to tell him/her that Jones murdered the heiress, the hearer could conclude who *the murderer* refers to. In effect, the defnp used in this way points out the attributional use of expressions which Donnellan has observed. The argument presented here is not only about the nature of focus; it is a statement of what information is sufficient to make up a description which can denote a unique entity. Viewing inferred co-reference defnps as attributions has an implication for a computational model which disambiguates such defnps. This model must be able to use an expression without knowing its referent and be able to link up the denotation to the referring expression if some knowledge makes that denotation available at a later point.

Implicit and inferred co-reference at first glance appear to be one in the same thing. The discourse below, from Karttunen [1968], will indicate just how the two differ.

D11-1 I was driving on the freeway the other day.

2 Suddenly the engine began to make a funny noise.

- 3 I stopped the car.
- 4 When I opened the hood, I saw that the radiator was boiling over.

With focus of *freeway* in D11-1, the relation of the engine can be found since vehicles are driven on freeways and vehicles have engines. The association chain here suggests that the connection between D11-1 and *the engine* involves a few inferences. These inferences are part of a hearer's general knowledge and true of the world. They are part of the knowledge in the association network. With D10, however, the inference about the murderer involves a supposition which is not necessarily true, since dying under mysterious circumstances does not necessarily imply murder. The distinction between implicit and inferred co-reference can thus be stated: implicit co-reference involves inferences which are true about the world, while inferred co-reference involves a supposition which the speaker has made which is not necessarily true.

Another kind of implicit focus-defnp relation exists in D12. I call this relation the set-element relation since *the clown with a unicycle* is an element of the set of clowns which the focus denotes.

- D12-1a I went downtown today,
- 1b and there were clowns performing in the square.
- 2 The clown with a unicycle did this really fantastic stunt.

As with inferred references, the focus does not make it possible to identify a specific denotation with the referring expression. Instead the focus is the set of which the referent of that phrase is a member. These cases are easier to distinguish than those of inferred reference because the head noun is the singular of the noun phrase represented in the focus. Unlike defnps using the EBC rule, set-element co-reference demands a modifier that distinguishes it from the focus. Without the modifier, there is no means of determining which member of the set is being discussed. Gross [1977] says of cases like the set-element relation that an inference is needed to establish additional properties of an object in focus. However, because the head noun phrase is the same as the focus, while the modifier is different, the relation can be established without the need for inference.

Another kind of focus relation, which I call computed reference, can be seen in D13. Here *the last meeting* does not refer to the meeting mentioned in the previous sentence, but that meeting can be used as a point for determining a last meeting, if one is known in the database, else it is a description of the entity required, as with inferred references and set-element references. Several modifiers - *first, last, next, second* and the other ordinals - are used in this way.

- D13-1 Aunt Het's Sewing Bee wants to have a meeting this week.
- 2 The meeting should be on Tuesday.
- 3 The last meeting, which was at 5, was too late, so schedule this one earlier.

From these examples, the nature of focus in discourse can be re-emphasized. It is the focus which

connects sentences of the discourse. In the process of determining the focus relation between a defnp and the focus, the link in the discourse is created. From these examples, one can predict that there ought to be cases of defnps which bear none of these relations to the focus and which result in unacceptable discourses. This prediction is accurate as D12' shows.

- D12'-1a I went downtown today,
- 1b and there were clowns performing in the square.
- 2 * I saw the chair

The difficulty with D12' is that a chair is not associated with clowns, and the discourse does not suggest any suppositions that would link chairs and clowns inferentially.

There are, however, cases of defnps which do not bear any relation to the focus and which are perfectly acceptable in discourse. Consider *the chairman of the math department* in D14 below.

- D14-1 George wants to have a seminar to discuss representation in frame-like languages.
- 2 He wants to invite the chairman of the math department.

The focus in D14 is the seminar of D14-1. The math department chairman is not directly related to the seminar. However, the focus does provide an important piece of information; it is the source of the ellided event to which the chairman is invited. Hence D14 is different from D9-6 or D12' where there is no link between the focus and the sentence in question. What can be concluded is that the focus is not a useful reference point for the referent of the chairman of the math department. The sentence is not odd because of focus links. Thus something outside of the discourse context must contain the needed denotation. This conclusion points to the limitation of the focus: it captures only those reference relations which are internal to the discourse. In a sense, the focus is a summary of the discourse context and what the hearer knows about it.

How can the denotation of *the chairman of the math department* be found? Since the denotation of *the chairman of the math department* lies outside the discourse context, a more global context such as that of the speaker's situation in time and space must be used to determine a context of possible referents. This context must be limited because there are potentially many math department chairmen in the speaker's and hearer's memories. I do not intend to describe just what such a context will look like, but I do want to indicate that it may be "grown" from a search through the associative network to other entities which are related to any of the objects in the discourse, including the speaker and hearer. The association net includes not only abstract representations of general classes of real world entities, it also contains representations of real world objects. The associations between real world objects can be gathered by a search method which collects associations close to the focus and then extends for other associations until one is found that matches the defnp in question. Of course, it is possible that no entity will be found: such a circumstance is yet another example of the hearer knowing that a defnp refers without being able to tell who the

speaker had intended as the reference. The implications of extra-discourse contexts for computational models is clear: models are needed of what the speaker assumes his/her hearer knows about, so that the speaker does not produce referring expressions which the hearer cannot disambiguate, and models are needed of what the hearer assumes the speaker has beliefs about so that the hearer can decide what to do with referring expressions which s/he cannot disambiguate.

Earlier in this paper I mentioned the use of such defnps as *the sun*, *the moon*, and *the Earth*. These defnps have default referents in initial sentences of a discourse. Inside a discourse, they can act in one of two ways: related to the focus as either a backwards co-reference, implicit co-reference, inferred reference, set-element reference or computed reference; or they may act as references to entities outside the context. The EBC rule predicts that such defnps will be taken as co-referring when the focus has the same defnp head. Thus a defnp like *the sun* or *the moon* will refer to its default value only if the focus does not predict a referent based on one of the five co-reference relations discussed here. These kinds of defnps are especially noteworthy because they are a clear example of a defnp that can be used in either role.

The four related co-reference relations specify ways in which a defnp can be predicted as co-reference to an entity associated with the focus. Other defnps either refer to objects outside the discourse context or the defnp fails to refer as intended by the speaker. The former case is marked by the presence of discourse links elsewhere in the sentence to the focus or by the use of default reference. In those cases where the reference as intended fails, the hearer may attempt to create a connection to the focus, as was suggested with D10-7, and thereby maintain the perspicuity contract. Alternatively, the hearer may fail to understand the referent.

Focus can be used for disambiguation of generic defnps in a manner similar to the cases presented above. As discussed earlier, the EBC rule predicts that a defnp which is u-ambiguous will be generic or not based on the focus, and that a strictly generic defnp is independent of the focus. The strictly generic defnp case, as shown in D15, presents one means of shifting focus in a discourse with defnps.

- D15-1 I got a new ASR 33 this week.
2 The ASR 33 is an old but reliable output device

If D15-3 were "It was available long before the newer electronic consoles," the focus would have shifted from the ASR33 which the speaker had gotten to the generic entity of ASR33 teletypes.

The strictly generic defnp used after a non-generic focus is just one case of implicit backwards co-reference using associations. Other associations occur as well. However, implicitly related defnps are not distinguishable as generic unless a full modifying noun phrase is attached, as D16 shows. *The time* as a simple defnp can be used only as an implicit co-reference to the focus of party. Only the complex noun phrase has the syntactic and semantic distinctions which

reflect the generic usage. The simple defnp used implicitly takes its generic/specific classification from the focus.

- D16-1 I want to have a party.
1a The time of a party is hard to decide on.
2b The time is hard to decide on.

Set-element implicit co-reference exists for generic foci as well as specific foci. A significant difference is that the foci for the generic case can be a singular defnp, or they can be a plural noun phrase with either a definite article or no article. The set membership is indicated by a distinguishing modifier, just as with specific set-element implicit co-reference. An example of generic set-element co-reference is given below with both a singular defnp focus and a plural unspecified focus.

- D17-1 The Australian aborigine represents an almost extinct hunter-scavenger social group.
1' Australian aborigines represent an almost extinct hunter-scavenger social group.
2 The aborigine in the southern sections of Australia sometimes gathers food, but the other aborigines do not.

Inferred generic co-references also occur. In D18, *the owner of a motorcycle* is a generic defnp:

- D18-1 Alronso was in an accident with a motorcycle last week.
2 I think the owner of a motorcycle ought to be required to take driving lessons.

The owner of a motorcycle is generically related to the first sentence by an inference of what happened to the agent. The same kind of machinery that is used for specific inferred co-references can be used for making these inferences as well. How can generic inferred co-references be distinguished from specific inferred co-reference? A strictly generic defnp as in D18-2 remains generic. Those defnps which are u-ambiguous at the sentential level, as with D10-2, can be disambiguated as specific because of the relation to the focus.

The use of a semantic network with a focus and inference mechanism results in a computational theory of co-reference which makes use of representation properties such as prototype of entities, hierarchical connections and associative links between entities in the representation. The use of focus for co-reference rules such as the EBC rules, implicit backwards co-reference rules and inferred co-reference relies on this net representation. An inference mechanism is necessary both to verify co-reference predictions and to test suppositions used in inferential co-reference. With the net and a focus, predictions about acceptable co-reference for noun phrases has been shown and verified with linguistic evidence. Psychological predictions, such as implicit co-reference requiring more time than explicit co-reference, can also be tested although the related psychological literature has not been discussed in this paper. The limits of focus as a co-reference mechanism suggest that focus is central for co-reference of noun phrases related to previous discourse. For noun phrases that refer outside the discourse, focus may also be used to generate a context of entities from which a co-referent may be chosen. Further research can extend the focus mechanism to rules involving

97

other types of noun phrases and the personal pronoun anaphors.

9. Acknowledgements

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under the Office of Naval Research under Contract Number N00014-75-C-0643.

10. References

Austin, J.L. [1962] *How to Do Things With Words*, Oxford University Press.

Bullwinkle, C. [1977] *Levels of Complexity in Discourse for Anaphora Disambiguation and Speech Act Interpretation*, Proceedings of the Fifth International Joint Conference in Artificial Intelligence, August 1977; also M.I.T. A.I. Lab Memo 413.

Charniak, E. [1972] *Toward a Model of Children's Story Comprehension*, M.I.T. A.I. Lab TR-266.

Cohen, Philip R. [1978] *On Knowing What to Say: Planning Speech Acts*, University of Toronto Department of Computer Science, T.R. No. 118.

Donnellan, Keith S. [1977] *Reference and Definite Descriptions in Naming, Necessity, and Natural Kinds*, S.P. Schwartz (ed.), Cornell University Press, Ithaca.

Fahlman, Scott E. [1977] *A System for Representing and Using Real-World Knowledge*, M.I.T. A.I. Laboratory AI-TR-450, December.

Grosz, Barbara [1977] *The Representation and Use of Focus in Dialogue: Understanding*. Stanford Research Institute Technical Note 151, Menlo Park, California

Halliday, M.A. and Ruqaiya Hasan [1976] *Cohesion in English* London, Longman press.

Karttunen, Lauri [1968] *What Makes Definite Noun Phrases Definite?* Paper p-3871, The RAND corporation, Santa Monica, California.

Kripke, Saul A. [1972] *Naming and Necessity*. in Semantics of Natural Language, Davidson and Harman (eds.) Reidel Publishing Co Boston.

Norman, Donald A. and David E. Rumelhart [1975] *Explorations in Cognition* W.H Freeman, San Francisco.

Rieger, Charles J. [1974] *Conceptual Memory: A Theory and Computer Program for Processing the Meaning Content of Natural Language Utterances*: Stanford Artificial Intelligence Lab Memo AIM-233.

Russell, Bertrand [1905] *On Denoting*, reprinted in Logic and Knowledge, R.W. March (ed.). Allen and Unwin, London, 1956.

Searle, J.R. [1969] *Speech Acts: An Essay in the Philosophy of Language*, Cambridge University Press.

Sidner, C. [forthcoming]. *A Computational Model of*

Co-reference Comprehension in English. Ph.D. dissertation, M.I.T.

Ninograd, Terry [1971] *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language*. M.I.T. dissertation.