

Focusing for Interpretation of Pronouns¹

Candace L. Sidner

Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, Massachusetts 02238

Recent studies in both artificial intelligence and linguistics have demonstrated the need for a theory of the comprehension of anaphoric expressions, a theory that accounts for the role of syntactic and semantic effects, as well as inferential knowledge in explaining how anaphors are understood. In this paper a new approach, based on a theory of the process of focusing on parts of the discourse, is used to explain the interpretation of anaphors. The concept of a speaker's foci is defined, and their use is demonstrated in choosing the interpretations of personal pronouns. The rules for choosing interpretations are stated within a framework that shows: how to control search in inferring by a new method called constraint checking; how to take advantage of syntactic, semantic and discourse constraints on interpretation; and how to generalize the treatment of personal pronouns, to serve as a framework for the theory of interpretation for all anaphors.

1. Introduction

Traditionally researchers have defined the problem of comprehending anaphoric expressions as one of determining *the antecedent* of an anaphoric expression, that is, determining to which word or phrase an anaphoric expression refers or "points". Recent studies in both artificial intelligence and linguistics have demonstrated the need for a theory of the comprehension of anaphoric expressions, a theory that accounts for the role of syntactic and semantic effects, as well as inferential knowledge in explaining how anaphors are understood. In this paper a new theory, based on the concept of focusing in the discourse, is introduced to explain the interpretation of pronouns.

Before a theory can be given, and before even the difficulties in interpreting anaphors² can be discussed,

¹ The research reported in this paper was supported in part by the Advanced Research Projects Agency under contract No. N0014-77-C-0378. Research reported here was also 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 ONR contract N00014-75-0643. I would also like to thank Bob Berwick, Barbara Grosz, David Israel, and the *AJCL* referees for their help in preparing this paper.

² I use the term "anaphor" for an anaphoric expression, and I use "anaphora" in its traditional meaning, that is, as the device of using a word or phrase "to point back."

we must first re-consider what an antecedent is. The traditional definition encounters difficulty right from the start; it is founded on the notion that one word in a sentence refers or points back to another word or phrase in the (same or another) sentence. But words don't refer back to other *words* [Morgan 1978]; people use words to refer to entities in the world. In particular they use pronouns to refer to entities which have already been mentioned in a discourse. Since an anaphoric phrase does not refer to an antecedent, one might want to claim that both the antecedent and the anaphor co-refer to the same entity. That description is adequate for sentence s1,

s1 I think *green apples* taste best and *they* make the best cooking apples too.

though not for discourse D1, where there is no antecedent phrase in the discourse that co-refers with the pronoun "they."

D1-1 My neighbor has a *monster Harley 1200*.

2 *They* are really huge but gas-efficient bikes.

Rather than view antecedence as co-reference, one might propose that antecedence is a kind of *cognitive pointing*, the kind of pointing that causes "they" and "green apples" to point (somehow) to the same class of entities in one's mind. This proposal is problematic for the same reason that co-reference is: people use pronouns when there is no other noun phrase in the

discourse that points to the right mental entity. In D1, "they" refers to bikes which are Harley 1200's as a group, while "a monster Harley 1200" mentions only some particular Harley 1200. "They" seems to be able to refer when used with the previously mentioned phrase "a monster Harley 1200" without the two phrases either co-referring or co-pointing.

If an anaphor does not refer to an antecedent phrase, and if it need not always co-refer with its apparent antecedent (as in D1), then anaphor interpretation is not simply finding the antecedent. Nevertheless the concept of antecedence as pointing back does seem to capture some aspect of the comprehension of anaphors, for when certain antecedent words are missing from a discourse, people often fail to understand what is being said.

Let us define the problem of interpreting and understanding an anaphor in the following way. The phrase "green apples" in s1, when syntactically and semantically interpreted, is said to *specify* a cognitive element in the hearer's mind. In the computational model of that process, this element is a database item, which might be represented by Apples2 in by the schema below:

```
Phrase76:
  string:      "green apples"
  context:    speaker1 think * tastes best
  specifies:  Apples2

Apples2:
  super-concept: apples
  color:       green
  used-for:   cooking
```

The speaker uses the information in a cognitive representation like Apples2 above to choose the phrase "green apples" in s1. The hearer then uses the phrase "green apples" plus the syntactic and semantic interpretation of the rest of the sentence to locate a similar cognitive element in his own mind; it may be slightly different because the hearer may not associate use in cooking with green apples. A cognitive element, such as Apples2, is called the *specification* of "green apples." These elements, present in the memories of speaker and hearer, are of course related to other cognitive elements in their memories.

What is the relation of specifications to the real world? One might like to claim that a reference relation exists between specified cognitive elements and objects in the world, but since referring is what people do with words, this relation is problematic for cognitive elements. Instead, specifications are said to *represent* the objects referred to; that is, they bear a well-structured correspondence to objects in the world. Apples2, the specification of "green apples," represents the objects that are green apples. For phrases such as "Santa Claus," where there is no real world

object to represent, a specification represents the mental schema to which are attached the properties normally associated with this imaginary person.

The phrase "they" in s1 also specifies a cognitive element, namely the same one that "green apples" does. Since the two bear the same relation to the representation Apples2, I say that they "co-specify" that memory element, or alternatively, that the interpretation of "green apples" is the *co-specifier* of the interpretation of "they." Co-specification, unlike co-reference, allows one to construct abstract representations and define relationships between them which can be studied in a computational framework. With co-reference no such use is possible, since the object referred to exists in the world and is not available for examination by computational processes.

Even if a phrase and a pronoun do not co-specify, the specification of the phrase may be used to *generate* the specification of a pronoun. For example, in D1 "they" does not co-specify with the apparent antecedent phrase "a monster Harley 1200," but rather it refers to the class of Harley 1200's of which the apparent antecedent is an instance. Thus anaphor interpretation is not simply a matter of finding the corresponding cognitive element that serves as the specification of the anaphor; some additional process must generate a specification for the anaphor from the related phrase "a monster Harley 1200."

The concepts of specification and co-specification capture the "pointing back" quality of antecedence, and also permit us to formulate an explanation of anaphor interpretation that avoids the pitfalls of the concept of antecedence. Anaphor interpretation can be studied as a computational process that uses the already existing specification of a phrase to find the specification for an anaphor. The process uses a representation of the discourse preceding the anaphor to encode the syntactic and semantic relationships in each sentence as well as co-specification relationships between phrases.³ These definitions in themselves do not constitute a theory of anaphor interpretation. They do, however, make possible a succinct statement of the problem: how does one determine the specification of a anaphor? Also, since we suspect that the specification of an apparent antecedent phrase plays some role in choosing an anaphor's specification, we may ask, just what is this role? We hope for a direct answer to these questions, but before one can be given, let us consider how a theory of interpretation ought to address these questions. A brief look at the

³ In the rest of this paper, I speak of a phrase co-specifying (or specifying) with another phrase, when what I really mean is that the relation is between representations of phrases that have been interpreted by some parsing process, which indicates the sentence syntactic relations, and by a semantic interpretation process, which computes semantic relations among words of the sentence.

difficulties of choosing co-specifiers suggests which issues our theory should cover.

Determining the co-specifier of an anaphor is difficult because there are a multitude of possible co-specifiers in a given discourse, and there is no simple way to choose the correct one. Yet human hearers and readers generally do recover the correct co-specifying phrase intended by the speaker. Human readers and hearers also fail to recover the co-specifying phrase in certain situations; this behavior is just as valuable an observation as garden path phenomena for theories of parsing. A theory of interpretation must predict the pattern of the hearer's and reader's correct and incorrect choices, as well as failures to understand, by a rule-governed account. In addition, a taxonomy of the cases in which specifications are used to generate other specifications must be given, as well as a means of predicting their distribution.

Two other aspects of communication make it difficult to find co-specifiers: the context of discussion and the inferences people make. People use the context surrounding an anaphor in understanding it. If a theory of anaphor interpretation is to capture understanding, it must include a means of encoding discourse context and whatever structure it has; the context must be distilled into a form that preserves its richness without adding overwhelming complexity to the interpretation process. In addition, researchers have discovered that anaphor interpretation involves making inferences, some of which can be complex, each of which must be chosen from a large base of knowledge about objects, people and things. The practical deployment of inferential capabilities for any task requires control: knowing what to infer when, and knowing when to stop. Since the general control problem is poorly understood, a theory of anaphor interpretation must provide solutions to the more specific problem of controlling inference in anaphor interpretation.

The major portion of this paper addresses the issues of determining the specifications of anaphors, with an emphasis on the role of context and inference. The theory of focusing on parts of the discourse is introduced, and a processor and its rules which rely on the focusing theory are discussed. Before turning to this discussion, the previous and extensive research on anaphora are reviewed.

2. Research on Anaphora

The role of context and inference, as well as syntax and semantics, on anaphor interpretation have been explored extensively. A brief look at these explorations indicates the necessity of a new approach. Research on anaphora falls into four broad categories:

1. General heuristics for finding antecedents [Winograd 1972]
2. Syntactic and semantic constraints on anaphora [Katz & Fodor 1963, Woods et al. 1976, Chomsky 1976, Lasnik 1976, Reinhart 1976, Walker 1976]
3. Use of inference to find antecedents [Charniak 1972, Rieger 1974, Hobbs 1976]
4. Analysis of relations among objects in a discourse context [Grosz 1977, Lockman 1978, Reichman 1978, Webber 1978, Hobbs 1979]

Rather than review each approach, I point out the contributions of each type to a theory of pronoun interpretation.

General heuristics, as a means of choosing antecedents, predict reliably in a large number of typical examples. However, no simple characterization fits the wide variety of cases where they fail (see Winograd 1972 and Hobbs 1977); furthermore, the heuristic approach is not theoretically grounded and cannot offer a unified approach to the phenomena.

Semantic selectional restrictions, based on the Katz-Fodor theory of semantic markers, and used by many computational linguists, can reduce the space of possible antecedents, but they cannot be used to eliminate all possibilities, as the example below illustrates (where feeling soft can be said of either a mud pack or one's face):

s2 Put *the mud pack* on *your face*. Notice how soft *it* feels.

Syntactic restrictions, on logical form [Chomsky 1976] and on constituent structure [Lasnik 1976, Reinhart 1976], stipulate conditions in which a pronoun and a noun phrase must have disjoint reference, as shown below.

s3 * Near *Dan*, *he* saw a snake.

s4 * The *man* whose house *he* bought went gold digging in Alaska.

These rules, however, do not stipulate the interpretation of a pronoun; in a general theory they serve only as a filtering condition on the class of possible co-specifications. Furthermore, syntactic restrictions must also stipulate the disjoint reference conditions on reflexive pronouns although no adequate account of these conditions has yet appeared.

Work by researchers in artificial intelligence on *inference* led to methods for forward and backward chaining of inferences to bind the pronoun, represented as a free variable, with some piece of knowledge; with this approach, the pronoun's interpretation was the value bound to the free variable. This approach revealed that inferences about world knowledge are often needed to interpret pronouns. However, these methods failed to control the inference process sufficiently. Charniak, attempting to resolve this problem,

proposed demons that would “wake up” in the appropriate situation (that is, processes which could themselves notice when they were to begin processing). But a large cache of demons would be required, and no assurance could be given that demons would exist in every situation. Most significantly, this proposal said nothing about the situation where two or more demons might apply (who gets control? how are the decisions made?). Furthermore all of the inference-based approaches to pronoun interpretation fail to offer any theoretical approach because they rely on a simple mechanism, (simple variable binding between pronoun and some other phrase) which does not apply in many uses of anaphors, such as D1.

Discourse approaches to anaphora include a technique similar to the inference method; one identifies sentence pairs and determines their semantic relationship as one of elaboration, similarity, contrast, parallel structure; the pronouns are interpreted by variable binding between items of the sentence pairs [Hobbs 1979]. Webber 1978, using a notion of “discourse identifications” (that contain certain semantic and discourse content) similar to the notion of specification, stipulates constraints on the representation of relations among items mentioned in a discourse.

Grosz [1977, 1978, 1981] illustrates how parts of a speaker’s knowledge, relevant to a discourse segment, are highlighted via focusing, a process that reflects what a speaker says and the nature of the knowledge in the space. She shows that the structure of a task affects what items will be focused on in the discourse. Reichman has expanded this paradigm by describing *context spaces*, delineated by their topics. Her analysis shows that within a context space only certain items may be pronominalized. She leaves open the questions: What is the recognition procedure for determining a context space? How does one identify its topic? How does the hearer determine the interpretation of a pronoun, i.e., how does a hearer decide which highly focused items act as the co-specification of a pronoun? All these approaches support the view that since hearers do not have privileged access to a speaker’s mind, other than through what a speaker says, imposing structure on the speaker’s discourse will provide a framework for establishing the interpretation of pronouns.

The remainder of this paper defines the concept of speaker’s foci and shows that they can be used to choose specifications for personal pronouns. The rules for choosing interpretations are stated within a framework that shows:

- how to control search in inferring by a new method called constraint checking,
- how to take advantage of syntactic, semantic and discourse constraints on interpretation,

- how to generalize the treatment of personal pronouns, to serve as a framework for the theory of interpretation for all anaphors.

3. Focus for Pronoun Interpretation

A simple discourse can be used to illustrate how pronouns indicate what the speaker is talking about in the discourse. In the sample below, the speaker mentions two dogs and tells us something about each.

- D2-1 I have two dogs.
- 2 (The) one is a poodle;
 - 3 the other is a cocker spaniel.
 - 4 The poodle has some weird habits.
 - 5 He eats plastic flowers and likes to sleep in a paper bag.
 - 6 It’s a real problem keeping him away from plastic flowers.
 - 7 The cocker is pretty normal,
 - 8 and he’s a good watch dog.
 - 9 I like having them as pets.

The speaker uses the pronouns “he” and “him” to indicate that the poodle is the actor of eating flowers and the object of an action of keeping. The speaker mentions his second dog with “the other” and then uses the pronoun “he” only in predicate adjective and predicate nominative constructions without reference to any events. Initially the speaker focuses attention on both dogs and then turns his attention to each for a part of the discourse. The process by which the speaker uses language to indicate his focus is called the *focusing process*. The items in focus are those that are talked about for a part of the discourse. Items in focus can be used in two ways, as *actor focus* and as *discourse focus*. The actor focus is a discourse item that is predicated as the agent in some event, so for example, the poodle is the agent in eating flowers, while the discourse focus is, roughly speaking, an item the speaker wishes to make several predications about.⁴

The focusing process, when explained from the viewpoint of a hearer, can be described as a process of *tracking* the speaker’s foci as they change over the discourse. A hearer does not have privileged information about what is in a particular speaker’s head, so the hearer must decide what the speaker is talking about on the basis of what the speaker uses as initial referring expressions and subsequent co-specifying anaphoric expressions. The hearer follows a focus and checks to see if the anaphoric expressions which the speaker uses either co-specify with the hearer’s representation of that focus or specify a representation

⁴ This description is not a definition of discourse focus. The discourse focus is defined by the effects of several focusing algorithms; for details, see Sidner 1979.

related to it. The hearer is said to be tracking the speaker's focus because she or he can determine what is being talked about only after the speaker has said something. Focus, as it is used here, is akin to Grosz's notion [1977] of *immediate* focus.

A process model of focusing and focus tracking consists of three distinct processors. The first chooses foci based on what the speaker initially says. The algorithm for this choice depends on many phenomena – see Sidner 1979 for details. Then an interpreter (the *pronoun interpreter*) uses these foci and a set of rules of pronoun interpretation (hereafter “pi-rules”) to interpret the anaphoric expressions in the discourse. This interpreter, like a human hearer, must keep in mind whatever other newly mentioned elements the speaker has introduced, since sometimes an anaphor may co-specify with one of those instead of the elements in focus. A third process updates the foci by decisions that depend on anaphoric interpretations chosen by the pronoun interpreter. During this last phase, the updating process moves one of the foci to a new element of the discourse, if some anaphor co-specified with that new element while no anaphor co-specified with the phrase already in focus. The three processes taken together sketch a simple model of focus tracking; the model appears to behave like its human counterpart in the way it interprets anaphors and in the instances in which it fails to “understand.”

So far, the process model looks circular. Foci are chosen, they are used to determine how anaphors co-specify with the foci and then the co-specifiers are used to determine what's in focus. In fact, the model is not circular, but its steps are cyclic. The processor cycles through the three processes for each sentence in the discourse. This cycling differs from logical circularity because it depends on new information presented over time, the time of each sentence of the discourse.

For example, reconsider the first few sentences of D2, given previously. Initially the speaker mentions two dogs, which become the discourse focus, and then he extends the discussion of them using “one...the other” anaphors. The actor focus is the speaker. Using a definite noun phrase anaphor “the poodle,” the speaker turns his attention to the that dog, and it becomes the discourse focus and the actor focus. Using another definite noun phrase anaphor “the cocker,” the speaker changes the discourse focus to the cocker. The speaker's concluding sentence expands the discourse focus back to both dogs with the use of “them.”

In Sidner 1979 a machine for choosing and updating the discourse focus and actor focus is given. The machine chooses items in the discourse to serve as “expected” foci after the first sentence of a discourse and uses the interpretation of pronouns as well as

non-pronoun anaphors to determine whether its expected choices were correct. The machine updates the discourse and actor focus after each sentence and changes foci when the anaphoric expressions no longer co-specify with the item in focus. This simple machine tracks the foci from the hearer's point of view and shows how the hearer can recognize foci and changes of foci. In the remaining discussion, it is assumed that the focus machine can determine a discourse and an actor focus. Our attention is directed at the following concerns. How can these two foci be used to interpret discourse pronouns? What rules are used by the pronoun interpreter? How do these rules make use of the constraints on the theory of pronoun interpretation discussed earlier in this paper?

The focusing theory of pronoun interpretation can be outlined as follows. At any time in the discourse after the first sentence, there are foci that are the prime candidates for co-specifying with a pronoun. Also available is a list of alternate candidates, called *potential foci*, for each focus. Unless the focus candidate is ruled out by certain criteria (to be discussed below), the pronoun interpreter uses it to determine the specification of the pronoun. Most of the discussion in this paper concerns the criteria that must be encoded in the rules for choosing among the foci as well as the criteria for rejecting a focus in favor of one of the potential foci.

The proposal made here contains two implicit processing assumptions, (1) serial processing, and (2) end-of-sentence processing. By “serial processing”, I mean that the interpreter checks a focus as a candidate in interpreting the pronoun, and then if that focus is unacceptable, checks alternate candidates in turn. By “end-of-sentence processing”, I mean that pronouns are not interpreted until the entire sentence has been syntactically and semantically interpreted. Both of these criteria could be given up without undermining the focusing theory. One could envision processing in parallel by checking the foci and alternates and then determining the pronoun's specification from an ordering of all those candidates that meet the criteria of choice. To interpret pronouns before the end of the sentence, the interpreter could decide using available information and then review its choice as more of the sentence is processed. These two implicit processing assumptions have been made because they simplify the account of focus and because they reflect an implemented version of one focus system. Further research will indicate whether these assumptions are too strong – if so, the focusing theory may be re-tested under parallel and mid-sentence interpretation assumptions.

4. Using Focus for Pronoun Interpretation Rules

The pronouns in the example discourse in the preceding section can be interpreted with a simple rule:

R1: Choose the discourse focus as the co-specifier of the pronoun. But if the pronoun appears in the agent position of the verb case frame structure, then choose the actor focus.

R1 is a naive formulation and fails to predict accurately over a wide range of phenomena. It can be revised to include criteria from syntax, semantics and pragmatics, as well as criteria about discourse characteristics. In the next several pages, I present criteria that make use of all four sources of knowledge and incorporate them into R1 to form a series of pi-rules; these rules make use of the discourse and actor foci, the potential foci, and the processing of the focus machine.

Syntax, semantics and pragmatics can be included in an initial way by the modification R2:

R2: If a pronoun appears in a verb case frame relation other than agent, choose the discourse focus as co-specifier unless any of the syntactic, semantic and pragmatic knowledge constraints rule out the choice. If the pronoun appears in agent position, choose the actor focus as co-specifier in the same way.

Ruling out a co-specifier on the basis of syntactic and semantic constraints is accomplished by computing the various syntactic relationships and restrictions (such as Lasnik's disjoint reference rules) and by use of semantic selectional restrictions on case frame categories. How one uses general knowledge constraints is not so obvious; one wants to avoid the problems of control on inference mentioned previously.

The focus rule above, although only a skeleton of the full rules needed, contains the basis for control of inferences. Under the method of rule R2, inference is no longer used to find a binding for the pronoun acting as a variable by forward and backward chaining. Instead inference is a constraint-checking process. The context up to the sentence under interpretation is integrated into memory; when a co-specifier is chosen for the pronoun, inferring serves only to find a contradiction in the database. The database will also include an interpreted form of the sentence with the pronoun, with the pronoun replaced by its specification as predicted from the phrase in focus. Only when a contradiction is found, is the suggested specification given up in favor of another choice. The foci of the discourse can be additionally helpful in database inference because only that part of the database that concerns what is in focus needs to be explored for contradictions. Some research on truth maintenance systems [Doyle 1978, McAllester 1978] has experimented with constraint-checking and developed algorithms for efficiently finding and undoing contradictions; however,

no one has considered how to explore only certain "sub-sections" of a database for contradictions.

On this account of inferring and pronoun interpretation, the inference machine must still infer propositions and reach contradictions. However, this method reduces the search needed to make sense of the new sentence because specific choices for the pronouns are given; when a contradiction is reached, a new choice is made rather than either blindly binding the pronoun to some other object or blindly searching for some other proposition to change, so that the contradiction is eliminated.

When a suggested co-specifier for a pronoun must be given up, the pronoun interpreter must use the potential foci as possible co-specifiers. However, the interpreter's actual choices require consideration of several matters, and before any rules can be stated, several motivating cases must be presented.

D3-1 Alfred and Zohar liked to play baseball.

2 They played it everyday after school before dinner.

3 After their game, Alfred and Zohar had ice cream cones.

4 The boys thought they tasted really good.

5 Alfred always had the vanilla super scooper,

6 while Zohar tried the flavor of the day cone.

7 After the cones had been eaten,

8 the boys went home to study.

In D3, Alfred and Zohar are the initial actor focus while baseball is the initial discourse focus. D3-2 contains two pronouns, "they" and "it" which, according to R2, co-specify respectively with Alfred and Zohar, and baseball. D3-3 uses "their," which co-specifies with Alfred and Zohar, but is not accounted for by rule R2. Furthermore "they" in D3-4 does not co-specify with baseball but with ice cream cones. What has happened?

Suppose in addition to the ongoing actor and discourse focus, the hearer can consider temporarily any new entities mentioned in the last sentence that the hearer has heard. These entities, the *potential foci*, are dropped out of processing use if the speaker fails to mention them a second time; such alternatives are elements the speaker may want to say more about, but their importance to the speaker cannot be determined from one sentence. Thus in D3-2, foci include the event of playing, and the times "everyday", "before dinner," and "after school". None of these potentials survives long into the discourse as D3-3 fails to pick up on any of them. D3-3 in turn introduces some other potential foci, and one of them, ice cream cones, is discussed in D3-4; notice also that D3-4 fails to mention baseball in any way.

R2 must be extended so that it uses the potential foci whenever criteria from syntax, semantics or pragmatic knowledge rule out the current actor or discourse focus. Thus since "they" in D3-4 cannot co-specify with baseball (on both syntactic⁵ and semantic⁶ grounds), a potential focus is chosen. Potential foci are ordered according to their thematic relation⁷ in a sentence, with semantic case objects first, and agents last. The first potential focus that meets all the constraints is chosen as the co-specifier; in D3, ice cream cones is the first acceptable potential focus that meets all the necessary constraints as a co-specifier for "they."

Use of this modified pronoun rule follows hand-in-hand with the focus machine discussed previously. The machine updates its discourse model after each sentence by tracking, among other things, pronoun use. When a pronoun is used to co-specify with a phrase that is a potential focus, and no phrase co-specifies the current focus, either the discourse focus or actor focus moves to that potential focus. Which of the two foci moves depends on whether the pronoun fills the agent case in the verb frame, and in the case of multiple agent cases, whether the ongoing actor focus is re-mentioned. For example, after D3-4, the discourse focus changes to ice cream cones because "they" co-specifies with the ice cream cones; the boys remain the actor focus, since "they" is not an agent case for "taste" and since the boys were already the actor focus.

So far pi-rules take into account the movement of focus and constraints on syntax, semantics and pragmatics. However, some of the semantic and pragmatic criteria must be expanded beyond the simple case frame semantics and representation of noun phrases given earlier.

As the reader may have concluded, the interaction between actor and discourse foci is a complex one; to supplement the case frame semantics, some functional notion of theme (in the sense of functionalists such as Halliday 1967) is needed. For example, the proper co-specifier for the pronoun in D4-3 cannot be chosen with R2 and only simple case frame relations.

D4-1 I haven't seen Jeff for several days.

2 Carl thinks he's studying for his exams.

3 But I think *he* went to the Cape with Linda.

Although Carl is the actor focus after D4-2, and "he" in D4-3 is the actor case in the embedded sentence, the proper choice for the co-specifier is Jeff. Howev-

⁵ since "they" is a plural pronoun and baseball is singular,

⁶ since the discourse items filling the object case of taste should be tastable items,

⁷ This type of ordering is motivated by the determination of an initial focus. See Sidner 1979 for details.

er, with simple case frame semantics, the interpreter must consider "Carl" as the co-specifier and then fail to rule it out (because no syntactic, semantic or knowledge constraints can eliminate it).

I have designed the pi-rules to use a semantics that marks verbs like "thinking" or "talking" with who the thinking or talking was about. D4-3 would indicate that the speaker thinks about X, that X went to the Cape with Linda. With this semantics, the interpreter follows a rule to prefer the discourse focus as the co-specifier of any pronoun filling the theme position in such a verb. Since the actor focus is associated with the agent case, use of such a theme-based rule is consistent with the framework. With such a rule, Jeff is chosen as the co-specifier of "he," not because Carl couldn't have gone to the Cape with Linda, but because the speaker is talking about Jeff and his thoughts about Jeff. In essence this approach takes the discourse focus as primary, the discourse focus being what the speaker is talking about so far, while the actor focus is the locus of information about actions in the discourse.

The theme-focus rule does not indicate what to do when the discourse focus and actor focus are both animate, have the same gender, number and person, and are both established during the same sentence of the discourse. People sometimes have difficulty choosing interpretations in such circumstances. In D5-2a below, "he" co-specifies with "John" (the actor focus) but if D5-2b followed D5-1, "he" may co-specify with either John or Mike (the discourse focus).

D5-1 John called up Mike yesterday.

2 a He wanted to discuss his physics homework.

b He was studying for his driver's test.

In these cases, native speakers report that the co-specifier for the pronoun is ambiguous. If the pronoun fills an agent case, the actor is preferred, but this preference is not a strong one. It appears that in such cases the ambiguity may not be easily resolved unless additional information about the two foci is known that stipulates that the sentence is true of only one.

Potential foci can be sub-categorized as either potential discourse foci or as potential actor foci. A *potential actor* is a noun phrase which specifies a database element marked as animate and which does not occur in agent position. In most sentences, the noun phrase in agent position contains a descriptive word or name in the head, and specifies a database element; it becomes the actor focus. But when the noun phrase in agent position is a pronoun, it may co-specify with either the actor focus (if one exists), or a potential actor. Ambiguities occur when an actor and one potential actor are both present in a previous sentence and when the discourse focus is a non-actor entity. An example is given below.

Consider the examples where s5 follows each of s6, s7, s8 and s9.

- s5 He knows a lot about high energy physics.
- s6 Prof. Darby will tell Monty about the neutron experiment.
- s7 Prof. Darby will lecture Monty on the neutron experiment.
- s8 Prof. Darby will help Monty with the neutron experiment.
- s9 Prof. Darby will teach Monty about the neutron experiment.

Some native speakers find all of these sentence pairs ambiguous, while some native speakers find only the pair s6 followed by s5 ambiguous. These examples are surprisingly similar to D5. How do some speakers decide that “he” co-specifies with Monty or Prof. Darby? It appears that they make a comparison and choose between the actor focus and the potential actor on the basis of evidence for their preferred interpretation. When that evidence is not forthcoming, informants are confused. Such a behavior suggests that the inference mechanism should be able to judge preferences between a given actor and one potential actor.

A computational system that makes such judgments must have a very rich knowledge base (e.g., to know that Monty is a male name, and that professors may be male) and must be able to infer which actor is preferable from that base. A computational framework for carrying out such subtle judgments is still beyond the state of the art, although Marcus 1980 has proposed a semantic choice mechanism that must also weigh evidence for prepositional phrase attachment; his parser, when attaching prepositional phrases, asks the semantic processor about its preferences for where to attach the phrases.

In summary, the interpreter can use the following condition for these ambiguities.

POTENTIAL ACTOR AMBIGUITY CONDITION: Whenever a pronoun may co-specify the actor focus, and a single potential actor exists, expect a possible ambiguity. To resolve,

1. If there is evidence supporting the actor focus as the co-specifier, but not the potential actor, then the actor focus is the co-specifier.
2. When evidence supports the potential actor but not the actor focus, select the potential actor as the co-specifier.
3. However, if there is evidence for both, select the actor focus but indicate ambiguity.

A summary of a full set of pronoun rules is given in the appendix to this paper. These rules represent what can be said about pronoun interpretation in the absence of any additional information in knowledge representation beyond that suggested in the discussion of

co-specification. To interpret certain pronouns, such as those where a co-specifying phrase does not precede the pronoun in the discourse, as in D1, we must consider how knowledge is structured and represented.

5. Focus and Knowledge Representation

Focusing theory with the syntactic and semantic representations discussed initially does not include an account of the representation of sentence pairs such as D6 where the sentence is ambiguous due to scope of quantifiers.

D6-1 Wendy showed each girl Bruce knows a cat.

2 a She had found it at Farmer John's.

b They were all from the same litter.

At present the focusing rules predict only that “a cat” is the expected discourse focus, but they say nothing about the scope relations for representing the semantics of the phrase as part of D6-1. Without scope relations, D6-2b is ruled out by the number agreement criteria when in fact most speakers say that “they” co-specifies with the set of cats that may be evoked in D6-1. Webber 1978 not only argues in detail for representing scope relations in the mechanisms for interpreting anaphors, but also gives rules for determining scope in a representation of phrases that may be co-specifiers.

On Webber's analysis, D6-1 will have two representations for “a cat” and for “each girl Bruce knows.” The one for “a cat” can be paraphrased as:

1. R1: the cat associated with D6-1 such that Wendy showed it to each girl Bruce knows.
2. R2: the set of cats, the members of which are associated with D6-1 such that Wendy showed (one of) them to each girl Bruce knows.

Webber has suggested *what* the representation must be, and what remains to be determined is *when* it is used. Let us suppose that D6-1 is ambiguous,⁸ and no syntactic or semantic processing of it adjudicates the two readings R1 and R2 for “a cat.” When the pronoun interpreter seeks to co-specify “it” in D6-2a with the discourse focus, both readings will be available. The set reading R2 may be eliminated immediately because of constraints that the co-specifier represent a singleton, so R1 is left. In contrast, for D6-2b the reading R1 will be ruled out since “they” requires a plural co-specifier. This account both explains how the pronouns may be understood, and also is consistent with Van Lehn's findings [1978] that people do

⁸ A question, often raised by Martin (see, for example, Martin 1979), is whether a sentence that is ambiguous among several readings must be represented by several different structures, one for each reading. He offers a semantic representation that preserves ambiguity until some processor demands a refinement. Whether this approach or an alternative representation containing both readings is best is still an open question.

not normally disambiguate certain ambiguous quantifier scopes during sentence understanding for sentences such as D6-1. It is compatible with Van Lehn's findings because alternative interpretations of scope are not considered until additional discourse material beyond the single sentence is presented.

Webber's rules also allow the phrase "each girl Bruce knows" to be interpreted as a prototype, with the result that "a cat" is also. Most speakers find this reading very odd, and one can ask how such a reading is eliminated. Possibly it is never generated, and Webber's rules need modification. Alternatively it is produced, and then is ruled out, either in understanding D6-1 because showing prototypic cats is bizarre, or because in interpreting "it" for D6-2a, the inference mechanism balks at Wendy having found a prototypic cat at Farmer Brown's. I think that such readings never occur in the first place, simply because most speakers, who do not report this reading themselves, have difficulty understanding the reading when they are told about it.

Another case of semantic ambiguity, similar to the one in D6, is illustrated in D7.

D7-1 Sally wanted to buy a vegomatic.

2 She had seen it advertised on TV.

"A vegomatic" may be interpreted specifically, to mean that there is one particular vegomatic or non-specifically, to mean that it is one of the many vegomatics.⁹ The focusing rules do not distinguish between the two after D7-1 because, like D6-1, D7-1 is ambiguous, and neither interpretation can be chosen with certainty. When "it" is resolved for co-specification in D7-2, the inference mechanism must decide that Sally does not want the very one used in the advertisement on TV. "A vegomatic," whether understood as specific or non-specific, specifies a different representation than "it" does. Therefore the pronoun's specification as a specific vegomatic must be generated from the ambiguous use.

Suppose, for a moment, that D7-1 is interpreted so that a representation that maintains ambiguity is available. When the pronoun interpreter processes a subsequent sentence with a pronoun, it need only rule out readings if the inference machine discounts as contradictory one of the readings (Sally didn't want to buy the very one she saw on TV). If no reading is ruled out, the co-specifier would remain ambiguous, so that both the indefinite phrase and the pronoun would have

⁹ The terms "non-specific" and "specific" are traditional semantic expressions that bear no relation to "specify" and "specification." A non-specific reading of "a dog" would be interpreted to produce a representation of an instance of the prototypic dog; what is represented is a dog which has the characteristics of the prototypic dog - e.g., an animal with four legs, a tail, medium size, brown, friendly, barks, and the like.

ambiguous co-specifications. As the next example shows, there is some evidence for this behavior.

Consider the case shown in D8.

D8-1 Sally bought a vegomatic that had a broken cutting blade.

2 She had seen it advertised on TV.

"A vegomatic that had a broken cutting blade" in the context of D8-1 usually means some particular vegomatic that Sally bought. However, "it" is ambiguous among the vegomatic Sally bought (the broken one), some vegomatic (possibly not broken), and a vegomatic that is an instance of prototypic vegomatic. Thus "it" is three ways ambiguous.

To understand D8-2, the pronoun interpreter does not distinguish among the three readings, since it accesses the one provided by the specification in D8-1, which is the specific reading. To find the specification of "it," the inference mechanism must discover that it is slightly odd (1) for Sally to have seen the vegomatic with a broken blade which she bought being advertised on TV, and (2) for Sally to see any broken vegomatics on TV, and (3) for Sally to have seen the very one she bought on TV. Then if no other choices for co-specification are available, the specification of vegomatic from D8-1 must be used to generate an appropriate specification for "it". Since only unbroken ones not bought by Sally are appropriate, the specification of "it" must be generated using only part of the phrase from D8-1.

This example seems problematic because it places much weight on the inference machine to decide that the reading is odd. However, this is likely to be just where the weight of the decision ought to be; many native speakers find D8 slightly bizarre because their first reading is that Sally had seen the vegomatic with the broken blade advertised on TV. In fact, it appears that when a specific indefinite noun phrase such as "a vegomatic" is introduced, and the speaker wants to turn attention to the non-specific reading, a plural pronoun is used as shown below:

s10 She had seen them advertised on TV.

The plural non-specific reading as in s10 is mentioned in the pronoun rules found in the appendix, but the generation of a specification as in D8-2 is not.

Examples such as D8 are perplexing for another reason; they are examples of what I call, following Fahlman 1977, the "copy phenomenon." The ambiguity centers around the fact that there can be many copies of an abstract prototype. Automobiles, computer programs, airplane flights and money are other common cases of entities that exhibit the copy phenomenon. In D9, the interpretation of "it" depends on whether the speaker is referring to a particular flight or the normal Sunday flight, a copy of which

occurred on “this Sunday” because of the use of “usually”.

D9-1 TWA 384 was so bumpy this Sunday I almost got sick.

2 It usually is a very smooth flight.

Note that the “it” *cannot* co-specify with the particular flight on “this Sunday”. However, it is possible that the speaker used “TWA 384” to refer to a particular flight; if this is so, the speaker mixed the specific and non-specific interpretations for the co-specifier of “it”. Just as in D8, a specification for the pronoun must be generated — in this case, a non-specific one from the specific reading.

Another characteristic anaphora is the bound variable case described by Partee [1978, 1972]. In D10 below, “him” co-specifies with Archibald, while if the reflexive “himself” were used, it would involve a variable bound to the quantifier from “every man.”

D10-1 Archibald sat down on the floor.

2 Every man put a screen in front of him.

In linguistic theory, bound variables are assumed to be represented in sentence semantics; when used in conjunction with syntactic disjoint reference rules, pronouns within the scope of the quantifier can be distinguished from non-scoped ones. Since the pronoun interpreter takes account of these conditions, it can easily choose a proper co-specifier for “him” in D10 in terms of the focus, but for “himself” it will recognize the bound relation to “every man.” It is crucial to these cases that the representation of the interpretation of a sentence includes scope of quantification, especially when the scope is unambiguous.

6. Focus Restrictions on Co-specification

There are other restrictions on co-specification that result from the processing of the focus machine. The focus machine includes a stack on which old foci are stored when the focus changes. In addition to co-specifications with current actor and discourse foci, a speaker may use a pronoun to co-specify a discourse element that was once in focus but is no longer; Grosz 1977 described and illustrated this behavior for anaphoric noun phrases in task-oriented dialogues.

A pi-rule that selects candidates from the focus stack can capture this behavior. However, the rule must be further constrained. An anaphor that is intended to co-specify with a stacked focus must not be acceptable as a co-specifier with either the foci or potential foci. An example from a literary text¹⁰ illustrates how.

Was that old lady evil, the one Saul and I had seen sitting on the porch? I had dreamed about her. When the trolley car took me and Saul past her house again this morning, she was

gone. Evil, it had a queer sound to it in English.

[Here the narrative moves on to an incident in a school classroom. A discussion between the speaker and a male teacher ensues for five paragraphs. Then the following paragraph begins:]

She had worn an old brown coat and a green scarf over her head.

In this example, “she” co-specifies with the old lady discussed previously. If Potok had told of a discussion between the speaker and a female teacher, it would no longer be possible to tell that “she” was co-specifying with the old woman. While interpreting the reading of “she” as teacher might be a bit surprising because the teacher’s clothing was not relevant to the previous conversation, no inference can rule out that reading because teachers may wear old brown coats. Hence the pronoun interpreter must reflect these facts, and it does so with *the stacked focus constraint*.

The stacked focus constraint is not stated directly within the pi-rules. Instead it is implicit in their function. The condition is as follows: A pronoun cannot be used to co-specify with a stacked focus when a current focus is an acceptable co-specification since that current focus will be taken as the interpretation and the stacked focus will never come into consideration. The stacked focus constraint is a consequence of the movement of focus in the focus machinery.

The stacked focus constraint, however, may be overridden. Grosz (see Deutsch 1975, 1974) has identified several examples, involving pumps and bolting operations, that show that a pronoun may co-specify with the stacked focus even when intervening material contains possible co-specifiers. Some complicated set of inferences about what can and cannot be bolted to what, or what can and cannot be loosened might be able to rule out the intervening possible co-specifications. However there appears to be too little delay in understanding for people who read her excerpts (there is no means of determining whether the original speaker and hearer experienced any delay in understanding) to suggest that they are ruling out multiple possibilities in interpreting the pronouns.¹¹ Hence one may conclude that the speaker and hearer are taking advantage of their knowledge of the task to provide a discourse context in which the focusing machinery can be applied.

How many such discourse interpretation mechanisms exist? While this paper does not address this question directly, some speculation is possible on the

¹⁰ *In the Beginning* by Chaim Potok, page 212, chapter 4, Fawcett Publications, Inc., Conn., 1975.

¹¹ This informal evidence needs to be tested out in a psychological laboratory. I have not done so, but the results of such experimentation would be revealing.

basis of research that has been reported elsewhere (see Sidner 1979, Robinson 1981). In general, it appears that discourses permitting violations of the stacked focus constraint must contain an implicit structure of task completion that guides the hearer to a context in which the foci for that context may be re-used. Without this structure the hearer cannot decide that the speaker intended for the pronoun to co-specify with something other than the representation of the object currently under discussion.

7. Pronouns Which Have No Co-specifiers

The previous discussion has assumed that a pronoun is always preceded by a co-specifying phrase. However, this is not always the case, and a complete theory of pronoun interpretation must address cases where the co-specifying phrase appears *after* the pronoun, and where *no co-specifier exists*, but the discussion implies a specification for the pronoun.

Pronouns used with their co-specifiers appearing after the occurrence of the pronoun have been called *backward anaphora* in the linguistic literature; I refer to them as *forward co-specifiers*. Two such examples, s11 and s12, are given below.

s11 If *he* comes before the show, give *John* these tickets and send *him* to the theatre.

s12 Near *him*, *Dan* saw a snake.

In general the pronoun co-specifies with some noun phrase interpretation, but the phrase is placed forward in the discourse. The types of sentences in which this behavior can occur are limited. In general it seems to be permitted for fronted sentential prepositional phrases (as in s12), subordinate clauses fronted on another sentence (as in s11), and for sentences containing co-ordinating conjunctions.¹² However, it is excluded, as far as I can tell, in the following cases:¹³

s13 * I heard about *her* job from *Mary*.

s14 * I spoke about *him* with *John's* wife.

Extensive research in linguistics on forward co-specifiers (Solan contains a good review) gives reliable evidence that structural constraints govern it; in particular, syntactic rules can be stated that determine when forward co-specifiers are not permitted. The most recent formulation, by Solan, called the backward anaphora restriction, fails on the cases below,

¹² Solan 1978 cites the example:

s15 Penelope grabbed his cane and beat Peter with it.

¹³ There is some disagreement on s13; Solan claims it is acceptable. All the native speakers I have asked said they interpreted the pronoun as "referring to Mary" only because there were no other usual choices and because you could hear about a person's job from that person. Since all the people I asked informally told me it was a strange sentence for them, I am assuming that it is deviant.

but some working modification of it may yet be forthcoming.

s16 In *her* room *Mary* saw a ghost.

s17 I heard about *Mary's* job from *her*.

Whatever the best formulation of the syntactic rules for forward co-specifiers, they are preferable only in initial sentences of a discourse. For example, when s11 occurs in mid-discourse, if a speaker has been talking about Henry, and just begun mention of Charles, native speakers will interpret "he" as co-specifying with Henry, or Charles (if Henry can be ruled out on basis of some special pragmatic knowledge). The pi-rules using focus behave in exactly this way. They will permit a forward co-specifier only if Henry *and* Charles can be ruled out by syntactic, semantic and world knowledge criteria. Were s11 to occur when Lydia was in focus, the forward co-specifier would be possible, and pi-rules mirror this behavior.

When faced with a pronoun that has no preceding co-specifier and is not a forward co-specifier, the pronoun interpreter relies on a condition that is called the *missing co-specifier condition*. In the remainder of this section, I describe that condition. The pi-rules include a rule for recognizing that a pronoun is missing its co-specifier, but they do not offer an interpretation for such pronouns.

In each of the cases below, the pronoun specifies some cognitive element that is related to one of the entities mentioned in either the previous sentence or the same sentence. These examples¹⁴ are different from the non-specific prototypic readings of pronouns discussed earlier because no phrase which can be used as a generator exists.

D11-1 I saw Mr. Smith the other day; you know, *she* died last year.

2 John is an orphan. He misses *them* very much.

3 Pro-Castro people don't believe *he* is a monster.

4 I went to a concert last night. *They* played Beethoven's ninth.

5 I want to meet with Bruce next week. Please arrange *it* for us.

6 I used to be quite a tennis player. Now when I get together with the young guys to play, I can hardly get *it* over the net.

With the exception of D11-1, most speaker-hearers are able to say which is the intended specification of the pronoun in the cases above. D11-1 can be understood

¹⁴ The examples given here are from several sources; the first three are from Postal 1969, the fourth from Chafe 1975, the fifth from dialogues collected for the PAL system [Sidner 1979], and the last was spoken by a lecturer at a presentation I attended.

if the hearer is informed that Mr. Smith had a wife. However, some of these, especially 1 and 2, are so odd for most hearers that at first they fail to comprehend the pronoun. Hearers are divided on the acceptability of 3, and most hearers find 4 and 5 acceptable. Furthermore, such examples, as far as I can tell, do not occur naturally in written samples.

This paper does not give an account of just how such cases are understood. However, the focusing approach provides some basic structure that may be useful in generating an explanation for such situations. In all the multi-sentence cases, the pronoun specifies something which is closely associated with the focus. What is problematic is the fact that some speakers seem unable to understand a pronoun which specifies a database element that, while related to the item in focus, represents something that no longer exists, such as John's parents in light of John's orphanhood. Whatever the manner in which hearers recover specifications for such pronouns, some principles are needed which govern why some uses of pronouns are acceptable and others are not.

8. The Problem of Parallelism

The pi-rules give incorrect predictions for certain uses of pronouns, uses that are difficult to define. Intuitively, they may be characterized as occurring when there is a *parallel structure* between sentences of a discourse.¹⁵ In many of these cases the pi-rules predict the wrong co-specifier. To understand what is meant by parallel structure, two simple cases, one in which the pi-rules do predict correctly, and another in which they fail, are discussed. In D12, the pronoun co-specifies with the mud pack, as the pi-rules would predict. The parallelism of these sentences is reflected in the semantics of "put on" and "pull off" as well as in the similarity of the syntactic structure of the two sentences, each being in imperative mood.

D12-1 Put the mud pack on your face.

2 After 5 minutes, pull it off.

The pi-rules predict the proper co-specifier in D12 because the thematic relations of the verb follow the similarity of structure. However, in D13 below, the pronoun "it" co-specifies with rose and not with the green Whitierleaf (the pi-rule choice). The initial focus after the first sentence is Whitierleaf, but the parallel syntactic structure of the sentences seems to govern a different choice of co-specifier. To summarize, between similarity of structure and the pi-rules, similarity is preferred as a means of choosing a co-

specifier, so when each gives a different prediction, similarity of structure must be used.

D13-1 The green Whitierleaf is most commonly found near the wild rose.

2 The wild violet is found near it too.

At first glance it appears that the pi-rules could be "fixed" by simply observing that the initial focus is wrong and that a potential focus should be chosen. However, no such option is available, for such a "fix" requires that the inference machine reject the initial focus. To do so, the inference mechanism needs some knowledge about the world that indicates the unacceptability. For D13 no such knowledge could possibly be forthcoming since all the flora involved are found near one another. There is no knowledge to the effect that violets are found near wild roses and not near Whitierleaves.¹⁶

Another example of parallel structures is shown in D14. The parallel structures again are reflected in the similarity of the syntactic forms as well as the semantics of "most" and "mine". After D14-1, the initial focus is the car radiator that is associated with each car of "most cars." Using the focusing rules, the pronoun interpreter will take "it" to co-specify with that radiator. But this prediction is incorrect; "it" co-specifies with the radiator of the speaker's car.¹⁷

D14-1 On most cars the radiator has a free bolt hook.

2 But on mine, it has a floating bolt hook.

The use of "it" here is similar to the instance of a prototype for "it" in D8. The two discourse examples differ because D14-2 has an underlying semantic form that parallels D14-1. D14-1 specifies a universal set of cars and says something about one of the parts for those cars; D14-2 specifies a set of one thing, the speaker's car, and says something about a part of it; the speaker's car is related to the universal car by instantiation. Thus "it" in D14-2 is not pointing to some instance of the prototypic radiator; it co-specifies with the radiator of the speaker's car, but "it" is related intensionally to the "radiator" in D14-1. The similarity in the underlying semantics of D14-1 and D14-2 must be used in interpreting the pronoun uses.

One might wish to construct some special-purpose mechanism that looks for similarities in structure between two sentences. This method is doomed for two reasons. First, parallelism exists in many aspects of language, and it happens at arbitrary levels of struc-

¹⁵ Hobbs 1979 defines a parallel relation between sentences S0 and S1 of a discourse as occurring when propositions P0 and P1, which follow from S0 and S1 respectively, have identical predicates, and arguments that are similar (p. 76). This concept of a parallel relation seems related to my intuitive characterization, but as defined does not capture the cases I discuss.

¹⁶ In certain cases a special audience may have different responses to the parallelism above. For example, botanists who know what flowers are near others might behave differently. But even special audiences must sometimes use general techniques. Such is the case in the D13 example, because Whitierleaves exist only in my imagination.

¹⁷ I thank R.C. Moore for suggesting this example.

ture rather than just syntactically. Second, at any given level, the problem of recognition of parallelism has plagued computational models of language since such models were first suggested. For example, parsing of English sentences containing conjunction is as yet an unsolved problem. Methods tried, such as those of Woods 1973 in LUNAR, fail because of over-generalization. Proper computational recognition of parallelism is still beyond the state of the art.

The fact that interpretation of parallelism has failed for other aspects of computational models of language only indicates that the problem is a deep one. The example below is especially surprising because it demonstrates the use of parallel constructions that may be found between whole paragraphs in a discourse. The interpretation of "the schedule," used anaphorically in D15-7, is not ambiguous between "SOL," which is a kind of schedule, and the transmission schedule; readers recognize that the anaphor is unambiguous presumably because the process described in lines 5-7 parallels the one in lines 2-4. Whatever the proper account of parallelism for pronominal anaphora, it must also be generalizable to account for this kind of example.¹⁸

- D15-1 The SOL is searched for an entry for the subscriber.
- 2 If one is found, the subscriber's relative transmission time is computed according to formula-1.
 - 3 The subscriber's clock transmission time is computed according to formula-2.
 - 4 When the transmission time has been computed, it is inserted as the primary entry in a transmission schedule.
 - 5 For each RATS entry, the RATS's relative transmission time is computed according to formula-1,
 - 6 and the RATS's clock transmission time is computed according to formula-2.
 - 7 The RATS transmission times are entered into the schedule.

One possible consequence of these observations could be that the focus mechanism should be abandoned in favor of some as yet unspecified mechanism that is able to determine parallel relations among sets of sentences in a discourse. However, methods for interpreting pronouns from parallel sentences and paragraphs offer no constructive way of interpreting the pronouns in most of the examples presented in this paper. Many cases of co-specification occur where there is no similarity of structure other than the com-

mon subject-verb-object pattern typical of English sentences. Since what is being talked about appears in many constituent positions in sentences of a discourse, the S-V-O pattern is too gross a level to specify similarity. Hence, while parallelism is needed to deal with a certain set of cases for which the pi-rules predict incorrectly, the pi-rules are effective for many other cases of co-specification where parallelism would not be helpful. One may conclude that focus mechanisms account for one aspect of pragmatic anaphora, and that some different mechanism is needed to encode similarities in structure that sometimes occur in discourse. This paper does not give an account of such a mechanism. Rather, the examples above provide some additional observations about the nature of parallelism in interpreting pronouns in natural languages.

9. Conclusions

Two claims have been substantiated in this paper. First, I have shown that focusing is compatible with linguistic rules for disjoint reference, semantic selectional restrictions, and with representations delimiting quantifier scope. Furthermore, these sources of information are necessary for the pronoun interpretation rules.

Second, I have demonstrated that focusing helps control the inference process needed to interpret pronouns. The pronoun interpreter predicts a co-specifier and then asks the inference machine for confirmation; when the machine finds a contradiction rather than consistency, the focusing process produces a new candidate co-specifier. In effect, the focusing process and the inference machine collaborate on pronoun interpretation. In previous AI natural language systems, interpretation resulted from binding of free variables when making inference; the inference process could be characterized as one of proving a consequent from a set of premises. However, because the consequent of an inference rule contained free variables, many inferences had to be drawn and then "undone" during the search for a complete chain of inferences. The focusing approach eliminates this kind of blind binding and unbinding as well as shortens the inference chain search.

A part of this paper has described and illustrated how the pronoun interpretation rules predict the co-specifiers of pronouns with both the actor and the discourse focus, and indicated that both actor and discourse foci are necessary. The rules show how many constraints — syntactic, semantic, and pragmatic¹⁹ — affect the choice of specification for pronouns; they also show how many kinds of informa-

¹⁸ For easier understanding, the reader must know that the SOL has both subscriber and RATS entries. This example comes from Balzer et al. 1977, "Information in Program Specification," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, p. 394.

¹⁹ Phonological constraints have not been discussed in this paper. However, rules for contrastive stress and the like could be incorporated in the way that the other three classes are.

tion about the world of the speaker and hearer play a part in distinguishing the co-specifications of pronouns. By means of the discourse and actor foci, the pi-rules differentially apply the constraints; the predictions which result are tested in the database representation of the speaker and hearer's world. Focusing captures the effects of foregrounding (cf. Chafe 1975) since focusing accounts for the co-specifications of pronouns by means of the foci, and focus movement indicates how new entities may be foregrounded and pronominalized.

In this paper I have demonstrated the types of anaphoric uses that cannot be explained by the focusing theory. These uses, the focus "popping" cases described by Grosz, and the parallelism cases, illustrate that in anaphor interpretation, other kinds of processes that are computationally realizable and controllable are needed. Furthermore, an adequate syntactic theory of forward co-specifiers, which remains to be discovered, must be incorporated in the syntactic criteria for the pronoun interpretation rules.

This paper further specifies the nature of focusing as it relates to a theory of pronoun interpretation. A focus-based theory with stipulations for syntax, semantics and inferential knowledge, provides a predictive and explanatory theory of pronoun interpretation. The theory is predictive because it stipulates legal and illegal pronoun uses as well as their interpretations; the limits of its predictive power have been demonstrated as well. The theory is explanatory because several observations about anaphora in English, all crucial to focusing theory, can account for each of the rules: that pronouns are signals of what is being discussed, that whenever the discussion changes, pronouns must signal the new element of discussion without confusion, and that hearers make use of syntactic, semantic, and pragmatic knowledge in a controlled way in understanding anaphors. A theory of anaphor interpretation different from the focusing theory needs either to use these observations in an account of the facts or to demonstrate why these observations are not relevant to the interpretation process.

Appendix. Pronoun Interpretation Rules

The rules in this appendix assume the availability of an actor focus, discourse focus, potential actor and potential discourse foci. The pronoun interpreter applies these rules in the order given.

Agent position pronouns

1. When a sentence in which a pronoun occurs is the second sentence of a new discourse, the recency rule may be applied. *Recency rule:* When a pronoun is in subject position and is the initial phrase in a sentence, and if a member of the potential (discourse or actor) foci occurs as the last phrase in the previous sentence, test the pronoun for co-specifying with that potential focus.

2. *Theme rule:* When the pronoun occurs in an embedded sentence, if the embedded sentence is marked as having a theme that is either the discourse or actor focus, test the focus in that theme position as the co-specifier of the pronoun.
3. *Potential actor ambiguity condition:* Whenever a pronoun may co-specify the actor focus, and a single potential actor exists, expect a possible ambiguity. To resolve, (1) if there is evidence supporting the actor focus as the co-specifier, but not the potential actor, then the actor focus is the co-specifier. (2) When evidence supports the potential actor but not the actor focus, choose the potential actor as the co-specifier. (3) However, if there is evidence for both, choose the actor focus but indicate ambiguity.
4. *Pronominalized actor focus rule:* When the actor focus was last mentioned with a pronoun, if the agent position pronoun is of the same gender and person as the actor focus, it must co-specify with the actor focus. If it does not, a potential actor focus may be chosen, but the pronoun use is odd.
5. *Plural rule:* If the pronoun in question is plural, while the actor focus is singular, test items in the following list for co-specification: a non-specific reading of the focus (only for foci that specify non-human entities), actor focus and potential actor foci, all the potential actor foci together, the discourse focus, and the potential discourse foci.
6. *Basic rule:* Test the actor foci as a co-specifier with a pronoun in agent position followed by potential actor foci. If these fail, check the discourse focus, potential discourse foci and actor focus stack.
7. *Closure rule:* Should all other rules fail, if the pronoun occurs in an introductory clause, expect a forward co-specifier. If the pronoun is not in an introductory clause, an instance of the missing co-specifier condition has occurred.

Non-agent position pronouns

1. When a sentence in which a pronoun occurs is the second sentence of a new discourse, the recency rule may be applied. *Recency rule:* When a pronoun is in subject position and is the initial phrase in a sentence, and if a member of the potential (discourse or actor) foci occurs as the last phrase in the previous sentence, test the pronoun for co-specifying with that potential focus.
2. *Basic rule:* Check the discourse focus, followed by potential discourse foci, followed by the actor focus.
3. *Plural pronoun rule:* If the discourse focus is singular and the pronoun is plural, test the non-specific reading (only for foci that specify non-human entities) for co-specification, then potential discourse foci, followed by the actor focus.
4. *Focus related item rule:* If some discourse entity has been related to the focus during the discourse, test it for co-specifying with the pronoun in question.
5. *Focus stack:* Check items in the focus stack for co-specifying in last-in first-out order.
6. *Closure rule:* When all the above rules fail, if the pronoun occurs in an introductory clause, expect a forward co-specifier. Otherwise the pronoun is an instance of a missing co-specifier.

References

- Chafe, W. 1975 "Structures and Human Knowledge." In J.B. Carroll and R.O. Freedle (eds.), *Language Comprehension and the Acquisition of Knowledge*. V.H. Winston and Sons, pp. 41-69.
- Charniak, E. 1972 "Toward a Model of Children's Story Comprehension." MIT AI Lab TR-266.
- Chomsky, N. 1976 "Conditions on Rules of Grammar." *Linguistic Analysis* 2(4), pp. 303-351.
- Deutsch, B. 1975 "Establishing Context in Task-Oriented Dialogues." *Proceedings of the 13th Annual Meeting of ACL. AJCL Microfiche* 35.
- Deutsch, B. 1974 "Typescripts of Task Oriented Dialogs." Unpublished manuscript, SRI International, Menlo Park, Calif.
- Doyle, Jon 1978 "Truth Maintenance Systems for Problem Solving." MIT AI Lab AI-TR-419, January.
- Fahlman, Scott E. 1977 "A System for Representing and Using Real-World Knowledge." MIT AI Lab AI-TR-450, December.
- Grosz, Barbara 1977 "The Representation and Use of Focus in Dialogue Understanding." SRI Technical Note 151, Menlo Park, California.
- Grosz, Barbara 1978 "Focusing in Dialog." *Tinlap-2: Theoretical Issues in Natural Language Processing*. ACM and ACL, New York, 96-103.
- Grosz, Barbara 1981 "Focusing and Description in Natural Language Dialogues." In Joshi, Sag and Webber (eds.) *Proceedings of the Workshop on Computational Aspects of Linguistic Structure and Discourse Setting*. Cambridge University Press, Cambridge, pp. 84-105.
- Halliday, M.A.K. 1967 "Notes on Transitivity and Theme in English." *Journal of Linguistics* 3. pp. 117-274.
- Hobbs, Jerry R. 1976 "Pronoun Resolution." Research Report #76-1, City College, City University of New York.
- Hobbs, Jerry R. 1977 "38 Examples of Elusive Antecedents from Published Texts." Research Report #76-2, City College, City University of New York.
- Hobbs, Jerry R. 1979 "Coherence and Co-reference." *Cognitive Science*. 3(1):67-90.
- Katz, Jerrold J. & Fodor, Jerry A. 1963 "The Structure of a Semantic Theory." *Language*. 39(2):170-210.
- Lasnick, Howard 1976 "Remarks on Co-reference." *Linguistic Analysis*. 2(1):1-22.
- Lockman, Abe D. 1978 "Contextual Reference Resolution in Natural Language Processing." Dept. of Computer Science TR-70, Rutgers University, New Brunswick, N.J.
- Marcus Mitchell 1980 "A Theory of Syntactic Recognition for Natural Language", MIT Press, Cambridge, Mass.
- Martin, William 1979 "Roles, Co-Descriptors and the Formal Representation of Quantified English Expressions." Laboratory for Computer Science, TM-139, MIT, Cambridge, Mass.
- Morgan, Jerry L. 1978 "Toward a Rational Model of Discourse Comprehension." *Tinlap-2: Theoretical Issues in Natural Language Processing*. ACM and ACL, New York, 109-114.
- McAllester, David A. 1978 "A Three Valued Truth Maintenance System." MIT AI Lab Memo 473, Cambridge, Mass., May.
- Partee, Barbara H. 1972 "Opacity, Co-reference and Pronouns." in Davidson and Harman (eds.) *Semantics of Natural Language*. Reidel, Boston, pp. 415-441.
- Partee, Barbara H. 1978 "Bound Variables and Other Anaphors." *Tinlap-2: Theoretical Issues in Natural Language Processing*, 79-85, ACM and ACL, New York.
- Postal, P. 1969 "Anaphoric Islands." *Proceedings of the Fifth Meeting of the Chicago Linguistic Society*. Chicago, 5:205-239.
- Reichman, R. 1978 "Conversational Coherency." *Cognitive Science*, 2(4):283-327.
- Reinhart, Tanya 1976 "The Syntactic Domain of Anaphora." unpublished Ph.D. dissertation, Department of Foreign Literature and Linguistics, MIT, Cambridge, Mass.
- Reiger, Charles J. 1974 "Conceptual Memory: A Theory and Computer Program for Processing the Meaning Content of Natural Language Utterances." Stanford AI Lab Memo AIM-233.
- Robinson, Ann 1981 "Determining Verb Phrase Referents in Dialogs," *Am. J. Comp. Ling.* 7(1):1-16.
- Sidner, C.L. 1979 "Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse." MIT AI Lab Technical Report #537.
- Solan, Lawrence 1978 "Anaphora in Child Language." Ph.D. dissertation, Dept. of Linguistics, University of Massachusetts, Amherst (available through the Graduate Linguistic Student Association, University of Massachusetts).
- Van Lehn, Kurt A. 1978 "Determining the Scope of English Quantifiers." MIT AI Lab AI-TR-483, Cambridge, Mass.
- Walker, Donald E. (ed.) 1976 "Speech Understanding Research." Final Technical Report, Project 4762. AI Center, SRI International, Menlo Park, Calif.
- Webber, Bonnie Lynn 1978 "A Formal Approach to Discourse Anaphora." Technical Report 3761, Bolt Beranek and Newman Inc., Cambridge, Mass.
- Winograd, Terry 1972 *Understanding Natural Language*. Academic Press, New York.
- Woods, William 1973 "An Experimental Parsing System for Transition Network Grammars." in Rustin (ed.) *Natural Language Processing*. Algorithmics Press, New York, pp. 111-154.
- Woods, William et al. 1976 "Speech Understanding Systems: Final Technical Progress Report - 30 October 1974 through 29 October 1976." Technical, Report 3438, Bolt Beranek and Newman Inc., Cambridge, Mass.

Candy Sidner is a Scientist in the Department of Artificial Intelligence at Bolt Beranek and Newman Inc. She received the Ph.D. in computer science from M.I.T. in 1979.