

# **THE FLEXIBILITY OF INFERENCE IN TRIGGERS FOR INFERABLE ENTITIES: EVIDENCE FOR AN INTERPRETABILITY CONSTRAINT** (IN PRAGMATICS AND THE FLEXIBILITY OF WORD MEANING, PP. 51-77)

*Sharon A. Cote, James Madison University*

## **1 INTRODUCTION**

Determining the meaning of an utterance can require a great deal of inference. For example, the intended meaning of a polysemous word in a particular utterance is generally not explicitly stated; it is inferred from situational and linguistic information. Determining the antecedent of a pronoun, connecting a particular discourse referent with the use of a definite noun phrase, and even deciding whether an indefinite noun phrase has specific or non-specific meaning are also inferential processes. In addition, a hearer must make time and place inferences about events, as well as inferences about the rhetorical relations between pairs of utterances. In each case, the inferential process focuses at least in part on a hearer's understanding of a speaker's intentions.<sup>1</sup>

Speakers, therefore, must take into account the inferences that they expect hearers to make in choosing linguistic forms and expressions. For example, one of the most studied aspects of discourse structure is how speakers use different types of referring expressions to lead hearers to make different inferences. A robust model of how speakers make this choice must, at the very least, include a good characterization of the different possible mental activation statuses a

---

<sup>1</sup> Although this work refers to "speakers" and "hearers" for convenience, most of the argument is as applicable to writers and readers as it is to participants in spoken discourse.

discourse entity may have and a means for limiting the possible felicitous referring expression choices based on these statuses.<sup>2</sup>

One particularly complex and interesting part of this study concerns the classification of entities that achieve an increased activation status based on an inferential relationship to another discourse entity. Most activation statuses are characterized in terms of whether an entity has or has not been used in the current discourse before, as well as in terms of whether the speaker has any expectation that the hearer has previous knowledge of this entity. Inferable entities, in contrast, meet neither of these criteria and yet are still at least somewhat activated in the mind of the hearer. Specifically, they have a relevant relationship to some other activated entity. For example, in the discourse segment below, the expression *those people* refers to an inferable entity.

1) I got another call today from somebody trying to sell me long-distance services. *Those people* drive me crazy.

In this particular case, the speaker expects that the hearer will be able to infer, from the mention of a particular long-distance salesperson, that there is a set of long-distance salespeople. (This particular relationship will be categorized later in the chapter.)

The general informational state of inferable entities, as well as various subsets of more specific inferable relationships, has been recognized repeatedly by researchers in linguistics, psychology and computer science. Unfortunately, though not surprisingly, a wide variety of different names have been used to represent some or all of these entities, including “Inferrables” (cf. Prince 1981a, 1992), “indirect anaphors” (cf. Gundel and Erku 1987), and “inferentially accessible entities” (cf. Lambrecht 1994).<sup>3</sup> Inferable entities have also been directly or indirectly subsumed by more general studies of inference, such as the treatment of inferable information in Birner 1997 and, to some extent, the comprehensive idea of ostensive-inferential communication in Sperber and Wilson 1995. Partially as a result of these two practices, there is not yet a widely accepted theory of how inference works to give these inferable entities their meaning (reference). While it may well be that certain general inferential principles extend to the treatment of inferable entities, as I will show in this work, there is much work to be done before we completely understand even the specific inferences that are being made.

In the remainder of this chapter, I will characterize many of these more specific inference types based on the types of “triggers” that can make an entity inferable. In doing so, I will

---

<sup>2</sup> I am by no means claiming that level of activation is the only factor relevant to a speaker when choosing a referring expression. The interaction between activation and other aspects of information status, syntactic and lexical constraints, and other discourse and sociolinguistic factors, while not relevant to the purposes of this paper, must ultimately be part of any discourse model that claims psycholinguistic adequacy.

<sup>3</sup> There are even more different terms for various subsets of inferable entities; studies of some particular types of inferable entities, for example, can be found under discussions of “implicit participants” (cf. Mauner and Koenig 2000 ), “role-filling links” (cf. Garrod and Terras 1999), “one-substitution” (cf. Kuno 1987), and even, I’d argue, in the many casual references to so-called “generic *they* and *you*” (cf. Lambrecht 1994).

demonstrate that triggers for inferable entities do not necessarily provide enough information for a hearer to actually disambiguate the meaning of a referring expression. I argue that this is not a unique property of inferable entities. Comparing inferable entities to previous work on null subjects and implicit objects, it is clear that we need to define the reference of referring expressions (both null and overt) not in terms of identification of co-referents, binding, or recoverability of antecedents but in terms of “interpretability.”

## 2 PREVIOUS WORK ON INFERABLE ENTITIES

Prince (1992) characterizes a speaker’s perception of the informational state for “inferrables” in the following way:

Minimally, the speaker must have a warrant for believing (A) that the hearer already has the belief that the entity in question is plausibly related to some other ‘trigger’ entity [...], where the trigger entity is not (or, minimally, would not be) at the relevant point in time, hearer-new, and (B) that the hearer is therefore able to infer the existence of the entity in question. (9)

A trigger, then, may serve as an alternative to a co-referent in the set of discourse entities. It is an entity that the speaker believes is, minimally, known to the hearer (though not necessarily, under this definition, evoked in the current discourse) and somehow inferentially linked to the speaker’s target entity. Because the trigger substitutes for a co-referent, some of the referring expression forms that are used for entities with co-referents in the set of discourse entities may also be used for inferable entities. In what is perhaps the most straightforward case (specifically, Prince’s “Containing Inferrables”), an inferable entity is represented by a noun phrase that actually contains the inferable entity’s trigger. In example 2 below, taken from Prince (1992), *the door* is an inferable entity and its trigger is *the Bastille*. The hearer does not have to seek out a trigger for this inferable entity, the linguistic form explicitly ties the door to the building, and the hearer only need infer, trivially, that the speaker expects her to know that buildings have doors.

2) The door of the Bastille was painted purple.

Though it is perhaps arguable, therefore, that the information status of inferable entities with contained triggers may be signaled by the syntactic form of the noun phrase used,<sup>4</sup> many inferable entities do not have contained triggers in their referring expressions. We cannot rely on particular referring expressions to signal inferable entities either. In English at least, there is

---

<sup>4</sup> There are even problems with this argument because what Prince calls “Brand-New-Anchored” entities can use similar referring expressions; there is just no inference that the connection between the two entities should be expected by the hearer. I think there probably is, in fact, something to be gained by labeling these two types of entities differently, but the benefits are not visible within a single utterance.

no special type of referring expression used just for inferable entities<sup>5</sup>; the set of referring expressions available for reference to entities previously introduced into the discourse context (particularly the various definite noun phrase types and pronominal expressions) are also used to refer to some inferable entities. On occasion, as noted in Prince 1992, it is even felicitous for an inferable entity to be represented by an indefinite noun phrase. Notice the contrast in examples 3 and 4 below.

- 3) One day, while walking around outside on a blustery afternoon, I noticed *a feather* floating in the wind.
- 4) The brightly-colored bird shook its wings, and *a feather* floated down from the branches of the tree.

In example 3, the feather that is introduced into the discourse is a brand-new<sup>6</sup> discourse entity; there is no particular feather that a hearer might infer anything about in this context. In example 4, on the other hand, though an indefinite referring expression is also used, *a feather* is triggered by the discourse entity “bird”. The speaker can expect that the hearer will infer that the feather being observed is a feather from that particular bird. (This particular trigger relationship will also be characterized below.) The difference can be seen by looking at the relative acceptability in these two contexts of a subsequent utterance that depends on this inference. In examples 5 and 6 below, there is a clear difference in the acceptability of the utterance *It was odd to see it separated from the others*. This utterance is felicitous only in the latter discourse segment, because the feather in that context has been inferred to be a member of a set of feathers belonging to the bird.

- 5) One day, while walking around outside on a blustery afternoon, I noticed *a feather* floating in the wind. #It was odd to see it separated from the others.
- 6) The brightly-colored bird shook its wings, and *a feather* floated down from the branches of the tree. It was odd to see it separated from the others.

To further complicate matters, there is some variation in which linguistic forms constitute felicitous references to particular inferable entities. (Notice, for example, that the indefinite noun phrase *a feather* in example 6 would not work as an inferable entity if it was changed to a definite referring expression.) The apparent lack of lexical or syntactic distinctions between inferable entities and other entities, as well as the wide variation in how inferable entities are represented themselves, has made it difficult to fit them neatly into any ontological explanation for the different types of referring expressions.

Prince (1981a) suggested, therefore, that inferable entities may need to be broken down into different groups, and she laid out some examples based on types of conceptual relationships

---

<sup>5</sup> But see Birner (1997) for a discussion of the special intonational characteristics of inferable entities in certain non-canonical syntactic constructions.

<sup>6</sup> A brand-new entity, in Prince’s terminology, is essentially an unactivated entity. (Cf. Hajicová, E. and J. Vrbová 1982 for another discussion of the hierarchy of discourse entity activation.)

between entities. Another effort to deal with more specific types of inferable entities can be found in Gundel and Erku (1987), who argued, for example, that inferable entities represented by definite referring expressions are governed by the principle that "...a definite noun phrase is an instruction to the hearer to locate the referent in a shared set." (540) They also argued that all inferable entities are either included in sets that are evoked by another entity (inclusive), included in sets of which another discourse entity is a part (exclusive), or related to propositions or events described in the previous discourse (created).<sup>7</sup> There is something intuitively appealing here, but sets are very malleable concepts and the idea of sets does not, by itself, help to identify inferable entities in a reliable way. The interesting question is, "How do particular sets become relevant?"

## 2.1 Some Specific Types of Inferable Entities

Building on Prince (1981a), I have suggested that there are a number of very specific types of conceptual relationships that hearers have available to them as sources of inferences about new entities. In (Cote 1988, 2000a), I proposed that an inferable entity might, for instance, be linked to the lexical-semantic class of its trigger, either through an *INFERABLE CLASS* relationship or through an *ADDITIONAL INSTANCE OF CLASS* relationship.

Example 7, below demonstrates an inferable class relationship. The reference to *a Dalmatian* is the trigger. It makes available to the hearer the concept of a class of Dalmatians, and this is what permits the use of the referring expression *they* in the second utterance. There is no antecedent for this pronoun, but there is a trigger, which makes it interpretable.

- 7) The judge awarded "Best of Show" to a Dalmatian this year. I have always believed that *they* are great animals.

Example 1 above was another instance of a trigger of this type.

In example 8, lexical-semantic class is also relevant. In this case the trigger relationship is an additional instance of the class "brothers." The first utterance introduces the class, so *my brothers* serves as a trigger for *hers* in the subsequent utterance.

- 8) My brothers are workaholics. Marjorie says that *hers* is too.

It is worth noting here that the inferable entity in both these previous examples was referenced with a pronoun. While this is not the only or most common choice of referring expression, it is

---

<sup>7</sup> Again, there have been discussions of very specific types of inferable entities, characterizing them essentially as inferential oddities. I'm limiting the discussion here to research that recognizes the bigger pattern and is concerned with accounting in a systematic way for the linguistic choices that relate referring expressions and discourse statuses.

acceptable with these types of trigger relationships.<sup>8</sup> With other types of inferable entities, a pronominal reference may not be felicitous, as in 9 and 10 below.

9) We're planning a trip to Argentina in December. *The weather*/\**It* should be wonderful.  
(for "it" = the weather in Argentina in December)

10) The first grape harvest of the year makes a very dry wine. *The later harvests*/\**They* yield much sweeter wines.

I have argued that the trigger relation in examples like 9 is a *DISCOURSE DOMAIN ELEMENT* relation in which more general world knowledge comes into play. In order to successfully interpret *the weather* in the second utterance in example 9, the hearer must understand that Argentina, evoked in the previous utterance, is a geographic location, and that geographic locations have weather patterns. The hearer must also understand that taking a trip involves going to that location; there is a scenario or "domain" in which the weather in Argentina would be relevant. Notice that there is no way for a pronoun to refer to the weather in this particular sentence (despite the existence of a "weather-*it*" in English).<sup>9</sup> The interpretation of *it* would be the trip. Nonetheless, the discourse entity "a trip to Argentina" is a trigger for *the weather*.

In example 10, the trigger is *the first harvest*, and the fact that this particular entity is ordered (i.e., *first*) triggers, in particular, a *SEQUENTIAL* relationship with *the later harvests*. Using Gundel and Erku's terms, there is an exclusive shared set here, the set of harvests. We only know this, however, because a sequence is triggered.

Two other trigger relations discussed in Cote (1988, 2000a) are *PART-WHOLE* and *POSSESSIVE*, both of which depend on knowledge about the conceptual properties of certain things in the world.<sup>10</sup> Examples 10 and 11 both demonstrate a trigger that has a part-whole relation to a target inferable entity. In the first example, the trigger, *a chicken*, is introduced in the first utterance by person A. In person A's next utterance, *the head* can be inferred to be the head of the chicken. The trigger is the whole, and the inferable entity is a part. In example 12, the trigger, *a little white button*, is a part that is most commonly associated with a particular type of clothing, i.e. shirts. The target inferable entity then, *the damaged shirt*, can be inferred to be the whole to which this part belongs.

11) A: You ever seen anybody kill a chicken?

B: No, but I, I've, I've heard stories uh, I've heard stories.

A: And they put *the head* under a tin tub and chop the head off.

12) Glen found a little white button on the floor but still hasn't identified *the damaged shirt*.

---

<sup>8</sup> Gundel and Erku (1987) argue that pronouns are not possible with inferable entities, but these examples seem to fit their criteria.

<sup>9</sup> Of course, a small change making it clear that the speaker is now talking about Argentina rather than the trip is all it takes to make the weather-*it* acceptable (e.g., "It should be nice there.")

<sup>10</sup> I also discussed the possibility that clause-evoked entities are a type of inferable entity. Event reference of this type will be discussed briefly later in this chapter.

Examples 13 and 14 demonstrate the possessive relation. A hearer of the utterance in 13 can infer that “Johnny” has certain other entities that are associated with him. As a human being, for example, he has a mother. The reference *his mother* is therefore acceptable here. Similarly, in example 14, the speaker is an activated discourse entity in his own discourse and, though he has not mentioned that he is married, he expects the hearer to be able to infer that an adult living in suburban America is quite likely to have a wife. He can therefore say *my wife* without further presentational information. In contrast, he could not have reasonably expected his hearer to make the inference in example 15 below, because the hearer has no world knowledge that would make her expect that the speaker possesses a wallaby.<sup>11</sup>

13) Johnny got upset when *his mother* told him that it was bedtime.

14) Uh, and I cook a little bit now. What I like to do mostly is stir-fries and *my wife* normally says, oh, Tom, why don't you make a stir fry tonight.

15) Uh, and I cook a little bit now. #What I like to do mostly is stir-fries and *my wallaby* particularly likes my vegetarian stir-fry.

16) Wendy stubbed *her toe* on the pool ladder.

Finally, in example 16, a possessive noun phrase is again used to refer to an inferable entity, but the inference seems to be based on the part-whole relation. Toes are generally part of the whole human body. The only possessive relation would be that human beings do, in fact, have bodies. It is well known that, though the possessive form is used in English for this type of inferable entity, the preferred referring expression is not universally the possessive. The exact connection between part-whole relations and possessive relations bears further scrutiny.

### 3 A CORPUS STUDY

In Cote (2000b), I tested these inferable entity trigger relation types with a corpus study. In the process, I found some additional relation types and a number of complicating factors. The source for most of the data was the Switchboard Corpus of telephone conversations. Twelve discourses were examined, all of which involved two adult participants (“A” and “B”) discussing some randomly assigned topic (ranging from hobbies, to air pollution, to public service, to the weather), and all of which lasted about 10 minutes. Inferable entities were identified wherever a referring expression referenced an entity that was not already either evoked in the discourse or situationally present but that met Prince’s two minimal criteria. (See

---

<sup>11</sup> Of course, it doesn't seem absolutely out of the question for Tom to say something like this, but it would immediately be marked by the hearer as an unusual and humorous choice. It seems, in fact, that Tom would be intentionally flouting a conversational maxim (cf. Grice 1975). One might then predict the hearer, in fact, to focus on this referential choice, saying something like *You have a wallaby?*, rather than something like *That's sweet*.

discussion above.) A total of 227 inferable entities were then examined for types of relationships to their triggers.<sup>12</sup>

### 3.1 Triggers in the Discourse Context

190 of these switchboard tokens were either inferable entities with triggers found strictly outside their noun phrases or possessives. A variety of different types of triggers were found in these tokens, and the types of referring expressions also varied.

Examples 17-20 below would be classified, based on the earlier terminology, as discourse domain triggers. The triggers for the inferable entities in these examples come from the discourse participants' understanding of either the current discourse domain or the specific scenario under discussion. In example 17, the speaker can refer to *the garage* because the hearer is likely to know that inspections take place in garages. In example 18 the hearer can infer which city is being evoked because the events are set in the area of DFW (Dallas Fort Worth) airport. The hearer in 19 can infer that *the stitch* is a technical choice related to knitting. Finally, in example 20, the trigger for *my degree* is the previously discussed piece of information that B went to college for four years -- in the normal scenario for attending college, a student receives a degree after four years, and the speaker can assume that the hearer is likely to know this. In all examples of this type, speakers used a definite full noun phrase. Using the categories adapted from Cote 1988, 2000, these inferable entities are all classified as Discourse Domain elements, though it should be becoming clear that there is a lot of variation in exactly how these entities fit into the discourse domains. Only in example 19 is there an overt referring expression representing the trigger (i.e., the gerund *knitting*). In examples 17,18, and 20, there is no particular explicitly evoked discourse entity that serves as the trigger for the inferable entity. In these examples, scenarios must first be created as discourse entities; specifically, the hearer must know to create the "getting an inspection" scenario, the "flying into DFW" scenario, and the "attending college" scenario. In other words, these examples are like the weather example in the previous section. There does not appear to be any difference in referring expression choices based on whether the discourse domain trigger is or is not overt, so it may be simply that there is an additional inferential step to get to the domain triggers in some situations. Nonetheless, based just on these few examples, it is clear that what at first seemed like a fairly straightforward type of inference (roughly equivalent to a case of Gundel and Erku's created sets) will eventually need to be formalized in more detail.

17) A: And, and a lot of, when you, a lot of places when you go get the inspection they just pull it out of *the garage* and then pull it back in and they don't really check anything, they just make sure *the lights* work.

---

<sup>12</sup> The total set of inferable entities in these discourses is actually larger than 227, but it eventually became clear, for example, that including all possessive noun phrases referring to family members and body parts was not going to add anything to this particular study. Still, a separate examination of all inferable entities of this type will eventually need to be done for completeness and for the purposes of comparing part-whole triggers and possessive triggers in more detail.



18) A: I noticed once flying into D F W there was just a, a brownish-orange haze,

B: Yeah that's it,

A: over *the city*.

19) A: Okay. I, uh, I started knitting awhile ago, I knitted I didn't know, even know what *the stitch* is called. I just had this, uh, uh, piece of yarn and I wanted to start something so I remembered something I learned I think when I was five years old [laughter]. So I kept doing that, and now I have a little blanket but that's all I've done.

20) B: I mean, *my degree* is absolutely worthless.

As just one more example of complications in determining discourse domain inferences, look at example 21 below. The second italicized referring expression in this discourse segment (I will return to the first shortly) is an inferable entity that also involves a trigger in the discourse scenario, but the relationship is not quite the same. A bus and a bus stop are elements of sequential steps in the process of taking a bus. First, one *goes to a bus stop*, and then one *gets on a bus*. Not only is the relevant set ordered, but the inferable entity is not actually a set member; it is just an internal component of a set member. Note that a “bus” is not a part of a “bus stop”, so the part-whole relation does not work. Nor is there a particular bus, “the bus” associated with the domain of a bus stop. We need the whole park-and-ride scenario, with its time frame and purpose, as well as the set of steps one follows in this scenario, to understand what bus is being discussed. The bus that is inferred is the bus that would be the reason for going to the bus stop, not just any bus. In other words, one might consider this a *THEMATICALLY PARALLEL STAGE ELEMENTS* relationship.

21) A: And they go all around, that it's just real easy to get around once you're downtown, the problem is getting downtown and they have some, uh, Park and Ride, uh, expresses, where you go to the, *the bus stop* and you get on *the bus* and it takes you directly downtown, but they don't have enough of *them* and they're not convenient enough –

Now let's examine the first italicized referring expression in this discourse segment for a moment, i.e. *the bus stop*. Notice that this entity might also be treated as inferable, with the trigger again being the park-and-ride scenario. Alternatively, a Park and Ride may just be thought of as an entity itself, a particular type of bus stop. In this case, the expression the bus stop is just a lexical variant representing the same entity while emphasizing or introducing different characteristics of this entity. Certainly, alternative full noun phrase characterizations of an entity are not uncommon in discourse (epithets being one good example), but the problem is that it is sometimes difficult to tell whether an expression is referring to the exact same entity or just to a closely related but distinct and inferable entity. Since the entity certainly can be inferred from the scenario, it may be that this is the safest way for a hearer to treat it. On the other hand, if the speaker was expecting hearers to understand that a Park and Ride is a kind of bus stop, this lack of understanding could have implications later in the discourse.

There were many examples in this corpus involving inferable entities with triggers that are, in fact, themselves discourse entities. A number of these fit rather neatly into one of the trigger

relation types discussed above. For example, the part-whole relationship in example 11 of this chapter was actually taken from the Switchboard corpus. In example 22 below, however, the reference to *all the healthy food* is inferentially triggered for the hearer by a combination of the scenario (shopping) and an entity that can be classified as a purchasable food item (sweets). Like sweets, healthy food is a purchasable food item but this relationship is not quite the same as an additional instance of the lexical-semantic class “sweets”. There are differences between the two entities beyond the fact that they are distinct instances; they are not members of the same basic conceptual class but are instead members of a class (purchasable food items) that is evoked as much by the scenario (shopping) as by the mention of a particular food group. I will therefore tentatively define inferable entities of this type as a new relation in which the target entity is *SITUATIONALLY PARALLEL* to its trigger.

22) A: but we just went shopping and we came back with, uh, with, uh, sweets, you know, chocolate covered peanuts and --

B: Ugh.

A: -- uh, we came back with sweets. We didn't bring *all the healthy food* back too.

Yet another type of inferable entity is introduced in example 23. In the discourse from which this fragment is taken, the discourse participants have been discussing fishing and fishing locations. The entity “the garage” is not obviously inferable from anything that has come before in the conversation. Nonetheless, it is clear to us, and probably was meant to be clear to the listener, that A must have a garage at the place where he lives and that this is the garage to which he is referring. Why is this garage inferable? It seems reasonable to suggest that discourse participants not only use the discourse domain/scenario as a source of inference triggers, but are also prepared to infer the existence of entities that they know are prevalent in their everyday world. Possessive inferable entities like the earlier examples of family member inferences (i.e., examples 13-14) clearly work this way, and examples like 23 may be related. What this means though, is that it is possible to have a discourse domain trigger with the target inferable entity expressed as a possessive noun phrase (as in example 20), and a possessive trigger relation with a target inferable entity expressed as a definite referring expression (though not commonly in English.) There is also something intuitively different about the inference that a child has a mother and the inference that a man in the suburbs has a garage. In the former case, the two entities are actually equal members of a set (“family”), but there is no such set for a man and his garage. In Cote (2000b), I suggested that these latter triggers be separately referred to as *WORLD KNOWLEDGE* triggers, which is really just a suggestion that these kinds of inferences may not be completely categorizable without a better understanding of how we organize world knowledge in general.

23) A: Yeah. Now I've got most of mine hung up in *the garage* now. Like I say, it's been a long time since I've gotten to go, you know. [...]

### 3.2 Contained Triggers

Before moving on to inferable entities evoked with pronominal referring expressions, I would like to briefly explore the complications that arose in categorizing even the 37 tokens in this

study that could be classified as examples of inferable entities with triggers contained inside their noun phrases (not counting possessives). Simple examples of these are given in 24-26:

24) A: Um, I guess in Colorado, I'm trying to think of *the place we went* was in Pagosa Springs, and we went up kind of in *the Southwest corner of Colorado*.

25) A: I remember coming back into *the, uh, port where we left* and I had a cooler on my head as a hat.

26) A: I've got my fishing tackle in *the trunk of my car*, so if the urge ever [laughter]...

The types of triggers that can be contained within referring expressions are, at the very least, related to the types that can be discourse triggers. A larger study of contained triggers is needed to confirm the exact relationship; I will, however, point out here several factors that clearly complicate the identification of inferable entities with contained triggers even in this small set. The next three examples illustrate some of the more complex issues that can make inferable entities with contained triggers difficult to identify.

27) A: But, it's, uh, you know, I, I got into that mode where I was buying lures and, and rods and reels, and just all kinds of stuff. I had *one of those little two-man, little two-man boats that*,<sup>13</sup>

B: Um, um.

A: *used to go around a lot*, use that a lot, but like I say, I kind of got out of the fishing business when the kids got up big enough so, I,

In example 27, the noun phrase headed by the indefinite quantifier *one* refers to a member of the set of *those little two-man boats*. In other words, *those little two-man boats* triggers the availability of a subsequent reference to a single member of that class. This seems to be a fairly straightforward inference related to class membership. What makes this example more complicated is that the trigger is itself the head of a noun phrase that seems to contain a trigger clarifying which set of little two-man boats is being referred to (i.e., those *that used to go around a lot*). So, the trigger appears to be being introduced as an inferable entity too. It is not at all clear, however, that the speaker believes there is some other set of little two man boats that do not have this characteristic. He seems perhaps to be simply justifying why he thinks the little two man boats are already known to the hearer, or simply reminding the hearer of something that she already knows. On one level, the speaker therefore seems to be assuming that the boats are inferable entities and, on another, by his choice of an essentially vacuous trigger, to be assuming that they are hearer-old. One could explain this contrast as being the result of a rhetorical device, but that explanation would not necessarily determine which information status should be assigned to this entity.

---

<sup>13</sup> The observant reader may have noticed that there is another inferable entity (also with a contained trigger) in this discourse fragment: *that mode where I was buying lures and ...* It is not the intended example here, but I left it in rather than chopping up speaker A's turn.

Example 28 below raises a different question. The trigger for *the air conditioning* is *the car*. Note that this trigger is, at least superficially, not inside the noun phrase. Instead, it is extraposed. The question then is the following: Does the fact that the speaker chose to extrapose the prepositional phrase correspond to an assumption on his part that the scenario or domain of the discourse already provided a trigger? One could argue that this utterance would have been equally felicitous if the words *on the car* had been left out. Under that argument, the trigger for *the air conditioning* is instead the domain information that the scenario of “riding” (at least under the conditions described in this discourse) involves something that gets ridden in, something which quite likely would have air conditioning. On the other hand, the speaker did choose to include the prepositional phrase *on the car* and did not choose to use intonational phrasing that would suggest this phrase was an afterthought (although this does happen with some contained triggers). If the speaker intended all along to include this apparently unnecessary information, why was it extraposed? The significance of this example is that it shows that the decision to use a contained trigger may be marked in some way that we have not as yet identified.

28) A: Yeah, [laughter] my wife and I, the last day we were in Florida, we, my visiting my, uh, my parents, and my brother and sister, we were, we were down there and, the last day, right before we left, we had to, I had to ride with my sister and *the air conditioning went out on the car* [laughter],

Finally, example 29 raises the interesting possibility that a contained trigger may not completely obviate the need for some inference based on the broader discourse context. The context for this discourse segment is the following: speaker A has already talked about the last time he went camping. The prepositional phrase, *before that* is therefore crucial to understanding the event he is now introducing. He is not discussing “the last time”, he is discussing “the last time before the time already discussed.” So, there is a sequential trigger here as well. This example seems to suggest that triggering an inferable entity may be a process that requires more than one stage.

29) A: And then, I guess *before that, the last time I went camping*, was up on, backpacking up on the Continental Divide.

### 3.3 Inferable Entities and Pronominal Reference

As mentioned earlier, not all inferable entities are expressed with full noun phrases, whether they be definite or indefinite, containing or non-containing. Of the tokens examined in Cote (2000b), 37 occurred with pronominal referring expressions, and not all of these were clearly inferable entities. The breakdown of specific pronominal forms was as follows:

<b>Personal</b>		<b>Demonstrative</b>		<b>Possessive</b>		<b>Indefinite</b>	
<i>We:</i>	7	<i>That (or All that):</i>	4	<i>Mine:</i>	2	<i>One:</i>	1
<i>Us:</i>	1	<i>Those:</i>	1			<i>You:</i>	3
<i>It :</i>	1						
<i>They:</i>	13						
<i>Them:</i>	4						

Notice that I have tentatively included uses of the so-called “indefinite” or “generic” *you*<sup>14</sup> on this list because it seems that speakers intend hearers to constrain the set of possible referents for this pronoun based on subsets triggered by the discourse domain. For instance, in example 30, the speaker expects the hearer to know that the *you* must be people who are DISD employees.

30) A: [...] Well, I work for D I S D and I don't know what you're talking about when you say health insurance, but Dallas doesn't pay. It pays most of mine, you know, now, after *you* work ((I think)) five years, they begin to pay most of it. But then, for the, the family, you know, I put my kids on my policy. It's like two hundred dollars --

There are at least two questions that we can ask about these pronominal expressions. Which types of triggers allow them, and how can we fit them into theories of anaphora resolution that normally depend on the existence of antecedents.

3.3.1 Pronominal Reference and Inferable Co-Referents. It is generally accepted that one of the defining characteristics of personal pronouns is that they have co-referents, either antecedents provided by the discourse context or situationally available entities. Pronouns are used felicitously only when they can be identified with a co-referent, and even then there are additional constraints. (Cf. Walker *et al.* 1994 for one discussion.) In some tokens from the Switchboard data, an inferable entity represented by a pronominal linguistic form did, in fact, have an inferable co-referent.

In particular, as predicted earlier, class and additional instance of class triggers provided particular co-referents for inferable entities, allowing pronominal reference to these inferable entities. For example, the pronoun *those* in example 31 below refers to the class of jalapeño peppers; the trigger is the specific set of jalapeño peppers introduced in the previous utterance. This is a trigger though, not an antecedent, because the pronoun does not refer to these particular peppers. In example 32, *mine* refers to speaker A's fishing tackle, and the trigger is speaker B's reference to his own fishing tackle. In example 33, B's use of the pronoun *it* refers to the humidity (in Orlando) in general, and the trigger is A's discussion of the humidity in that city on a particular day.

---

<sup>14</sup> See, for example, Lambrecht (1994) who, along with many others, suggests that generic *you* and *they* may be used for entities that do not have active status in the discourse.

31) B: Yes. We planted, um, potatoes and onions and bell peppers.

A: Um, boy.

B: Uh, what else did I get in there, and jalapeño peppers

A: Oh, yeah, always got to have *those*.

32) B: I've got my fishing tackle in the trunk of my car, so if the urge ever [laughter],

A: [Laughter].

B: ever bites, you know, I'm ((kind of)) available but [laughter],

A: Go across the bridge, across the water, if it looks good, you can just pullover and start it then.

B: Yeah, right.

A: Yeah. Now I've got most of *mine* hung up in the garage now. Like I say, it's been a long time since I've gotten to go, you know. [...]

33) A: and this, it was just, the humidity was like eighty plus, eighty percent plus and it was just killing us.

B: Oh, yeah, *it's* just terrible. Orlando is the only place I've ever been where I've seen a car sweat.

There is another example that should be considered at this point, one which may or may not belong with this group. The expression *all that* in example 34 refers to the class of practical camping items, and the trigger is the several items introduced as a related list. This last example is the least clear case. As with the “purchasable food items” back in example 22, no one trigger here creates a class which can be referred to as a whole. While this complicates the idea of class triggers, it does seem plausible that the grouping of multiple related entities creates a class in the mind of hearer. On the other hand, the definition of class triggers gets much less straightforward if inferences like these are include. Though I have no final answer to offer here, another possibility is that this example is best explained in terms of a set that is introduced by a domain trigger. If so, the explanation for why there is a pronominal reference here gets more complicated (and perhaps has something to do with the quantificational nature of this particular expression).

34) A: There's a few things that you just kind of have to have or you can't go camping.

B: Uh-huh.

A: Depending on what kind of camping you're doing, like a stove or sleeping bag, and a tent, and *all that*.

In Cote (2000b), I incorporated an analysis of pronominal inferred entities resulting from either of the two class trigger relationships into a revised model of attentional state in local discourse structure (i.e., the Centering Theory discussed in some detail in Grosz *et al.* 1995). The major change required to handle the preferred interpretation of pronouns in these examples (as well as other phenomena<sup>15</sup>) was to make information from Lexical Conceptual Structures (cf. Jackendoff 1990, 1993) available for discourse processing. Assuming this information is

---

<sup>15</sup> Cf. Cote (1993, 1996) for a discussion of the need for lexical conceptual information to account for the interpretation of certain types of implicit objects in English.

available, pronominal references to inferable entities of this type do not force a revision of our concept of pronouns.

There are other kinds of triggers that provide co-referents for pronominal inferable entities, and some do clearly use the pronoun *that*. For instance, in example 35 below, *that* refers to the inferred result of a cooking process. These types of references, based on inferences about the results of a process, have been observed before, both with pronominal and definite reference, but the issue that remains is how they are related to other types of inferable entities. We could perhaps propose another kind of trigger, based on a *COMPLETED PROCESS*, and it may even be that these inferable entities are co-referential with implicit roles in the lexical conceptual structure of certain process verbs, but further study is needed.<sup>16</sup>

35) B: Then you, uh, pour *that* in there, you know, to make the gravy. And you let it come to a boil again and then you let it simmer and you add, uh, about a tablespoon of, what I'm using now is Jamaican curry powder.

Only a few of the remaining tokens of inferable entities with pronominal reference found in the switchboard corpus were of this last type. In the remainder, the pronouns did not behave as predicted under theories of co-reference.

3.3.2 Pronominal Reference and Elusive Co-Referents. In particular, one of the more interesting features of some of these inferable entities is that the hearer cannot always infer the existence of an unambiguous co-referent for the pronoun. I argue that, in some cases, a speaker may not intend a hearer to make an exact identification of the referent. Furthermore, there are cases where even the speaker may not even be able to pin down the referent. Yet these examples do not seem infelicitous.

Consider example 36 below. In this example, the speaker uses what is sometimes called a generic or arbitrary “they”. Though the speaker believes there is some specific group of people who have done the cracking down, and that these people should be identifiable if necessary, they are nonetheless not identified here. It is not even clear that B knows who the *they* are.

36) B: *They*'ve really cracked down up here.

Similarly, in example 37, the speaker is referring to fishing authorities of some sort whom he believes he has encountered and who are individuals who would know how to use the term “jig and a pig” correctly. The speaker does not really care, however, for his purposes in this discourse, who fits in this group.

37) B: One thing, uh, in fishing that I have not ever, have not done and I'm not really sure what *they* mean, I guess is when they're talking about using a jig and a pig.

---

<sup>16</sup> There were only three such examples in the Cote (2000a) data; clearly, a larger sample is needed.

In example 38, the speaker could have chosen a referring expression that would disambiguate between all the possible groups that could be *we* (the U.S. as a group, the speaker's state or hometown, the Western world, and so forth) in this utterance, but she chooses not to.

38) B: But, *we* have so much traffic now, so many cars. You know they're trying to fix it, uh, with all this emission control and everything.

Finally, in 39, the speaker is including himself in a group (*us*), but he is really talking about his own situation and is not concerned with whether the hearer interprets the group as the speaker and his schoolmates, the speaker and his family, the speaker and his father, the speaker and other people on limited budgets, or any other plausible option.

39) A: So, I mean, it's, for, for *us*, it's like, I, I can spend, I can spend that money but it's not, it's not what I want to spend it on, you know.

#### 4 INEXACT TRIGGERS AND INTERPRETABILITY

Though there are differences between the examples in the last section, they could, perhaps, all be classified as having an *Inexact Trigger* relationship to the previous discourse. In other words, while the trigger does give the speaker a warrant for believing that the hearer can infer the existence of such an entity, meeting Prince's criteria, it does not allow the hearer to actually identify this entity.

One reason why inferable entities with inexact triggers are expressed as pronouns may be that it is awkward and often difficult to create full noun phrases that have the same effect. For example, notice that difference between examples 40 and 41 below and examples 36 and 37.

40) *Some or all of the people in my area who make decisions about and enforce emission laws relevant to our discussion* have really cracked down around here.

41) One thing, uh, in fishing that I have not ever, have not done and I'm not really sure what *those people who write fishing articles or otherwise inform us about their expertise in fishing and who use the particular term I am about to mention* mean, I guess is when they're talking about using a jig and a pig.

It is not actually easy to come up with a noun phrase that captures all the possibilities that the speakers can mean in their succinct reference to *they*, nor is it clear that speakers themselves could tell you what they meant. In fact, the inferred entity in these cases is not an important part of the intended communication in their utterances. (It is interesting that these triggers are all human entities; this shared feature may ultimately also provide part of the explanation for why a pronominal reference is possible.)



## 4.1 The Interpretability Constraint

The actual requirement for felicitous use of these pronouns might be defined in the following way:

### **Interpretability Constraint:**

A hearer must be able to assign as much meaning to a pronoun as is needed to avoid causing a speaker to fail to achieve his discourse purpose.

The means for achieving interpretability would be dependent on the intentions of the speaker. While the “normal” uses of pronouns, where the entity referred to by the pronoun is identified with another entity in the discourse, would certainly meet this constraint, there are many other possibilities. A possible set of means would include at least the following options:

## 4.2 Means for Achieving Interpretability with Pronouns<sup>17</sup>

Here is a non-comprehensive list of ways speakers satisfy the interpretability constraint:

- i. Discourse and situational structure features combine with semantic features and syntactic co-reference constraints to determine a set of possible co-referents for a referential pronoun and, when there is more than one possible co-referent, to create a ranking of the possibilities in terms of likelihood in order to facilitate disambiguation.
- ii. Lexical conceptual features of another referring expression lead to the inference of a possible referent.
- iii. The pronoun is one which can take arbitrary reference and this interpretation is consistent with the discourse context.
- iv. The pronoun is one which can be expletive, and this interpretation is consistent with the discourse context.
- v. Semantic features of the pronoun itself combined with the discourse context and/or world knowledge lead to the inference of a loosely defined set of plausible antecedents, and the discourse intentions of the speaker do not require that this ambiguity be resolved.

The second possibility on this list distinguishes the inferred co-referents for inferred entities with class and additional instance of class triggers from interpretability based on overt discourse antecedents. The third and fourth allow for a natural treatment of truly arbitrary or expletive pronouns under the same constraint. The final option accounts for felicitous pronominal references with inexact triggers.

---

<sup>17</sup> I first started talking about interpretability with respect to null subjects in English, which are not always recoverable in felicitous uses. (Cote 1996) There are a slightly different set of options for achieving interpretability with null arguments, but the principle is the same.

## 5 OTHER EVIDENCE FOR THE INTERPRETABILITY CONSTRAINT

If the interpretability constraint were needed only for the treatment of inferable entities, one might suggest that there is simply a different process at work for inferable entities and entities evoked in other ways. There are, however, other motivations for this constraint.

### 5.1 Null Subjects

Cote (1996) first proposed the interpretability constraint as a way to explain the interpretation of null subjects in English and other languages. Though null subjects are not as prevalent in English as they are in so-called “pro-drop” or “discourse-oriented” languages, they are, in fact, not uncommon in conversational or informal written English. They can even be found occasionally in somewhat formal written English registers. Naturally-occurring tokens of these English null subjects can be found in examples 42-45 below.

42) This is Sid. *Ø* thought I'd call you up.  
(conversation in Hopper 1992:35)

43) Oops -- *Ø* won't hear me complaining.  
(television commercial)

44) Are you sure you wanna change it? *Ø* Looks kinda sexy to me.  
(television commercial)

45) *Ø* Sounds as if Mr. X has a lot in common with those nuts who parade around the countryside at night in white sheets, ...  
(a letter to the editor of a newspaper)

While null subjects as a group in English cannot be accounted for by late phonological reduction processes, and can be shown to have specific discourse constraints,<sup>18</sup> they do not always meet the recoverability condition on empty categories.<sup>19</sup> Examples 46-47 include tokens of this type. In example 46, the null subject could be referring to the speaker as an employer, to the speaker and her office colleagues, to the offices in her area in general, and so forth. An exactly defined subject simply is not necessary for the speaker's purposes here. Similarly, in example 47, it is not relevant to the speaker's point that the hearer know whether it is she, her husband, or both of them together making the stop for the video.

---

<sup>18</sup> See Cote (1996) for a detailed discussion of the grammatical and discourse properties of English null subjects, including a large corpus study (also from the Switchboard corpus) showing that the discourse constraints observed in these null subjects are statistically different from the constraints on overt pronouns and other referring expressions in English.

<sup>19</sup> See Roberge (1990) for one discussion of the recoverability condition.

46) A: We're having a lot of allergies down here right now.  
B: Uh-huh.  
A: Everything blooming, and the weather.  
B: Uh-huh.  
A: And I think a lot of people have contracted spring fever too, so.  
Ø Had a lot of people out at work, you know, for fishing and, and uh,  
B: [Laughter].  
A: and golfing, reasons and things like that.

47) A: We, uh, got a new baby in the house --  
B: Oh, yeah.  
A: -- and, and, she just turned a year old, so it just kind of put the --  
B: Yeah.  
A: -- clamp on things. Uh,  
B: So you watch videos.  
A: Yeah.  
B: Then, uh,  
A: Ø Stop by and get them at, you know, for ninety-nine cents, and bring them home but,

In each of these cases, it is impossible to determine the full set of features associated with the null subject. In other words, there is no overt pronoun in English that could even be reliably substituted for these null pronouns. With these ambiguities, the subject cannot be 'recovered' in the traditional sense because there is no overt subject that can maintain these particular ambiguities. The utterance can, however, still be interpreted by a hearer acting under the assumption that the missing features are irrelevant to the information the speaker wishes to convey.

This aspect of interpretability is compatible with the Gricean maxim of quantity; even overt pronouns may be underspecified in certain ways. For example, the use of *we* in the following example could mean "my family", or just "my spouse and I". The distinction is simply not crucial to the speaker, and a listing of the included members therefore would be non-cooperative.

48) The old mini-van finally died, so *we* got a new station wagon last week.

For this reason, and because there is so much variation in the use of null arguments cross-linguistically that has not been adequately accounted for,<sup>20</sup> Cote 1996 argued that a more robust hypothesis might be that there is a continuum from null argument to non-null argument languages rather than a number of preset language types, and that the only absolute constraint

---

<sup>20</sup> For example, like the various null arguments found in a number of other configurational, non-discourse-oriented languages, English null subjects also do not meet any of the agreement conditions on the licensing of null pronouns in pro-drop languages.

on null arguments is interpretability. A language may, in principle, obey this constraint by using null arguments that have one (or more) of the characteristics below:

- i. Recoverable from grammatical features in sentence
- ii. Recoverable from discourse/situational context
- iii. Recoverable because of limited feature variation in null pronouns in that language
- iv. Arbitrary in reference
- v. Expletive
- vi. Identification of exact referent of subject unnecessary

Within the range of these possibilities, we can try to determine whether there is any correspondence between general discourse constraints and/or functions associated with the types of null arguments available in a particular language and the way null arguments satisfy interpretability in that language. In other words, we can hope to learn if the functions of null arguments are partially or completely dependent on the patterns of interpretability.

## 5.2 Null Objects

Cote (1996) also discussed interpretability with respect to null objects, a largely lexical phenomenon in English. Null objects in English have been shown to be constrained by choice of verb, both in terms of availability and in terms of how they are interpreted. Certain verbs, for example, require that a null object establish a new discourse entity (though it may ultimately be shown to be identical to an already existing one). Hence, verbs like *eat* assign partial interpretations to their null objects directly. In other words, there is yet another way that a null argument can satisfy the interpretability constraint:

- vii. Sufficient information can be extracted from lexical constraints

Two examples of null object utterances are given below in 49-50.

49) The young man always wrote very carefully.

50) At Sunday's picnic, the children ate sitting on blankets.

*Write* and *eat* are both verbs that allow what Cote (1996) calls "Indefinite Null Objects." In 49, the null object is the discourse entity corresponding to what the young man wrote. It is constrained only to be a set of things that the intended hearer should not consider abnormal for the young man to be writing in whatever context is provided. Similarly, in 50, the null object is the discourse entity corresponding to what everybody ate, which is constrained only to be something that the intended hearer should not consider abnormal for the children at the picnic to be eating. Other types of null objects would have different interpretability constraints.

### 5.3 Event Reference

Anecdotal evidence suggests that interpretability may turn out to be the only absolute constraint on event reference too. For instance, in example 51 below, the event to which *that* refers is unclear. It could be the whole series of actions proposed by the first speaker or just some part. Subsequent utterances could disambiguate, as in the two alternative discourse progressions in example 52. Yet, if the speaker is unconcerned about this ambiguity, the reference may never be resolved, as in the possible progression in example 53.

51) A: I just can't stand the attitude at my office anymore. I'm just gonna quit, open up a used book store, and feed stray cats all day.

B: *That* won't work.

52) B1: *That* won't work. Used book stores don't make any money. You should open a used CD store instead.

B2: *That* won't work. It's never good to run away from your problems. You should try to change the attitude at your office by introducing "Casual Fridays."

53) B: *That* won't work. Life is so depressing. Let's go get hot fudge sundaes.

If a speaker does not need a hearer to disambiguate to make his point, it should not be surprising if he does not go out of his way to provide the information needed for disambiguation. These event references may, in fact, be very similar to certain kinds of discourse domain triggered inferable entities, but this is material for future work.

In general, hearers are not surprised and do not perceive anything infelicitous about loosely interpretable referring expressions.<sup>21</sup> Of course, there can be a conflict between what the speaker intends to communicate and what the hearer hopes to learn. Interestingly, this conflict seems to come into play in the use of pronominal event reference in written contexts. Anyone who has read a college-level essay knows that speakers (actually writers in this case) may attempt to hide sloppy thinking behind loosely interpretable pronominal event references. Hence the ban in some classes on the use of pronominal *this* and *that* in written work.

## 6 RELATED ISSUES

The interpretability constraint is clearly related to other discourse processing issues. Although an extensive discussion of these connections is beyond the scope of this work, I will point out in this section just a few issues that need to be considered.

---

<sup>21</sup> If you're skeptical, consider this. When's the last time you thought it odd when someone wrote "Hope to see you soon" in a card, letter, or email message? Did she mean herself, her whole family, or some other subset?

## 6.1 Failed Anaphora Resolution

There is some experimental evidence that hearers sometimes themselves choose not to make the effort to resolve anaphors, even when the speaker does provide the necessary information. For example, Levine *et al.* (2000) did a study in which they determined that readers sometimes simply stop trying to resolve definite noun phrase anaphors that have antecedents that are difficult to retrieve. The difficulty arose when they provided a “distractor” along with an antecedent. The distractor was semantically plausible antecedent but did not play the discourse role of the antecedent. A simplified example of this type of context is reproduced below.

54) Wanda was throwing a surprise party for her best friend John. John had just been promoted to Vice President of the company and some of his close friends wanted to congratulate him. Wanda even made him a *tart*.

She felt a little pressured because her daughter’s graduation dinner was the next day and she needed to prepare for that as well. She also still had to get decorations and stop at the bakery for a *cake*. Her daughter loved chocolate *cakes*.

Wanda hoped John’s party would be fun. The guests arrived right on time. As everyone sat down to eat, Wanda said to leave room for *the dessert*.

*Anaphora resolution task: What was the dessert?* (Levine *et al.* 2000, 612)

As the authors themselves mention, the reader could simply be treating *the dessert* as a new entity (inferable, really), but, technically, it is not. The discourse context does in fact include the information that the dessert at John’s party was a tart.

At least two questions relating to interpretability arise. How much of the intended interpretation is actually enough? How much do hearers really care about speakers’ discourse purposes? These questions have impacts for other aspects of discourse study as well.

## 6.2 Relevance

The interpretability constraint also seems to be very compatible with the principle of relevance discussed in Sperber and Wilson (1995). Relevance is, according to this work, crucial to all “inferential communication.” For example, hearers make referential hypotheses because references generally are not recoverable by linguistic decoding alone. More generally, Sperber and Wilson talk about semantic incompleteness, but then suggest that identification of propositional form is still an essential part of the process. It seems that it would take little to say that full identification is not an absolute constraint in inferential communication.

## 6.3 Local Discourse Models

Finally, the interpretability constraint has substantial implications for models of pronoun resolution and local discourse coherence. Current models treat co-indexing with an antecedent

as a goal in and of itself rather than as a means to achieve interpretability. If the interpretability constraint is, in fact, the only absolute constraint, then the design of these models needs to be modified.

For example, Centering Theory (cf. Grosz *et al.* 1995, Walker *et al.* 1994, Cote 1998), is a model of local discourse coherence which, as it is currently discussed and implemented, makes the following assumptions:

- i. At any point in a discourse segment, there is a backward looking center that serves as the crucial point of continuity with the previous utterance, and
- ii. a partially ordered forward-looking centers list of discourse entities that may become the next backward-looking center.
- iii. The ordering of the discourse entities on this list is an indication of how likely each of them is to be the next backward-looking center.

A hearer may, therefore, choose between ambiguous interpretations of an utterance by considering these features.

Even in robust versions of the theory, which incorporate information status and lexical conceptual information as well as grammatical information in the template used to form forward-looking center lists, the antecedent for pronouns is still sought on these lists. Ultimately, theories like this will need to allow for inexact triggers and other linguistic phenomena that allow incomplete identification of an antecedent. In addition, they will have to be modified to handle a wide variety of still largely under-formalized triggers for inferable entities.

## 7 CONCLUSION

In this work, I have attempted to demonstrate that inferable entities are complex phenomena and that, while some specific classification choices can improve our ability to predict the availability of inferable entity references, we do not yet have all the tools we need to describe these phenomena. I have discussed some of the observations that arise out of even a relatively small data study and proposed areas where more data need to be collected. Perhaps the most interesting observation is that not all triggers for inferable entities serve to identify those entities with unique co-referents.

More generally, I have argued that an interpretability constraint on various types of discourse entity inferences, including some inferable entities, correctly represents the real inferential expectation that speakers attempt to fulfill for hearers. Other supposedly absolute constraints are actually common means for adhering to the interpretability constraint, but certainly not the only means. This seemingly small change in perspective can have significant impacts on our understanding of anaphoric references and of the structuring of local discourse.

## REFERENCES

- Birner, B. J. (1997). The linguistic realization of inferrable information. *Lang. & Comm.*, **17**, 2, 133-147.
- Chafe, W. (1976). Givenness, contrastiveness, definiteness, subjects, topics and point of view. In : *Subject and Topic* (C.N. Li, ed.) pp. 25-55. Academic Press, New York.
- Cote, S. A. (1988). *Just what Are you inferring?: Toward an analysis of inferable entities in discourse*, ms.
- Cote, S. A. (1993). The ranking of forward-looking centers. Presented at the workshop on centering theory in naturally-occurring discourse, IRCS, University of Pennsylvania, May 8-10.
- Cote, S. A. (1996). *Grammatical and discourse properties of null arguments in English*. Ph.D. Dissertation, University of Pennsylvania.
- Cote, S. A. (1998). Ranking forward-looking centers. In Marilyn Walker, Aravind K. Joshi, and Ellen F. Prince, (eds.), *Centering theory in discourse*. Oxford University Press.
- Cote, S. A. (2000a). Inferable entities in discourse and in narrative. Presented at Linguaging 2000, University of North Texas, March 3-4.
- Cote, S. A. (2000b). Lexical Conceptual Structure and Inferable Entities in Discourse. Presented at the 7<sup>th</sup> International Pragmatics Conference, Budapest, July 11.
- Garrod, S. and Terras. M. (1999). The contribution of lexical and situational knowledge to resolving discourse roles: bonding and resolution. *J. Mem. & Lang.*, **42**, 526-544.
- Grice, H. P. (1975). Logic and Conversation. In: *Syntax and Semantics 3: Speech Acts* (Cole, P. and Morgan, J., eds.), 41-58. Academic Press, New York.
- Grosz, B., A. K. Joshi, and S. Weinstein (1995) Centering: A framework for modelling the local coherence of discourse. *Comp. Ling.*, **21**, 2, 203-225.
- Gundel, J. and Erkü, F. (1987). The pragmatics of indirect anaphors. In: *The pragmatic perspective: Selected papers from the 1985 international pragmatics conference* (J. Verschueren and M. Bertuccelli-Papi, eds.), pp. 533-545, John Benjamins, Amsterdam.
- Hajicová, E. and J. Vrbová (1982). On the role of the hierarchy of activation in the process of natural language understanding. In: *COLING 82* (J. Horecky', ed.), 107-113.
- Hopper, R. (1992). *Telephone Conversation*. Indiana University Press, Bloomington.
- Jackendoff, R. (1990). *Semantic structures*. The MIT Press, Cambridge.
- Jackendoff, R. (1993). On the role of conceptual structure in argument selection: A reply to Emonds. *NLLT*, **11**, 279-312.
- Kuno, S. (1972). Functional Sentence Perspective: A case study from Japanese and English. *LI*, **3**, 269-320.
- Kuno, S. (1987). *Functional Syntax*. The University of Chicago Press, Chicago.
- Lambrecht, K. (1994). *Information structure and sentence form*. Cambridge University Press: Cambridge.
- Levine, W. H. and A. E. Guzmán, and C. M. Klin (2000). When Anaphor Resolution Fails. *J. Mem. & Lang.*, **43**, 594-617.
- Maurer, G. and Koenig, J.-P. (2000). Linguistic vs. conceptual sources of implicit agents in sentence comprehension. *J. Mem. & Lang.*, **43**, 110-134.
- Prince, E. (1981a) Toward a taxonomy of given-new information. In: *Radical pragmatics* (P. Cole, ed.), 223-255. Academic Press, New York.



- Prince, E. (1981b) On the inferencing of indefinite-*this* NPs. In *Elements of discourse understanding* (A. K. Joshi, B. L. Webber, and I. A. Sag, eds.), 231-250. Cambridge University Press, Cambridge.
- Prince, Ellen (1992) The ZBG letter: Subjects, definiteness, and information-status. In: *Discourse description: Diverse analyses of a fundraising text* (S. Thompson and W. Mann, eds.), 295-325. John Benjamins, Amsterdam.
- Roberge, Y. (1990). *The Syntactic Recoverability of Null Arguments*. McGill-Queen's University Press, Montreal.
- Sperber, D. and D. Wilson (1995). *Relevance: communication and cognition*. Blackwell: Oxford.
- Walker, Marilyn, Masayo Iida, and Sharon Cote (1994) Japanese discourse and the process of centering. *Comp. Ling.*, **20**, 2, 193-232.
- Walker, Marilyn A. and Ellen F. Prince (1996) A bilateral approach to givenness: A hearer-status algorithm and a centering algorithm. In: *Reference and referent accessibility* T. Fretheim and J. Gundel, eds.), 291-306. John Benjamins, Amsterdam.