References in Narrative Text

Janyce M. Wiebe
Department of Computer Science
University of Toronto

Abstract

The *propositional content* of a reference is the proposition attributing to the referent the properties that correspond to the nouns and modifiers in the reference (for example, the propositional content of 'Mary' is that the referent is named 'Mary'). During language comprehension, the hearer or reader must determine the set of beliefs with respect to which the propositional content of a reference is to be understood. In the prototypical case, this set consists of the propositions that she believes that the speaker or writer believes that she and the speaker or writer mutually believe. This paper identifies two contexts in which the propositional content of a specific reference is not understood with respect to this set—*subjective* and *objective* sentences in third-person fictional narrative text—and identifies some implications of this for understanding specific references in these contexts.

1 Introduction

Specific references are references to particular entities, for example, 'a car' in (S1) and 'the car' in (S2), as opposed to 'a car' in (NS1) or 'the car' in (NS2):

S1. John bought a car yesterday.

S2. John couldn't get to work today because the car wouldn't start.

NS1. John wants to buy a car, but he hasn't chosen one yet.

NS2. The car is a four-wheeled vehicle.

A specific reference has a *propositional content*—the proposition attributing to the referent the properties that correspond to the nouns and modifiers in the reference. For example, the propositional content of 'Mary' is that the referent is named 'Mary', the propositional content of 'the car' is that the referent is a car, and the propositional content of 'he' is that the referent is male. During language comprehension, the recipient (the hearer or reader) must determine the set of beliefs with respect to which the propositional content of a reference is to be understood. In the prototypical case, this set consists of the propositions that she believes that the producer (the speaker or writer) believes that she and the producer mutually believe. This paper identifies two contexts in which the propositional content of a specific reference is not understood with respect to this set—*subjective* and *objective* sentences in third-person fictional narrative text—and identifies some implications of this for understanding specific references in these contexts.

This work has grown out of the work in philosophy (e.g., Castañeda 1970), linguistics (e.g., Fodor 1979), and artificial intelligence (e.g., Rapaport 1986) on the interpretation of references in *opaque contexts*. Opaque contexts are the objects of sentences with third-person subjects and with main verbs that are either propositional-attitude verbs, such as 'believe', 'know', and 'hope', or other psychological and perceptual verbs, such as 'want', 'hate', and 'see'. Under one interpretation of such a sentence (the *de dicto* as opposed to the *de re* one), the propositional content of references in the object are understood with respect to the beliefs of the subject of the sentence. As we shall see, references in subjective sentences have a similar property, even in those that do not have the syntactic form described above.

In section 2, I provide background information about subjective sentences in third-person narrative text. Sections 3, 4, and 5 identify the set of beliefs with respect to which the propositional content of a specific reference is understood in conversation, objective sentences, and subjective sentences, respectively. In other work, I have developed an algorithm for distinguishing between subjective and objective sentences in third-person narrative text (Wiebe 1990a); this algorithm can be used to choose the proper set of beliefs with respect to which references should be understood in this genre. Section 6 shows some implications of this choice for computational approaches to the resolution of specific references. Section 7 considers the the *de re*/*de dicto* distinction in light of the issues raised in earlier sections, section 8 draws some parallels between the interpretation of references and the interpretation of *subjective elements*, and section 9 discusses some issues for future research.

Throughout this paper, it should be assumed that only specific references are being considered.

2 Subjective Sentences

Sentences of fictional narrative can be characterized as either *subjective* or *objective* (Banfield 1982). In contrast to sentences that objectively narrate events or describe the fictional world, subjective sentences present the consciousness of a character within the story. They express characters' evaluations, emotions, judgments, uncertainties, beliefs, and other attitudes and affects. Some present a character's thought or perception (*represented thought* or *represented perception*; Banfield 1982); others *report* a character's *private state*, which is a perceptual state such as seeing, a psychological state such as wanting, or an experiential state such as feeling ill—that is, a state that is not open to objective observation or verification (Quirk et al. 1985).[1] We shall call the character whose consciousness is presented by a subjective sentence the *subjective character* of that sentence.[2] Further, an *objective context* is a maximal block of objective sentences, and a *subjective context* is a maximal block of subjective sentences with the same subjective character.

The following passages illustrate different kinds of subjective sentences (throughout this paper, sentences in cited passages are indented to reflect paragraphing in the original texts).

(1)

[1.1]"What are you doing in here?" [1.2]Suddenly she [Zoe] was furious with him [Joe].

[1.3]"Spying, of course."

[1.4]"Well of all dumb things! [1.5]I thought you ran away." [1.6]Joe Bunch was awful. [Oneal, *War Work*, p. 130]

Sentence (1.2) is a private-state report: It reports Zoe's private state of being furious with Joe. Sentence (1.6) is a represented thought: It presents Zoe's thought, and it expresses her evaluation of Joe (that he is awful).

(2)

[2.1]Certainly, Dennys thought, anything would be better than this horrible-smelling place full of horrible little people.

[2.2]There was a brief whiff of fresh air. [2.3]A glimpse of a night sky crusted with stars. [L'Engle, *Many Waters*, p. 25]

Sentence (2.1) is Dennys's represented thought, and (2.2) and (2.3) are Dennys's represented perceptions.

The feature of subjective sentences that is important to the current work is that they reflect the subjective character's beliefs, which may be false in the fictional world (this has been noted by many in literary theory, such as Doležel 1973, Kuroda 1973ab, Uspensky 1973, Cohn 1978, and Banfield 1982). For example:

(3) This was David's boy. [Bridgers, *All Together Now*, p. 91]

This sentence is the represented thought of a character named 'Dwayne'; it is about a female character whom Dwayne incorrectly believes to be a boy. Another example is the following:

(4) She [Morgaine] was indeed a witch, he [Chei] thought. [Cherryh, *Exile's Gate*, p. 64]

Sentence (4) is the represented thought of Chei. Morgaine is not actually a witch, but Chei believes that she is one. (The word 'witch' is not used as a figurative term in (4) or in (8), below; Chei believes that there are such things as witches.)

Before leaving this section, it should be noted that a sentence of quoted speech, such as (1.1), is really a mixture of genres. First, it is a narrative sentence whose main verb phrase contains a communicative verb, and whose object is a quoted string (though the *discourse parenthetical*, such as "Zoe said", may be only implicit, as in (1.1)). Second, the contents of quoted strings are conversational utterances. The discussion in this paper of references in third-person narrative text applies to references in the discourse parenthetical (if there is one), not to any in the quoted string.

3 References in Conversation

In conversation, references are directed by the speaker toward the hearer. To understand a reference, the hearer cannot consider only her own beliefs, or only what she believes that the speaker believes; she may also have to consider what she believes that the speaker believes that she believes, and so on. Clark and Marshall (1981) and Perrault and Cohen (1981) describe some of the complexities of this process.

Clark and Marshall show the potential relevance of a conjunction of beliefs that comprise the speaker and hearer's *mutual beliefs*:

(MB) $\underline{S}$ believes that P, and $\underline{S}$ believes that $\underline{H}$ believes that P, and $\underline{S}$ believes that $\underline{H}$ believes that $\underline{S}$ believes that P ...

where '$\underline{S}$' stands for the speaker, '$\underline{H}$' stands for the hearer, and P is the proposition mutually believed. Clark and Marshall point out that speakers and hearers cannot really consider an infinite series of conjuncts during processing, and they suggest some heuristics that speakers and hearers might use to assess their mutual beliefs in a finite amount of time.

Perrault and Cohen show that beliefs acquired privately about objects, people, places, etc., can "override" some of the conjuncts of mutual belief. They give the following example, adapted from Donnellan 1966:

> (5)
>
> S and H are at a party. They watch together as water and gin are being poured in two identical glasses and given to women W1 and W2, respectively. Unbeknownst to H, S sees W1 and W2 exchange glasses. Later S tells H: "The woman with the martini is the mayor's daughter." [Perrault and Cohen 1981, p. 222]

The speaker is referring to W2, even though she herself does not believe that W2 is the woman with the martini. However, she believes that the *hearer* believes that W2 is the woman with the martini, and, since the hearer doesn't know that the speaker has changed her belief, the speaker also believes that the hearer still believes that the speaker and hearer mutually believe that W2 is the woman with the martini.

Of course, from the hearer's perspective, the reference does reflect the series of conjuncts in (MB). However, the hearer can also acquire a private belief that affects her understanding of the speaker's reference. Cohen, Perrault, and Allen (1982, p. 257) give the following example (I have replaced 'system' by 'hearer' and 'user' by 'speaker'):

> (6)
>
> Suppose that at first the hearer and speaker agreed that Kirk was the captain [of the Enterprise]. Then suppose that the hearer found out through direct, private access to the Enterprise that Kirk had been replaced by Spock. The hearer would therefore believe that Spock was the captain, while believing that the speaker believed that Kirk was. The speaker's utterance of "the captain of the Enterprise" still clearly identifies Kirk, and should be understood as such by the hearer.

The hearer understands that the speaker is referring to Kirk because although the hearer believes that Spock is the captain, she believes that the speaker believes that Kirk is the captain and also that the speaker still believes that the hearer and speaker mutually believe that Kirk is the captain.

Perrault and Cohen show through example that any finite number of the conjuncts of (MB) might not be true for a successful reference; what is required, taking the hearer's perspective, is some nested belief of the form

$\underline{H}$ believes that $\underline{S}$ believes ... (MB).

If a more nested environment than $\underline{H}$ believes (MB) is required, then the hearer understands the reference in spite of some private belief that she holds or some private belief that she believes that the speaker holds. (Note that Cohen 1978 specifies a finite representation of mutual belief.)

The importance of this discussion for our purposes is that the hearer cannot consider only what she believes or only what the speaker believes in order to understand references in conversation. She has to consider what she believes that the speaker believes that she believes, and she has to distinguish what she privately believes (and what she believes that the speaker privately believes) from what she believes they mutually believe.

Before leaving this discussion of conversation to consider narrative text, there is something obvious about conversation that needs to be noted, so that objective sentences can be contrasted with it. In conversation, a hearer always has the option of questioning the veracity (or sincerity) of the speaker's utterances. Even the most gullible listener with the greatest belief in the speaker's authority might question the truth of the speaker's utterances, if things begin to sound too preposterous. In particular, the hearer can question the propositional content of references. That is, the hearer may believe that a reference reflects a false belief of the speaker (as discussed previously in this section)—she may believe that the speaker herself holds the false belief, she may believe that the speaker holds a false belief about what the hearer believes, and so on.

4 References in Objective Sentences

Objective sentences narrate events independently of any character's consciousness. The major difference between objective sentences and conversation is that the reader does not question the truth of an objective sentence. This is because objective narration actually creates the fictional world, to which the reader has no other recourse but through the text itself:

> A fictional narrative statement is immune to judgments of truth or falsity; in fiction, [these judgments] are suspended. It is inappropriate to say that a fictional statement is false. Rather, it creates by fiat a fictional reality which can only be taken as fictionally true. [Banfield 1982, p. 258]

In conversation, people talk about a reality to which they have independent access, but, in objective narration, what is narrated is created in the fictional world (Hamburger 1973, Kuroda 1973ab, Banfield 1982).

For example, consider a reader's response to the following narrative statement, assuming that she interprets it to be objective:

(7) Shea Stadium was a large stadium in Chicago.

Even if she knows where Shea Stadium and Chicago are, she will not question the truth of (7), but instead will simply update her model of the fictional world. That is, even if she knows that, in the real world, Shea Stadium is in New York City, from (7) she understands that, in the fictional world being created, Shea Stadium is in Chicago. (Chicago and Shea Stadium may themselves turn out to be completely different in the fictional and the real worlds.)

A reader understands that the propositional content of references in objective contexts is (fictionally) true. This means that the reader does not have to question whether a reference reflects a false belief of a producer (either a producer's own false belief, or a producer's false belief about the reader's beliefs, or about the reader's beliefs about a producer's beliefs about the reader's beliefs, etc.). The reader cannot have private beliefs about what is true in the fictional world that would allow her to disagree with the propositional content of a reference (although she might have beliefs about the fictional world, for example, that it does not correspond to the real world). In contrast to conversation, in which the hearer has independent access to the world under discussion (the real world), the reader of a narrative text does not have any other access to the fictional world than that which is provided by the text.[3] Consider a use of the reference 'the king'. In conversation, it is possible that the hearer believes both that the referent is actually a usurper and that the speaker is expressing a false belief about the referent; it is also possible that the hearer believes that the speaker is expressing the false belief that the hearer believes that the referent is the king, etc. The reader of a narrative text, however, cannot have the private belief that the referent of 'the king' is a usurper, or that a producer has the mistaken belief that the reader believes that the referent is the king, etc. If the propositional content of a reference in an objective context is that a character is the king, then it is simply so.

5 References in Subjective Sentences

As noted above in section 2, subjective sentences reflect the beliefs of the subjective character, which may be false in the fictional world. This is true not only for the main proposition of the sentence, but for the propositional content of references as well. In the following passage, Morgaine is referred to as 'the witch' in Chei's subjective sentence, reflecting his false belief about her.

> (8)
> He [Chei] reckoned even that it was a spell *the witch* had cast over him, that from the time she surprised him with that look into his eyes, from that moment his soul had been snared. [Cherryh, *Exile's Gate*, p. 86, my italics]

As in conversation, the reader does not automatically assume that the propositional content of references in subjective contexts is true; this is in contrast to objective contexts, in which she does automatically assume this. As in objective contexts, however, subjective sentences are not understood with respect to the relationship between a speaker and a hearer; subjective and objective contexts differ from conversation in this respect. A subjective sentence is not directed by the subjective character toward an addressee (Kuroda 1973ab, Banfield 1982). Rather, the pragmatic situation might be thought of as the reader "overhearing" (or being "privy to") the character's thoughts, or the reader "perceiving" the fictional world through the character's senses (or a combination of the two; cf. Brinton 1980). The subjective character is clearly not addressing the reader (the reader does not even exist in the fictional world). The reader may believe that a reference reflects a false belief of the subjective character, and yet she will not try to understand the disparity between her own beliefs and the subjective character's in terms of beliefs about their mutual beliefs (in fact, she does not have any such beliefs). Subjective sentences are *propositionally transparent* (Castañeda 1970, pp. 167ff); they reveal, unobscured by any communicative relationship between the subjective character and the reader, the subjective character's model of the fictional world.

6 Understanding References

6.1 Representation

A single text can contain subjective contexts attributed to different characters, as well as objective contexts. So, simply to understand references in a narrative text, the reader must maintain (1) a model of the objects, places, characters, etc., and their properties that are actually in the fictional world (cf. Webber's (1983) *discourse models*, Kamp's (1984) *discourse representation structures*, and Rapaport's (1988) notion of the "mind" of an AI/natural-language-understanding system), and, (2) for each subjective character, another model of the objects, characters, etc., and their properties that the character believes are in the fictional world (cf. Fauconnier's (1985) *mental spaces*). Where one character's beliefs differ from another character's beliefs, or from what is true in the fictional world, the models are distinct, and where they correspond to one another, the models overlap.

Section 5 noted the most obvious situation for which these models must be maintained to understand references in subjective contexts, namely, when a character has a false belief about an individual (see passage (8)). Another situation in which the reader's and characters' beliefs about the fictional world must be distinguished is when the reader (and perhaps other characters) knows more than the subjective character knows. For example:

(9)

Perhaps the man understood. Perhaps he did not. [Cherryh, *Exile's Gate*, p. 24]

These sentences are the represented thoughts of the character Vanye about Chei. While the reference 'the man' does not reflect a false belief of Vanye's about Chei (Chei really is a man), it does reflect Vanye's limited knowledge about Chei: At this point in the novel, the reader, but not Vanye, knows Chei's name.

6.1.1 Representation of Proper Names

This and related work (Rapaport 1986; Wiebe and Rapaport 1986; Rapaport, Shapiro, and Wiebe 1986) shows that when representing an individual in a natural-language-understanding system, it is important to represent the individual in a way that is neutral with respect to any properties ascribed to it; in particular, the individual must be represented independently of attributes used to refer to it, even a proper name. For example, the reader may know that a character has two names and that some of the other characters know only one of the names. In Robert Ludlum's *The Parsifal Mosaic*, for example, the characters at the State Department know an individual by the name 'Arthur Pierce', but the reader knows that he is a Russian spy whose real name is 'Nikolai Petrovich Malyekov'. The individual cannot be represented simply as Arthur Pierce, or as Nikolai Petrovich Malyekov, but has to be represented as an individual whom the reader believes has both names, but whom the reader believes some of the other characters believe has only the name 'Arthur Pierce'. References to this individual can reflect these different beliefs that the reader herself holds and that she believes the other characters hold (in objective and subjective contexts, respectively).

6.1.2 Belief Spaces

The reader's and characters' models of the fictional world can be represented by *belief spaces* (as in Rapaport's (1986) computational model of *de re* and *de dicto* belief reports). A belief space is labeled by a stack, and represents those things that the bottom member of the stack believes that ... the top member of the stack believes. The reader is always the bottom member of the stack, reflecting the fact that other agents' beliefs are not directly represented, but rather what the reader believes that those other agents believe.

The belief space $BS[\textit{the reader}, a_1, a_2, \ldots, a_{n-1}, a_n]$ consists of the propositions P such that the reader believes that a_1 believes that a_2 believes that $\ldots$ a_{n-1} believes that a_n believes that P, together with the propositions that follow from the propositions P. As we shall see in section 6.2, the important belief spaces in this paper are $BS[\textit{the reader}]$ and $BS[\textit{the reader}, c]$, where c is a character. The former will be referred to as the reader's belief space, and the latter will be referred to as c's belief space.

Shapiro and colleagues (Maida and Shapiro 1982; Rapaport 1985; Shapiro and Rapaport 1987) have argued that the terms of an AI representation language, if it is to be used to represent the mind of a cognitive agent, should be interpreted as *intensional* entities (in particular, as Meinongian objects of thought (Meinong 1904)). Under this interpretation, there is not necessarily a one-to-one correspondence between entities actually in the world and concepts of those entities. For example, the cognitive agent might believe that two concepts are actually concepts of the same individual, that is, that they are *co-extensional*. Once the cognitive agent holds such a belief, then anything that she believes about one of the concepts she simultaneously believes about the other. Since we are concerned with nested beliefs, we need the following generalization: if the belief that two concepts are co-extensional appears within a belief space, then anything believed about one of them in that belief space is also believed about the other in that belief space.[4]

Note that a single proposition may appear in more than one belief space. An example is the situation from *The Parsifal Mosaic* described in section 6.1.1: although the proposition that the individual in question is named 'Nikolai Petrovich Malyekov' appears in the reader's belief space and not in the belief spaces of the characters at the State Department, the proposition that he is named 'Arthur Pierce' does appear in all of these belief spaces.

6.2 The Current Belief Space and References

In this section, we consider reference resolution in light of the observations made in sections 4 and 5 about references in subjective and objective contexts. The main idea of this section is that references should be understood with respect to the *current belief space* (*CBS*), which is the belief space of the reader while an objective context is being processed, and the belief space of the subjective character while a subjective context is being processed. This explains some interesting referential phenomena in third-person narrative text.

As an overview, we will begin with sr, a specific reference, and p, a predicate corresponding to the propositional content of sr. For example, if sr is 'the car', then p is the predicate car. The referent of sr will be a concept c such that the proposition $p(c)$ is in CBS. For example, if 'the car' appears in a subjective context of John, the referent will be a concept c such that $car(c)$ is in John's belief space.[5] Depending on the sort of reference sr is and upon what is currently believed, $p(c)$ may or may not already be in CBS before sr is resolved. We will first consider specific definite references (sects. 6.2.1 and 6.2.2) and then specific indefinite references (sect. 6.2.3). In order to focus on the specific phenomena I wish to illustrate, I will idealize reference resolution in these sections, ignoring many complexities.

6.2.1 Non-Anaphoric Specific Definite References

In general terms, a specific definite reference refers to an individual who is already known. Examples of definite references are 'Ellen', 'the mirror', and 'she'. Context may be required to identify the referent, so the referent should first be sought among the concepts recently evoked in the narrative (if it is found among them, the reference is *anaphoric*); this will be the topic of section 6.2.2. In the present section, we consider specific definite references that are *non-anaphoric*, that is, whose referents are not recently-evoked concepts.

Let sdr be a specific definite reference and p a predicate corresponding to the propositional content of sdr. Further, assume that the referent of sdr could not be found among the concepts recently evoked in the text. Then sdr should be resolved as follows:

i. If there is a concept c such that the proposition $p(c)$ is in CBS, then the referent should be chosen from among the set of such c.

ii. Otherwise, if there is a concept c such that $p(c)$ is in a belief space other than CBS and $p(c)$ is consistent with CBS, then the referent should be chosen from among the set of such c. Let $c1$ be the the concept chosen. Then $p(c1)$ should be added to CBS.

iii. Otherwise, a new concept c should be created to serve as the referent of sdr, and the proposition $p(c)$ should be added to CBS.

If there are possible referents given only the propositions in CBS, one of them should be chosen as the referent (branch i). In *The Parsifal Mosaic*, for example, a set of characters believe that there is someone named 'Parsifal', but each believes it is a different person. When 'Parsifal' appears in the subjective context of one of these characters, the referent is the person who the subjective character believes is named 'Parsifal'. (There would be a set of possible referents to choose from if more than one person were believed by the subjective character to have the name 'Parsifal'.)

If the test in (i) fails, then other belief spaces should be considered (branch ii). To see why, suppose that sdr is 'Ellen' and that it occurs in a subjective context of John. Since branch (i) was not taken, the belief that someone is named 'Ellen' is not yet in John's belief space. But someone may have been referred to earlier as 'Ellen', and, unless there is some reason to believe that John believes that this person is not named 'Ellen' (in which case, the proposition that she is named 'Ellen' would not be consistent with CBS), then that person should be taken to be the referent. (There are many cases in which there is a set of possible referents to choose from. The simplest one is if, for all c, any proposition $p(c)$ that appears in a character's belief space also appears in the reader's belief space, where p is as above. Then the choice is among the concepts c such that the reader believes that $p(c)$.)

Once we have the referent, $c1$, the proposition $p(c1)$ should be added to CBS. In the above example, the fact that 'Ellen' occurs in John's subjective context indicates to the reader that John believes that the referent is named 'Ellen'.

If the tests of both (i) and (ii) fail, then *sdr* is a definite reference to a new concept (branch iii). This kind of reference is acceptable in both objective and subjective contexts. In conversation, in contrast, a specific definite reference would not be completely acceptable to the hearer if she were not already familiar with the referent.

In objective contexts, a definite reference can be used to introduce something new into the fictional world. An example is 'the jockey' in the following, which is the first mention of the jockey in the story:

> (10)
>
> The jockey came to the doorway of the dining room, then after a moment stepped to one side and stood motionless, with his back to the door. [McCullers, "The Jockey", p. 58]

In a subjective context, a specific definite reference is understood to refer to an individual with whom the subjective character is familiar (Fillmore 1974). This is so regardless of whether or not the reader is already familiar with the referent. For example:

> (11)
>
> [11.1]She [Hannah] winced as she heard them crash to the platform. [11.2]The lovely little mirror that she had brought for Ellen, and the gifts for the baby! [Franchere, *Hannah Herself*, p. 3]

Neither the mirror, the gifts, the baby, nor Ellen have been mentioned before in the novel (the pronoun 'them' in (11.1) is an anaphoric reference to some luggage mentioned in the previous sentence). For each, a new concept is created to serve as the referent, and an appropriate proposition is added to Hannah's belief space.[6] That is, by virtue of understanding the references in (11.2), the reader comes to believe that Hannah (the subjective character) believes that there is a mirror, some gifts, a baby, and a person named 'Ellen'.

6.2.2 Anaphoric References

An anaphoric reference accesses a concept recently evoked in the narrative (typically, pronouns and definite descriptions such as 'the jockey' are anaphoric). We will first briefly consider anaphor resolution in computational linguistics (ignoring many details),[7] and then consider the role of CBS in this process.

The clearest case of anaphoric reference, the only one we will consider here, is when an anaphor co-refers with a noun phrase that explicitly appeared earlier in the text (called its *antecedent*).[8]

So that the referent of an anaphor can be recovered from the previous context, an ordered list of previously mentioned concepts is maintained (the *history list*). To find the referent of an anaphor, this list is searched for possible referents, starting with the first element. (Although recency of mention should partly determine how the elements of the list are ordered, other factors should also be taken into account. Cf., e.g., Sidner 1983.)

When any sort of reference is encountered (not only anaphoric ones), once the referent has been found, it is added to the history list in anticipation of later anaphoric references to it. If it already appears on the list, then the old entry is deleted and a new one is added.

There are syntactic constraints governing the relationship between an anaphor and its antecedent. For example, they must agree in number and gender. Thus, when a concept is added to the history list, any syntactic properties of the reference that may be relevant for resolving later anaphoric references should be included in the list entry.

Now consider finding the referent of an anaphoric reference. Let c be a concept on the history list, ar an anaphoric reference, and p a predicate corresponding to the propositional content of ar. Then c should be ruled out as a possible referent for ar if the syntactic properties stored with c on the history list and the syntactic properties of ar do not agree. (There may also be syntactic constraints from the sentence in which ar appears, such as those on the use of reflexive pronouns. In addition, of course, c is acceptable only if the sentence would make sense if c were the referent.) In addition, c might be rejected on the basis of the content of ar. It is here that CBS should be consulted—c should be ruled out if $p(c)$ is inconsistent with the propositions in CBS. However, if $p(c)$ is consistent with CBS (and no other constraints are violated), then we can accept c as the referent, even if $\neg p(c)$ is in another belief space. Consider the following example:

(12)

[12.1]His [Dwayne's] brain worked slowly through what he knew about this person [Casey]. [12.2]David's kid. [12.3]The name stumbled into place. [12.4]This was David's boy. [12.5]David was in the war, and here was his kid in the arcade scared of something. [12.6]He wasn't sure of what. [12.7]What in the arcade could scare a boy like that? [12.8]He rubbed his head under his baseball cap. [12.9]He could see tears in Casey's eyes. [12.10]He could tell they were tears because *his* eyes were too shiny. [12.11]Too round. [12.12]Well, it was all right to cry. [12.13]He'd cried when they took him to that place a few years back. [12.14]Now Casey was in a new place, too, feeling maybe the same as him. [12.15]If he just knew what to do about it.

[12.16]"Let's don't play that game anymore," he said. [12.17]"I don't like that one."

[12.18]Casey wiped *her* face on her sleeve ... [Bridgers, *All Together Now*, pp. 91-92, my italics]

The interesting references in this passage are 'his' in (12.10) and 'her' in (12.18)—both refer to Casey, even though one is masculine and one is feminine. Because 'Casey' can be either masculine or feminine, it is syntactically permitted to be the antecedent of both masculine and feminine references (but not of ones such as 'it').

Sentences (12.1)-(12.15) are Dwayne's subjective sentences, and (12.16)-(12.18) are objective. Although Dwayne correctly believes that Casey's name is 'Casey', he incorrectly believes that Casey is a boy. Suppose $c1$ is the concept of Casey. When 'Casey' in (12.9) is understood to refer to $c1$, an entry for $c1$ is made in the history list. When 'his' in (12.10) is encountered, CBS is Dwayne's belief space. Because $male(c1)$ is consistent with CBS, $c1$ is not ruled out as a possible referent of 'his', even though $\neg male(c1)$ is in the reader's belief space. Similarly, an entry for $c1$ is made when 'Casey' is resolved in (12.18). Since (12.18) is objective, CBS is the reader's belief space. Thus, $c1$ is not ruled out as a possible referent of 'her' in (12.18), because $female(c1)$ is consistent with CBS.

Suppose that a concept c is taken to be the referent of an anaphoric reference ar (that is, c is the first, or only, concept on the history list that is not ruled out as a possible referent).

Since $p(c)$, the propositional content of ar, need only be consistent with CBS, it may not already be in CBS. If $p(c)$ is not already in CBS, it should be added to it. The reason is the same as the one given for (iii) in section 6.2.1.

The question arises as to how long entities should remain in focus. In computational linguistics, this question has been related to discourse structure.[9] (The term 'discourse' in this context is neutral with respect to genre. It is a text or conversation considered as a structured whole.) It has been suggested that a discourse is composed of hierarchically-related *discourse segments*, such that within a segment, sentences are *locally coherent*. One aspect of this coherence is that within a discourse segment, anaphoric references can be resolved to concepts evoked in that segment, but once a segment has ended, those concepts are no longer available for anaphoric reference (but note that there can be *discourse returns* to segments, and *semantic returns* to concepts evoked in previous segments (Allen 1987)). As a consequence, each segment should have its own local history list.

Some phenomena I have observed in narrative suggest that each subjective and objective context is a discourse segment composed of sub-segments. Thus, following the approach outlined in the previous paragraph, each subjective or objective context has its own local history list (as does each sub-segment within a context) (Nakhimovsky 1988, Nakhimovsky and Rapaport 1988). Further, there may be a level of discourse structure at which a discourse segment can include more than one context; because such a segment would have its own history list, an entity could remain in focus across context boundaries. This would explain cases I have observed in which (1) a subjective context is followed by an objective one, (2) the subjective character of the subjective context is referred to by a pronoun at the beginning of the objective context, and (3) there is a reference to someone else in the subjective context that is a possible antecedent of the pronoun, and this reference is more recent than the last reference to the subjective character. An illustration is 'he' in (12.16). The following modified version of (12) is a more convincing illustration:

(13)

Dwayne wasn't sure what John was scared of. What in the arcade could scare a boy like that? He rubbed his head under his baseball cap. He could see tears in John's eyes. He could tell they were tears because his eyes were too shiny. Too round. Well, it was all right to cry. He'd cried when they took him to that place a few years back. Now John was in a new place too. Maybe that was why he was crying.

"I want to leave", he said.

Even though John is the last-mentioned male entity, the referent of 'he' in the last sentence is taken to be Dwayne, the subjective character of the first paragraph.

I am currently investigating this view of the discourse structure of narrative text and also the related processing issues, such as how it can be determined whether or not a new context is part of a larger discourse segment.

6.2.3 Specific Indefinite References

A specific indefinite reference, in general terms, refers to an individual who is not already known. In a subjective context, such a reference is understood to refer to an individual with whom the subjective character is not familiar (Fillmore 1974). Interestingly, even if the referent is actually an individual with whom the reader is familiar, the author must still use an indefinite reference—a definite reference would imply that the subjective character, too, is familiar with the referent (as discussed in section 6.2.1). In a subjective context, in general, whether a specific reference is definite or indefinite is determined by the familiarity of the subjective character with the referent, regardless of whether or not the reader is familiar with it.

The following is a strategy for understanding specific indefinite references in third-person narrative text (a situation in which another strategy might be better will be identified below):

> Let sir be a specific, indefinite reference and p a predicate corresponding to the propositional content of sir. Then a new concept c should be created to serve as the referent of sir, and the proposition $p(c)$ should be added to CBS.[10]

The reader's believing that the referent of sir is actually an individual she is already familiar with can be represented by including, in the belief space of the reader, the proposition that c and $c1$ are co-extensional, where $c1$ is the concept of that individual that already existed.

Suppose that sir appears in a subjective sentence. The reader may decide that the referent is actually an individual she is already familiar with after reading later sentences. For example, suppose that the following is a passage appearing in the middle of a novel:

(14) Zoe looked up. A man was coming toward her. My, he had shocking red hair.

Suppose that in earlier objective sentences, a man was referred to as 'John' and was described as having very red hair. Thus, there are already propositions in the reader's belief space that someone is named 'John' and has very red hair; let $c1$ be the concept of that person. Under the above strategy, a new concept c is created to serve as the referent of 'A man', and the proposition $man(c)$ is added to Zoe's belief space. After reading the third sentence, the reader may come to believe that c and $c1$ are co-extensional, given the clue that the man has shocking red hair.

However, suppose that it is at the time 'A man' is processed that the reader decides that the referent is actually John. This is a situation in which the above strategy may not be the best one. That is, rather than creating a new concept c to serve as the referent and then having to add to the reader's belief space the proposition that $c1$ and c are co-extensional, we may simply want to add $man(c1)$ directly to Zoe's belief space.

Another interesting feature of specific, indefinite references in subjective contexts is that they may imply that the subjective character is not sure what sort of thing the referent is. This may occur when the head noun is a superordinate term, such as 'plant' or 'vehicle', rather than either a basic-level term, such as ‘flower’ or ‘car’, or a subordinate term, such as ‘violet’ or ‘station wagon’ (Rosch and Lloyd 1978). Often, it should be understood from the use of the superordinate term that the subjective character cannot classify the referent at the basic level. For example:

(15)

Slowly Hannah raised her head and blinked her eyes. Small dots of purple covered the ground around her and she reached out to explore. Violets! [Franchere, *Hannah Herself*, p. 25]

When she first sees the violets, Hannah can identify them only as small dots of purple.

There are other reasons to use a superordinate term in a specific indefinite reference. For example, a superordinate term is used when the reference is plural and refers to a group composed of individuals in different basic-level categories (Murphy and Wisniewski 1989). In any event, the fact that a superordinate term was used is a source of information about the subjective character’s beliefs about the referent or referents. (Peters and Shapiro 1987ab and Peters, Shapiro, and Rapaport 1988 describe a representation for natural category systems in which the fact that something was classified only at a superordinate level can be expressed.)

The following passage illustrates the observations made in this section.

(16)

There they [the King and his men] saw close beside them a great rubbleheap; and suddenly they were aware of two small figures lying on it at their ease, grey-clad, hardly to be seen among the stones. [Tolkien, *The Two Towers*, p. 206]

The King and his men have come upon two hobbits, Merry and Pippin. The King and his men do not know the hobbits, but other characters also present in the scene do know them. In a subjective sentence attributed to the King and his men, the hobbits are referred to with an indefinite reference, ‘two small figures’, reflecting the fact that the King and his men do not know them. In addition, the superordinate term ‘figures’ reflects the fact that the King and his men do not know what the referents are (‘hobbit’ is a basic-level term in this novel; the kinds of sentient beings are hobbits, dwarves, elves, men, etc.).

Suppose that the reader does not decide until after processing 'two small figures' that it is Merry and Pippin whom the King and his men have come upon. Once 'two small figures' is processed using the strategy given above, the following are some of the propositions about Merry and Pippin that appear in various belief spaces. The belief spaces of the reader and of the characters in the scene other than the King and his men, include the following propositions about two concepts $c1$ and $c2$: $c1$ is named 'Merry', $c2$ is named 'Pippin', $c1$ is a hobbit, and $c2$ is a hobbit. The belief space of the King and his men include the following propositions about two other concepts $c3$ and $c4$: $c3$ is a small figure, and $c4$ is a small figure. When the reader decides that the referents of 'two small figures' are in fact Merry and Pippin, two more propositions appear in the reader's belief space: that $c1$ and $c3$, and $c2$ and $c4$, are concepts of the same individual.

6.3 Conclusion

This section has identified some consequences for reference resolution of the fact that although references in objective sentences reflect true information about the fictional world, those in subjective contexts reflect the beliefs of the subjective character.

7 De Re and De Dicto Reports

7.1 Introduction

As stated in the introduction, this work has grown out of the work in philosophy, linguistics, and AI on *de re* and *de dicto* belief reports. This section considers the *de re/de dicto* ambiguity with respect to some issues raised in earlier sections.

The terms '*de re*' and '*de dicto*' have been used in different ways; here, I use them to describe an ambiguity to which particular types of sentences, those with opaque contexts, are subject. (Another term for this is *ambiguity with respect to description* (Fodor 1979)). Recall from section 1 that opaque contexts are the objects of sentences with third-person subjects and main verbs that are either propositional-attitude verbs, such as 'believe', 'know', and 'hope', or other psychological and perceptual verbs, such as 'want', 'hate', and 'see' (we will use the term 'private-state' for such verbs). Below, we will only consider such sentences that have 'believe' as the main verb (although similar comments apply to those with other private-state verbs), and, following usual practice, we will call them 'third-person belief reports'.[11]

The ambiguity is generally viewed as follows. In a *de re* belief report, the speaker herself believes the propositional content of references in the subordinated clause; in a *de dicto* report, the speaker believes that the subject of the sentence believes the propositional content of those references. Consider (17):

(17) John believes that the morning star is beautiful.

If (17) is *de re*, then 'the morning star' reflects the speaker's belief that the referent is the morning star; if it is *de dicto*, then 'the morning star' reflects the speaker's belief that John believes that the referent is the morning star. In the first case, the speaker may not believe that John believes that the referent is the morning star; in the second case, the speaker may not herself believe that the referent is the morning star.[12]

Analyses of belief reports do not generally consider whether genre affects how reports are interpreted. In section 7.2, I suggest that analyses of *de re* reports in conversation should take into account the mutual beliefs of the speaker and hearer, and I contrast this with reports in objective contexts. In section 7.3, following Janet Fodor (1979), I suggest that "pure" *de dicto* interpretations are unlikely in conversation; given this and the observations made in previous sections, in section 8 I draw parallels between the interpretation of references and the interpretation of *subjective elements* (linguistic elements that express attitudes and affects).

7.2 The De Re Interpretation and Mutual Belief

As discussed in section 3, references in conversation are understood to reflect the speaker's beliefs about the hearer's and speaker's mutual beliefs; they are not always understood to simply reflect the speaker's beliefs about the entities referred to. I suggest in this section that this extends as well to references in the subordinated clauses of *de re* belief reports.

We will first reconsider references in conversation from the speaker's perspective. Even in non-opaque contexts, the speaker's beliefs may not in fact be reflected by her references. For example, because of the communicative situation, a speaker may use a reference whose propositional content the speaker herself does not believe, but which the speaker believes the hearer believes. One example, (5), was given above in section 3. Donnellan (1966) provides another example (pp. 290-291):

> (18)
>
> Suppose the throne is occupied by a man I firmly believe not to be the king but a usurper. Imagine also that his followers as firmly believe that he is the king. Suppose I wish to see this man. I might say to his minions 'Is the king in his countinghouse?'. I succeed in referring to the man I wish to refer to without myself believing that he fits the description.

The same can be true of a reference in the object of a *de re* belief report. That is, given the above scenario, I might utter the following, even if I do not myself believe the propositional content of 'the king'; I might do so if I believe that my hearers believe the propositional content of this reference:

(19) The queen believes that the king is in his countinghouse.

Now let us take the hearer's perspective. Even in the objects of *de re* belief reports, the hearer does not always understand references as simply reflecting the speaker's beliefs about the referent. Consider understanding (19). If the hearers don't know that I believe that the man on the throne is a usurper, then they will not detect my "dishonesty", and they will understand the reference as reflecting my belief that the man is the king. However, suppose that they do know that I believe that he is a usurper. Then they understand 'the king' as reflecting my belief that they believe that the man is the king, my belief that they believe that I believe that he is the king, etc., even though they do not believe that I believe that he is the king. In this case, the hearers do not understand the reference as reflecting my belief in the propositional content of the reference. Thus, in conversation, even references in the objects of *de re* belief reports are subject to the sorts of complexities that were identified by Clark and Marshall and by Perrault and Cohen.

Now consider (syntactic) private-state reports in objective contexts. One sort that can be objective is a report with a negated factive verb (such as 'know' or 'realize') and a propositional object. For example:

(20) John did not know that Mary was in the cell next to his.

Either this sentence is another character's reflection about John's lack of knowledge, or it is an objective sentence. Let us assume the latter. As with other references in objective contexts, the reader understands that the propositional content of the reference in the object of the report is true; the fact that the sentence is syntactically a knowledge report is irrelevant. The *de re/de dicto* ambiguity lies in whether or not this reference also reflects the beliefs of the subject of the report; if it is understood that it does, then the report is interpreted to be *de dicto*. Otherwise, it is interpreted to be *de re*. In either case, the propositional content of the reference is understood to be true.

7.3 De Dicto Reports and the Speaker's Beliefs

In my experience, it is difficult to find clear cases of "pure" *de dicto* reports in natural conversation. (A "pure" *de dicto* report is one in which a reference does not reflect the speaker's beliefs about the referent, but only the speaker's beliefs about the subject of the report's beliefs about the referent.) Janet Fodor (1979) argues that the speaker is always responsible for references, even in the objects of belief reports interpreted *de dicto*—that is, the speaker cannot use a description in the subordinated clause of a belief report that the speaker does not believe correctly describes the referent, even if the speaker believes that the subject believes that it does. The difference between a *de dicto* and a *de re* belief report, she claims, is that in the former, but not the latter, the speaker believes that the subject of the report believes the propositional content of references, too. (Note that Fodor, whose work predates Clark and Marshall (1981) and Perrault and Cohen (1981), does not consider the communicative relationship between the speaker and the hearer, so she does not consider, for example, the possibility that a reference reflects the speaker's beliefs about the hearer's beliefs.)

Fodor makes her argument by considering examples in which the speaker and subject disagree about what would be a correct description of the referent. Examples in which they agree would not reveal anything, she argues, because then there would be no way of determining who is responsible for the reference (p. 250).

I will not present her entire argument here, but just show her analysis of an example from Postal 1967.

(21) Charley believes that the book which *was* burned was *not* burned.

Both *de re* and *de dicto* interpretations of (21) ascribe a false belief to Charley.

By the definition of a *de re* report, the speaker believes the propositional content of 'the book which *was* burned'; thus, what the speaker disagrees with must be Charley's belief that the book was *not* burned.

A *de dicto* interpretation of (21) ascribes a contradictory belief to Charley—by the definition of a *de dicto* report, Charley believes the propositional content of the reference 'the book which *was* burned', and the main assertion of the sentence is that Charley believes that the book was *not* burned. It is Fodor's contention that a speaker cannot reasonably utter (21) unless the speaker herself believes the propositional content of 'the book which *was* burned', and, therefore, as in the *de re* case, what the speaker disagrees with must be Charley's belief that the book was *not* burned, not Charley's belief that the book *was* burned (p. 251). Fodor's evidence is that a speaker could reasonably utter (22) but not (23):

(22) Charley believes that the book which was burned was not burned. Silly old Charley. Of course it was burned. [p. 255]
(The speaker disagrees with the propositional object of 'believe'.)

(23) Charley believes that the book which was burned was not burned. Silly old Charley. Of course it wasn't burned. [p. 255]
(The speaker disagrees with the propositional content of the reference.)

"Whichever opinion as to the state of the book is expressed by the content of the noun phrase [in the object of the report] is the one to which the speaker is committed. He may be ascribing it to Charley, but he is also, in using it in his report, endorsing it himself" (p. 256). If the speaker has no opinion either way about the book, then neither the reference 'the book which was burned' nor 'the book which was not burned' can be used. Thus, (24) and (25) are both inconsistent:

(24) Charley assumed that the book which was burned was not burned, but I don't know whether it was burned or not. [p. 256]
(Speaker has no opinion about the propositional content of the reference.)

(25) Charley assumed that the book which was not burned was burned, but I don't know whether it was burned or not. [p. 256]
(Speaker has no opinion about the propositional content of the reference.)

We will now consider examples that are similar to the ones above, except that in these, a *de dicto* reading does not ascribe a contradictory belief to the subject. Even without the contradiction, we find the same phenomena. Thus, the speaker could reasonably utter (22′), but not (23′) or (24′):

(22′) Charley believes that the book which was burned was written by Shakespeare. Silly old Charley. It was obviously written by Burns. [p. 255]
(The speaker disagrees with the propositional object of 'believe'.)

(23′) Charley believes that the book which was burned was written by Shakespeare. Silly old Charley. Of course it wasn't burned. [p. 255]
(The Speaker disagrees with the propositional content of the reference.)

(24′) Charley believes that the book which was burned was written by Shakespeare, but I don't know whether it was burned or not. [p. 256]
(The speaker has no opinion about the propositional content of the reference.)

In contrast to (24′), (25) is fine, since it is the truth of the propositional object of 'believe' that the speaker has no opinion about:

(25) Charley believes that the book which was burned was written by Shakespeare, but I don't know whether he wrote it or not. [p. 256]
(No opinion as to the propositional object of 'believe'.)

Returning to third-person narrative text, there is no question that "pure" *de dicto* reports appear in subjective contexts. An example is (8) above.

8 Parallels with the Interpretation of Subjective Elements

We have seen that the language-use context of a reference (i.e., whether it occurs in a subjective context, an objective context, or in conversation) affects how the reference is understood. This is not an isolated issue. There are lexical and grammatical elements whose interpretations are similarly affected by language-use context—those that express *subjectivity*, as analyzed by Banfield (1982).

Banfield equates subjectivity with that which is expressed by *subjective elements*. Examples of subjective elements are exclamations, as in (15), which express emotion, and epithets, such as 'the idiot', which express evaluation of the referent. In any context, Banfield argues, the subjectivity expressed by such an element is attributed to the "subject of consciousness", or the "SELF". In conversation (or any context in which 'I' could appear), the SELF is the speaker, so the subjectivity expressed by subjective elements is attributed to the speaker.[13] Banfield shows, however, that subjectivity is not tied by definition to the speaker. In a narrative subjective sentence, the subjective character is the SELF, even when she is referred to in the third person. Consider the following sentence:

(26) The idiot was standing next to her.

If uttered in conversation, 'the idiot' expresses the speaker's subjectivity. But if (26) is a narrative subjective sentence, then 'the idiot' expresses the subjectivity of the subjective character (who is potentially the referent of 'her').

Given Fodor's contention that the speaker is always responsible for references, and the fact that references in subjective sentences reflect the beliefs of the subjective character, we can relate the issue addressed in this paper to Banfield's theory as follows: belief in the propositional content of a reference is attributed at least to the SELF (in the case of a belief report, the SELF may believe that the subject of the sentence believes it as well).[14] As mentioned in section 7.3, Fodor does not consider the communicative relationship between the speaker and hearer, so she does not consider the possibility that, for example, a reference reflects the speaker's beliefs about the hearer's beliefs. Thus, the above statement needs to be revised to account for this relationship. However, one might consider the following to be an example of an analogous phenomenon (note that a kinship term such as 'Daddy' is another type of subjective element): my uttering (27) with the intention of referring to the child's father rather than to my own.

(27) Daddy will be here to pick you up soon.

Thus, it may be that the communicative relationship between the speaker and hearer also potentially affects the interpretation of subjective elements.

Subjective elements cannot appear in objective sentences. The reason, Banfield argues, is that there is neither a speaker nor any other subject of consciousness in an objective sentence to whom subjectivity can be attributed. The absence of a SELF can also explain the fact that references in objective sentences are simply understood to reflect true information—since there is no SELF to whom to attribute belief, the question of belief in the propositional content does not arise.

9 Conclusion

I conclude by mentioning some open questions related to those addressed in this paper. First, in section 6, it is assumed that we are given fully fleshed-out belief spaces; given these, and the sort of reference to be understood, section 6 identifies some specific situations in which propositions are added to belief spaces as a result of understanding references. An important question is what strategies should be employed in general to add propositions to belief spaces. (Wilks and colleagues have dealt extensively with this question. Cf., e.g., Wilks and Bien

1983, Wilks and Ballim 1987, Ballim et al. 1991.)

Second, the content of a reference may be important in determining which belief space is the current one. For example, where p is a predicate corresponding to the propositional content of a reference r, and c is a concept, suppose that $p(c)$ is inconsistent with the propositions in CBS, but is a proposition in another belief space. This may suggest that we are wrong as to which belief space is the current one, not that c should be rejected as the referent of r.

Third, section 8 specified that it is in the absence of a speaker that subjectivity and belief in the propositional content of references can be attributed to an agent referred to in the third person. However, I did not define the term 'speaker', as used in that context. I will not attempt to define it here, but merely point out that containing a first-person pronoun is a sufficient but not a necessary condition for a sentence to have a "speaker". Thus, a sentence in an expository text in which the author is referred to as 'I' has a "speaker", even though the sentence is not verbally uttered. But in the absence of a first-person pronoun, what is important is whether a first-person pronoun could potentially appear, without thereby changing to whom subjectivity is attributed. (Banfield shows that deictic elements, tense, and aspect can place constraints on whether or not both a first-person and a second-person pronoun could appear in a sentence.)

Finally, in objective and subjective narrative sentences, the author is not a speaker or even a SELF, in Banfield's terms (note that by 'the author' I mean the actual writer of the text, not an overt narrator; see note 2). Experimental fiction aside, the author cannot be referred to in the first person, and subjectivity cannot be attributed to her. However, the author is clearly communicating with the reader, in some sense of the word 'communicating'. For example, part of understanding a narrative text is to recognize the author's intentions with respect to whether sentences should be interpreted as subjective or objective (Pollack (1986) notes that in narrative, one must infer both the author's and the characters' plans). This aspect of fictional narrative has important implications for computational models of pragmatic processing.

Acknowledgments. I am indebted to Mary Galbraith, Sandra L. Peters, William J. Rapaport, Stuart C. Shapiro, David Zubin, and the other members of the Narrative Research Group of the SUNY Buffalo Center for Cognitive Science and the SNePS Research Group for many discussions and ideas, and to William J. Rapaport, Sandra Peters, Toby Walker, Diane Horton, Peter Heeman, and anonymous referees for valuable comments on earlier drafts of this paper. This research was supported in part by National Science Foundation grants IST-8504713 and IRI-8610517, and the preparation of this paper was supported in part by the Natural Sciences and Engineering Research Council of Canada.

Notes

1 See Wiebe 1990a and the references therein for alternative categorizations of subjective sentences.

2 A text may take the point of view of an <u>overt narrator</u>, to whom beliefs, emotions, etc., are attributed (Chatman 1978). In this paper, texts with overt narrators are not considered. Note that a sentence with an overt narrator is <u>not</u> an objective sentence.

3 Actually, the author relies on the hearer bringing to the fictional world the knowledge she has about the real world. However, the reader accepts any discrepancies, and updates her model of the fictional world accordingly.

4 Beliefs that two concepts are co-extensional, especially nested beliefs of that sort, clearly pose hard problems for knowledge representation (cf., e.g., McCarthy 1979, Moore 1980, Asher 1986, Konolige 1986, Rapaport 1986, Shapiro and Rapaport 1987, and Maida, forthcoming, for some relevant work in knowledge representation). Since this paper is not about belief representation per se, the issues involved will not be discussed here.

5 There may be more than one predicate, $p_1, \ldots, p_n$, that together correspond to the propositional content of the reference. For example, if sr is 'the little red car', then there are three predicates, $little$, red, and car. In this case, the referent of sr will be a concept c such that $p_1(c) \wedge \ldots \wedge p_n(c)$ is in CBS. One sort of complexity we are ignoring are interactions among components of a noun phrase. For example, the referent of 'a fire alarm' is not something that is believed to be a fire and believed to be an alarm (cf. Hobbs et al. 1988 for a computational treatment of compound nominals).

6 We are considering only a small subset of the reader's beliefs. Some we are not considering are beliefs about the sense in which the referent is familiar to the subjective character and beliefs about the temporal extent of this familiarity.

A new concept that is the referent of a definite reference in a subjective context is new in the sense that it is created to serve as the referent when the reference is processed by the reader. However, the reader believes that, with respect to the current moment within the fictional world, the concept is not new to the subjective character.

7 Hirst 1981 provides an extensive survey of computational linguistics work on anaphor resolution, and Allen 1987 gives an excellent synthesis of the computational results.

8 An anaphoric reference may refer to an entity that is merely related to a previously-mentioned entity. For example: "John tried the lamp, but the lightbulb was burned out." One understands that 'the lightbulb' refers to the lightbulb of the previously-mentioned lamp. There are many similar sorts of complications; cf., e.g., Webber 1983 and Sidner 1983.

9 The ideas in this paragraph are due mainly to Reichman 1985 and Grosz and Sidner 1986. Two earlier works that influenced theirs are Grosz 1977 and Linde 1979.

10 Suppose that a specific, indefinite reference appears in a subjective context of Mary, and a concept c is created to serve as the referent. An interesting question is what the reader should believe about the temporal extent, within the fictional world, of Mary's unfamiliarity with c. In the following passage, the subjective character is familiar with the referent before the current moment in the fictional world: "Mary was curious. A man had been sitting in a car across the street, watching her house, for two hours." The current moment is the time at which Mary is curious. However, Mary has been aware of the referent of 'A man' for at least two hours. (I owe this point to an anonymous referee.)

11 In this section, 'report' refers to a syntactic form of sentences. However, I should point out that 'report' is also used in this work to refer to a type of subjective sentence (see section 2), and not all subjective sentences that are syntactically reports are interpreted to be the type of subjective sentence called a 'report'; cf. Wiebe 1990b.

12 According to Quine (1976), *de dicto* reports are referentially opaque and *de re* reports are referentially transparent—substitution of co-referential expressions into the opaque

context of the former does not necessarily preserve truth value, while substitution into the latter does. According to Castañeda (1970), *de dicto* reports are propositionally transparent and *de re* reports are propositionally opaque—in the former, the speaker believes that references in the opaque context convey the content of the believer's belief, but, in the latter, the speaker may not believe this.

13 This is somewhat a simplification of Banfield's theory. For example, see her *revised principle of the priority of the speaker*, developed to account for complexities that arise in the interpretation of *echo questions* and a type of subjective sentence not considered in this paper, *represented speech* (Banfield 1982, p. 132).

14 Even when the SELF is a subjective character, a reference in a (syntactic) belief report might reflect the SELF's belief that the subject of the sentence believes the propositional content of a reference as well. Suppose that the following sentence is interpreted to be the thought of a character Rick:

> John believed that Mary was rich.

In this case, Rick, not John, is the subjective character, so all references reflect Rick's beliefs. However, the sentence may be interpreted to be either *de dicto* or *de re* from the perspective of Rick. Under the former interpretation, but not the latter, Rick believes that John, too, believes the propositional content of 'Mary'. On the other hand, if the sentence is interpreted to be the type of subjective sentence called a "private-state report" (see note 11 and Wiebe 1990b), the sentence is a report of John's belief, John is the subjective character, and, since the subject of the report believes the propositional content of 'Mary', the sentence is simply *de dicto*.

References

1. Allen, James F. (1987), *Natural Language Understanding* (Menlo Park, CA: Benjamin/Cummings).

2. Asher, Nicholas (1986), "Belief in Discourse Representation Theory," *Journal of Philosophical Logic* 15: 127–189.

3. Ballim, Afzal; Wilks, Yorick; & Barnden, John (1991), "Belief Ascription, Metaphor, and Intensional Identification," *Cognitive Science* 15: 133–171.

4. Banfield, Ann (1982), *Unspeakable Sentences: Narration and Representation in the Language of Fiction* (Boston: Routledge & Kegan Paul).

5. Brinton, Laurel (1980), " 'Represented Perception': A Study in Narrative Style," *Poetics* 9: 363–381.

6. Castañeda, Hector-Neri (1970), "On the Philosophical Foundations of the Theory of Communication: Reference," *Midwest Studies in Philosophy* 2 (1977): 165–186.

7. Chatman, Seymour (1978), *Story and Discourse: Narrative Structure in Fiction and Film* (Ithaca, NY: Cornell University Press).

8. Clark, Herbert H., & Marshall, Catharine R. (1981), "Definite Reference and Mutual Knowledge," in A. Joshi, B. Webber, & I. Sag (eds.), *Elements of Discourse Understanding* (Cambridge: Cambridge University Press): 10–63.

9. Cohen, Philip R. (1978), "On Knowing What to Say: Planning Speech Acts," *Technical Report 118* (Toronto: University of Toronto Dept. of Computer Science).

10. Cohen, Philip R.; Perrault, C. Raymond; & Allen, James F. (1982), "Beyond Question Answering," in W. Lehnert & M. Ringle (eds.), *Strategies for Natural Language Processing* (Hillsdale, NJ: Lawrence Erlbaum): 245–274.

11. Cohn, Dorrit (1978), *Transparent Minds: Narrative Modes for Representing Consciousness in Fiction* (Princeton: Princeton University Press).

12. Doležel, Lubomir (1973), *Narrative Modes in Czech Literature* (Toronto: University of Toronto Press).

13. Donnellan, Keith S. (1966), "Reference and Definite Descriptions," *Philosophical Review* 60: 281–304.

14. Fauconnier, Gilles (1985), *Mental Spaces: Aspects of Meaning Construction in Natural Language* (Cambridge, MA: MIT Press).

15. Fillmore, Charles (1974), "Pragmatics and the Description of Discourse," in C. Fillmore, G. Lakoff, & R. Lakoff (eds.), *Berkeley Studies in Syntax and Semantics I* (Berkeley: University of California at Berkeley, Dept. of Linguistics and Institute of Human Learning): V1–V21.

16. Fodor, Janet Dean (1979), *The Linguistic Description of Opaque Contexts* (New York: Garland Publishing). Originally author's Ph.D. dissertation (MIT, 1970).

17. Grosz, Barbara J. (1977), "The Representation and Use of Focus in a System for Understanding Dialogs," *Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI-77, MIT)* (Los Altos, CA: Morgan Kaufmann): 67–76.

18. Grosz, Barbara J., & Sidner, Candace L. (1986), "Attention, Intentions, and the Structure of Discourse," *Computational Linguistics* 12: 175–204.

19. Hamburger, Käte (1973), *The Logic of Literature*, M. J. Rose (trans.) (Indiana: Indiana University Press).

20. Hirst, Graeme (1981), *Anaphora in Natural Language Understanding: A Survey* (Berlin: Springer-Verlag).

21. Hobbs, Jerry R.; Stickel, Mark; Martin, Paul; & Edwards, Douglas (1988), "Interpretation as Abduction," *Proceedings of the 26th Annual Meeting of*

the Association for Computational Linguistics (SUNY Buffalo) (Morristown, NJ: Association for Computational Linguistics): 131–138.

22. Kamp, Hans (1984), "A Theory of Truth and Semantic Representation," in J. Groenendijk, T. M. V. Janssen, & M. Stokhof (eds.), *Truth, Interpretation and Information* (Dordrecht: Foris): 1–41.

23. Konolige, Kurt (1986), *A Deductive Model of Belief* (Los Altos, CA: Morgan Kaufman).

24. Kuroda, S.-Y. (1973a), "Where Epistemology, Style and Grammar Meet: A Case Study from the Japanese," in P. Kiparsky & S. Anderson (eds), *A Festschrift for Morris Halle* (New York: Holt, Rinehart & Winston): 377–391.

25. Kuroda, S.-Y. (1973b), "Reflections on the Foundations of Narrative Theory—From a Linguistic Point of View," in T. A. van Dijk (ed.) *Pragmatics of Language and Literature* (Amsterdam: North-Holland): 107–140.

26. Linde, C. (1979), "Focus of Attention and the Choice of Pronouns in Discourse," in T. Givon (ed.) *Syntax and Semantics*, Vol. 12 (New York: Academic Press).

27. Maida, Anthony S. (forthcoming), "Maintaining Mental Models of Agents who have Existential Misconceptions," *Artificial Intelligence*.

28. Maida, Anthony S., & Shapiro, Stuart C. (1982), "Intensional Concepts in Propositional Semantic Networks," *Cognitive Science* 6: 291–330.

29. McCarthy, John (1979), "First-Order Theories of Individual Concepts and Propositions," in J. Hayes, D. Michie, and L. Mikulich (eds.) *Machine Intelligence*, Vol. 9 (New York: Halsted Press).

30. Meinong, Alexius (1904), "The Theory of Objects," English translation by I. Levi et al., in in R. M. Chisholm (ed.), *Realism and the Background of Phenomenology* (New York: Free Press, 1960): 76–117.

31. Moore, Robert C. (1980), "Reasoning about Knowledge and Action," *Technical Report 191* (Menlo Park, CA: SRI International).

32. Murphy, G. L., & Wisniewski, E. J. (1989), "Categorizing Objects in Isolation and in Scenes: What a Superordinate is Good For," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15: 572–586.

33. Nakhimovsky, Alexander (1988), "Aspect, Aspectual Class, and the Temporal Structure of Narrative," *Computational Linguistics* 14: 29–43.

34. Nakhimovsky, Alexander & Rapaport, William J. (1988), "Discontinuities in Narratives," *Proceedings of the 12th International Conference on Computational Linguistics (COLING-88, Budapest)* (Morristown, NJ: Assoc. for Computational Linguistics): 465–470.

35. Perrault, C. Raymond, & Cohen, Philip R. (1981), "It's For Your Own Good: A Note on Inaccurate Reference," in A. Joshi, B. Webber, & I. Sag (eds.), *Elements of Discourse Understanding* (Cambridge: Cambridge University Press): 217–230.

36. Peters, Sandra L., & Shapiro, Stuart C. (1987a), "A Representation for Natural Category Systems: I," *Proceedings of the 9th Annual Conference of the Cognitive Science Society (Seattle)* (Hillsdale, NJ: Lawrence Erlbaum Associates): 379–390.

37. Peters, Sandra L., & Shapiro, Stuart C. (1987b), "A Representation for Natural Category Systems: II," *Proceedings of the 10th International Joint Conference on Artificial Intelligence (Milan)* (Los Altos, CA: Morgan Kaufmann): 140–146.

38. Peters, Sandra L.; Shapiro, Stuart C.; & Rapaport, William J. (1988), "Flexible Natural Language Processing and Roschian Category Theory," *Proceedings of the 10th Annual Conference of the Cognitive Science Society (Montreal)* (Hillsdale, NJ: Lawrence Erlbaum Associates): 125–131.

39. Pollack, Martha E. (1986), "Inferring Domain Plans in Question-Answering," *Technical Note 403* (Menlo Park, CA: SRI International).

40. Postal, Paul M. (1967), "Restrictive Relatives and Other Matters ...," unpublished mimeograph, cited in Fodor 1979. Later published as Postal, Paul M. (1974), "On Certain Ambiguities," *Linguistic Inquiry* 5: 367–424.

41. Quine, W.V. (1976), "Quantifiers and Propositional Attitudes," in *The Ways of Paradox and Other Essays* (Cambridge: Harvard University Press): 185–196.

42. Quirk, Randolph; Greenbaum, Sidney; Leech, Geoffrey; & Svartvik, Jan (1985), *A Comprehensive Grammar of the English Language* (New York: Longman).

43. Rapaport, William J. (1985), "Meinongian Semantics for Propositional Semantic Networks," *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics (University of Chicago)* (Morristown, NJ: Association for Computational Linguistics): 43–48.

44. Rapaport, William J. (1986), "Logical Foundations for Belief Representation," *Cognitive Science* 10: 371–422.

45. Rapaport, William J. (1988), "Syntactic Semantics: Foundations of Computational Natural-Language Understanding," in J. H. Fetzer (ed.) *Aspects of Artificial Intelligence* (Dordrecht, Holland: Kluwer Academic Publishers): 81–131.

46. Rapaport, William J.; Shapiro, Stuart C.; & Wiebe, Janyce M. (1986), "Quasi-Indicators, Knowledge Reports, and Discourse" *Technical Report 86-15* (Buffalo: SUNY Buffalo Dept. of Computer Science).

47. Reichman, Rachel (1985), *Getting Computers to Talk Like You and Me* (Cambridge, MA: MIT Press).

48. Rosch, Eleanor, & Lloyd, B.B. (1978), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum).

49. Shapiro, Stuart C., & Rapaport, William J. (1987), "SNePS Considered as a Fully Intensional Propositional Semantic Network," in N. Cercone & G. McCalla (eds.), *The Knowledge Frontier: Essays in the Representation of Knowledge* (New York: Springer-Verlag): 262–315.

50. Sidner, Candace L. (1983), "Focusing in the Comprehension of Definite Anaphora," in M. Brady & R. Berwick (eds.), *Computational Models of Discourse* (Cambridge, MA: MIT Press): 267–330.

51. Uspensky, Boris (1973), *A Poetics of Composition* (Berkeley: University of California Press).

52. Webber, Bonnie L. (1983), "So What Can We Talk About Now?," in M. Brady & R. Berwick (eds.), *Computational Models of Discourse* (Cambridge, MA: MIT Press).

53. Wiebe, Janyce M. (1990a), "Recognizing Subjective Sentences: A Computational Investigation of Narrative Text," *Technical Report 90-03* (Buffalo: SUNY Buffalo Dept. of Computer Science).

54. Wiebe, Janyce M. (1990b), "Identifying Subjective Characters in Narrative," *Papers Presented to the 13th International Conference on Computational Linguistics (COLING-90, Helsinki)* (Morristown, NJ: Association of Computational Linguistics): 401–408.

55. Wiebe, Janyce M., & Rapaport, William J. (1986), "Representing *De Re* and *De Dicto* Belief Reports in Discourse and Narrative," *Proceedings of the IEEE* 74: 1405–1413.

56. Wilks, Yorick, & Ballim, Afzal (1987), "Multiple Agents and the Heuristic Ascription of Belief," *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87, Milan)* (Los Altos, CA: Morgan Kaufmann): 118–124.

57. Wilks, Yorick, & Bien, Janusz (1983), "Beliefs, Points of View, and Multiple Environments," *Cognitive Science* 7: 95–119.

Texts Cited

Bridgers, Sue Ellen (1979), *All Together Now* (New York: Knopf).

Cherryh, C.J. (1988), *Exile's Gate* (New York: DAW Books).

Franchere, Ruth (1964), *Hannah Herself* (New York: Thomas Y. Crowell).

L'Engle, Madeleine (1986), *Many Waters* (New York: Dell Publishing).

Ludlum, Robert (1982), *The Parsifal Mosaic* (New York: Random House).

McCullers, Carson (1931), "The Jockey," in M. Crane (ed.), *50 Great Short Stories* (New York: Bantam Books, 1952).

Oneal, Zibby (1971), *War Work* (New York: Viking Press).

Tolkien, J.R.R. (1965), *The Two Towers* (New York: Ballantine Books).