

A COMPUTATIONAL THEORY OF PERSPECTIVE AND REFERENCE IN NARRATIVE

Janyce M. Wiebe and William J. Rapaport

Department of Computer Science
State University of New York at Buffalo
Buffalo, NY 14260
wiebe@cs.buffalo.edu, rapaport@cs.buffalo.edu

ABSTRACT

Narrative passages told from a character's perspective convey the character's thoughts and perceptions. We present a discourse process that recognizes characters' thoughts and perceptions in third-person narrative. An effect of perspective on reference in narrative is addressed: references in passages told from the perspective of a character reflect the character's beliefs. An algorithm that uses the results of our discourse process to understand references with respect to an appropriate set of beliefs is presented.

1. **INTRODUCTION.** A narrative is often told from the perspective of one or more of its characters; it can also contain passages that are not told from the perspective of any character. We present a computational theory of how readers recognize the current perspective in third-person narrative, and of the effects of perspective on the way readers understand references in third-person narrative. We consider published novels and short stories, rather than artificially constructed narratives.

2. **BANFIELD'S THEORY.** Our notion of perspective in narrative is based on Ann Banfield's (1982) categorization of the sentences of narration into subjective and objective sentences. *Subjective* sentences include those that portray a character's thoughts (*represented thought*) or present a scene as a character perceives it (*represented perception*). *Objective* sentences present the story directly, rather than through the thoughts or perceptions of a character. The language used to convey thoughts and perceptions is replete with linguistic elements that make no sense unless they are interpreted with respect to the thinking or perceiving character's consciousness. Banfield calls them *subjective elements*; they appear only in subjective sentences and cannot appear within objective sentences. Banfield identifies perspective in narrative with subjectivity, which is expressible via subjective elements. We call the thinking or perceiving character of a subjective sentence the *subjective character*.

3. **A DISCOURSE-LEVEL APPROACH.** Our task of recognizing the current perspective is, therefore, to recognize subjective sentences and the subjective characters to whom they are attributed. However, we cannot take a sentence-by-sentence approach, deciding independently for each sentence whether it is objective or subjective, and, if subjective, who the subjective character is. First, although thoughts and perceptions are often reported (as by sentences beginning

with "He thought that . . ." or "She saw . . ."), and thoughts are often accompanied by narrative parentheticals (such as "he thought" or "he realized"), many thoughts and perceptions are not marked in these ways. Second, subjective sentences do not always explicitly indicate who the subjective character is. For example:

- (1) ¹He wanted to talk to Dennys. ²How were they going to be able to get home from this strange desert land into which they had been cast and which was heaven knew where in all the countless solar systems in all the countless galaxies? [L'Engle, *Many Waters*, p. 91]
- (2) ²¹But what [Muhammad] had seen in those few moments made him catch his breath in amazement. ²²On the floor of the cave, which curved back in a natural fault in the rock, there were several large cylindrical objects standing in a row. [John Allegro, *The Dead Sea Scrolls*]

Sentence (1.2) is a represented thought, and (2.2) is a represented perception, presenting what the character sees as he sees it; yet neither is explicitly marked as such. Also, neither indicates who the subjective character is. Finally, although a subjective element marks a sentence as subjective (cf. Section 4.2), not all subjective sentences contain subjective elements, and subjective elements do not in general indicate who the subjective character is.

However, subjective sentences that are not marked as such, or that do not indicate who the subjective character is, usually appear in the midst of other subjective sentences attributed to the same subjective character. That is, once a clearly marked subjective sentence appears for which the subjective character can be determined, unmarked subjective sentences attributed to the same subjective character often follow. Thus, to recognize subjective sentences in general, we need to consider subjectivity at the level of the discourse. For this reason, we extend the notions of subjective and objective sentences to the notions of subjective and objective *contexts*, which consist of one or more subjective sentences attributed to the same subjective character, or one or more objective sentences, respectively.

Our algorithm for recognizing the current perspective is a discourse process that looks for the boundaries of subjective contexts. During narrative understanding, it maintains a stack, called the *current perspective (CP)*. At the beginning of a narrative, the CP is initialized to be the reader. When a new subjective context is recognized, its

subjective character is pushed onto the CP. When the end of a subjective context is recognized, a character is popped from the CP. More precisely, since SNePS (Shapiro 1979), our knowledge representation system, is fully intensional, only the reader's concepts of the characters are represented (Maida and Shapiro 1982, Shapiro and Rapaport 1987). So, it is actually the reader's concepts of the characters that are pushed onto the CP.

4. RECOGNIZING SUBJECTIVE CONTEXTS. To recognize subjective contexts, our discourse process relies exclusively on linguistic signals of subjective contexts. In this, it is incomplete: If a subjective context appears in which these linguistic signals are not present, then the subjective context is not recognized.

4.1. Psychological Verbs, Actions, Adjectives, and Perceptual Verbs. Reports involving psychological verbs (e.g., 'think', 'wonder', 'realize', 'want', 'remember') or perceptual verbs (e.g., 'see', 'hear') signal that a subjective context will follow. So do predicate-adjective sentences with psychological adjectives (e.g., 'delighted', 'happy', 'jealous', 'scared') (cf. Banfield (1982), Cohn (1978), Doležel (1973)). In addition, we have identified what we call *psychological actions*—e.g., "he smiled to himself", "she gasped", "she winced"—which function in the same way as psychological verbs.

A sentence of one of these types is a typical way of establishing a subjective context. Examples (1) and (2), above, and (3), below, exhibit this pattern:

- (3) ^{3.1}She [Hannah] winced as she heard them crash to the platform. ^{3.2}The lovely little mirror that she had brought for Ellen, and the gifts for the baby! [Franchère, *Hannah Herself*, p. 3]

In each example, the first sentence is a psychological or perceptual report, and the second is a represented thought or represented perception, respectively; the subjective character of the second sentence is taken to be the subject of the first. In our discourse process, the subject of a perceptual or psychological report, or of a predicate-adjective sentence with a psychological adjective, is pushed onto the CP if a character isn't already on the top of it. If a character is already on the top of the CP, then no change is made, and the sentence is understood to be part of the already established subjective context.

4.2. Subjective Elements. Many subjective elements mark a sentence in third-person narrative as subjective because they are *expressive* in nature. Some that Banfield identifies are *exclamations*, which express emotion; *questions*, which express wonder; *epithets*, such as 'the bastard', which express some qualification of the referent; and certain *kinship terms*, e.g., 'Daddy', 'Mom', and 'Aunt Margaret', which express a relationship to the referent. She also identifies *evaluative adjectives*, which express an attitude

toward the referent, e.g., 'ghastly', 'surprising', 'poor', and 'damned', although some evaluative adjectives, such as 'poor' and 'damned', have their evaluative meanings only when they occur in certain parts of the sentence. Intensifiers such as 'too', 'quite', and 'so' are also evaluative (Banfield 1982), as in:

- (4) He could tell they were tears because his eyes were too shiny. Too round. [Bridgers, *All Together Now*, p. 92]

So are *emphasizers*, such as 'really' and 'just'. An example is 'really' in (5.3):

- (5) ^{5.1}Jody managed a frail smile. ^{5.2}She was a little bit ashamed. ^{5.3}She should really try to be more cheerful for Aunt Margaret's sake. ^{5.4}After all, Aunt Margaret had troubles of her own—she was the mother of that ghastly Dill. [Gage, *Miss Osborne-the-Mop*, pp. 16-17]

Modal verbs of obligation, possibility, and necessity are also expressive. For example, 'should', in (5.3), is a modal verb of obligation. So are many *content (or attitudinal) disjuncts*, which comment on the content of the utterance (Quirk et al. 1985). For example, 'likely', 'maybe', 'probably', and 'perhaps' express some degree of doubt:

- (6) Something jingled—car keys probably. [Oneal, *War Work*, p. 132]

Conjuncts, which comment on the connection between items (Quirk et al. 1985), can also be expressive. For example, 'anyhow', 'anyway', 'still', and 'after all' express concession (Quirk et al. 1985). An example is 'After all' in (5.4).

Other subjective elements are sentence fragments (Banfield 1982), such as (7.2),

- (7) ^{7.1}His brain worked slowly through what he knew about this person. ^{7.2}David's kid. [Bridgers, *All Together Now*, p. 91]

and the uses of 'this', 'that', 'these', and 'those' that Robin Lakoff (1974) has identified as *emotional deixis*. In conversation, they are "generally linked to the speaker's emotional involvement in the subject-matter of his utterance" (Lakoff 1974: 347); in third-person narrative, they are linked to the subjective character's emotional involvement in the subject matter of his thoughts or perceptions. Examples are 'this' in (8.1) and 'That' in (9.2):

- (8) ^{8.1}Ibrahim could remember every time this godless pig had patronized him . . . [Clancy, *Red Dawn Rising*, p. 13]

- (9) ^{9.1}As she watched, a wave of jealousy spread through her. ^{9.2}That insufferable stranger who had passed them on the road was receiving the welcome that she had been dreaming of all the way from Connecticut. [Franchere, *Hannah Herself*, p. 15]

In speech, the emotion, evaluation, etc., expressed by a subjective element is always attributed to the speaker; in third-person narrative, it is attributed to a character.¹ Clearly, many types of language-understanding abilities are needed to understand the range of subjective elements. Our purpose here is to show how our discourse process uses them as markers of subjective contexts, and how it determines the subjective character whose thoughts or perceptions they mark. However, recognizing the subjective character is always required before a subjective element can be understood.

When a subjective element is encountered in the narrative, our discourse process updates the CP according to the following algorithm:

- (A1)
- If there is currently a character on the CP,
 - ¹then do not change the CP
 - else if there is an actor focus at the start of the current sentence who is a character in the scene
 - ²then push him or her onto the CP
 - ³else create a new and indeterminate concept and push it onto the CP.

4.2.1. Discussion of branch 1. Branch 1 is taken when a subjective element continues the current subjective context. For example, the exclamation in (3.2), which is a subjective element, continues the subjective context established by (3.1). The subjective elements 'should' and 'really' in (5.3) and 'After all', 'Aunt Margaret', 'that', and 'ghastly' in (5.4) continue the subjective context established in (5.1).

4.2.2. Discussion of branch 2. The *actor focus* used in branch 2 is one of the foci that need to be maintained for the comprehension of definite anaphora (Sidner 1983). It is whoever is the agent of the current sentence. (Note that quoted speech has its own foci, which must be maintained separately. In this sense, quoted speech constitutes a separate discourse segment (cf. Grosz and Sidner (1986).) Consider the following example:

- (10) ^{10.1}In the kitchen she [Jody] set the basket down on the table. ^{10.2}She put the thermos and the cups in the sink and filled them with cool water to soak. ^{10.3}Then

¹ In some third-person novels, particularly in the 19th century, an *overt narrator* (Chatman 1978) uses subjective elements. We do not consider novels with overt narrators.

she tiptoed upstairs to her room. ^{10.4}Perhaps Aunt Margaret was taking a nap. ^{10.5}It wouldn't do to disturb her. [Gage, *Miss Osborne-the-Mop*, p. 25]

Since Jody is the actor focus at the beginning of (10.4) (she is the actor focus of (10.1)-(10.3)), and she is a character in the scene, the subjective element 'Perhaps' is attributed to her when it is encountered, and she is pushed onto the CP.

Sidner (1983) has shown that, in anaphora comprehension, the current actor or discourse focus can be rejected as the co-specifier of an anaphor on the basis of pragmatic factors. Similarly, the actor focus may be rejected as the subjective character to whom the subjective element is attributed, in favor of another character in the scene. The pragmatic factors involved appear to be which characters have been subjective characters in the past and whose thoughts or perceptions the sentence containing the subjective element is likely to be reflecting. Consider the following example, in which Adnarel, a seraph, has just appeared before Lemech and Sandy:

- (11) ^{11.1}Lemech greeted him [Adnarel] respectfully. "Adnarel, we thank you." ^{11.2}Then he said to Sandy, "The seraph will be able to help you. Seraphim know much about healing."
^{11.3}So this was a seraph. [L'Engle, *Many Waters*, p. 39]

Lemech is the actor focus of (11.1) and (11.2). However, it would be clear to someone who had read the novel up to this point that (11.3) is Sandy's thought. First, Sandy is a visitor to a strange world, of which Lemech is an inhabitant; so it is Sandy, not Lemech, who is likely not to have known what a seraph is. Second, prior to this passage, subjective contexts have been attributed to Sandy, but not to Lemech. We are investigating the reasoning required by the reader in rejecting the current actor in favor of another character.

4.2.3. Discussion of branch 3. Branch 3 is taken when the reader cannot identify the subjective character and must read further in the text in order to do so. In this case, an indeterminate, intensional concept is pushed onto the CP. When the reader finally identifies the subjective character, the information that this character and the indeterminate one are co-extensional is built (that is, it is asserted that they are concepts of the same individual; cf. Maida and Shapiro (1982), Shapiro and Rapaport (1987)). Naicong Li (1986) uses this approach in her pronoun resolution algorithm if the information needed to resolve a pronoun is supplied after the pronoun is encountered.

Often, the subjective character is identified by a narrative parenthetical, as in the following example:

- (12) ^{12.1}What was holding Washington up? the colonel asked himself. ^{12.2}All he needed was a simple yes or no. [Clancy, *Red Storm Rising*, p. 170]

Sentence (12.1) begins with a question, which is a subjective element. It occurs just after a shift in scene, so it does not continue a current subjective context, and there is no actor focus. When 'What' is encountered in (12.1), branch 3 pushes a new concept onto the CP: At this point, the reader has recognized someone's thought, but does not yet know whose. When the reader encounters the parenthetical, she identifies the subjective character as the colonel, and builds a proposition that the colonel and the new concept are co-extensional (that is, she comes to believe that the question was the colonel's thought).

4.2.4. Comparison of evaluative and psychological adjectives. Before leaving our discussion of subjective elements, it will be useful to contrast the ways that predicate-adjective sentences with psychological and with evaluative adjectives are treated by our discourse process. Compare the following sentences:

- (A) Jody was delighted.
- (B) Jody was ghastly.

Sentence (A) contains the psychological adjective 'delighted', and (B) contains the evaluative adjective 'ghastly'. In (A) (assuming no previous subjective context), Jody is pushed onto the CP, and (A) establishes a subjective context attributed to Jody. In (B), algorithm (A1) determines whose attitude toward Jody is being expressed, and it does not choose the subject of the sentence. Thus, psychological adjectives can establish the perspective of the subject, whereas evaluative adjectives express an attitude *toward* the subject.

5. RECOGNIZING ENDING BOUNDARIES OF SUBJECTIVE CONTEXTS. Recognizing the ending boundaries of subjective contexts is a more difficult problem than recognizing the beginning boundaries. While not all subjective contexts are signaled in the ways discussed in Section 4, it is very common that they are. However, we have not found equally reliable or common signals for the ending boundaries. It appears that the reader often has to reason about the content of the current sentence and confirm that it can continue the subjective context; if it cannot, then the ending boundary has been found.

Nevertheless, we have identified two reliable ways of recognizing the ending boundaries of subjective contexts. One way subjective contexts are ended is by a shift in scene, as in the following example:

- (13) He [Sandy] wanted to talk to Dennys. How were they going to be able to get home from this strange desert land into which they had been cast and which was heaven knew where in all the countless solar systems in all the countless galaxies?

<Chapter Break>

Dennys was sleeping fitfully when he heard the tent flap move. [L'Engle, *Many Waters*, pp. 91-92]

Dennys and Sandy are not at the same place. The shift in scene at the chapter break ends the subjective context attributed to Sandy.

A second way is by a negated perceptual report whose subject is the subjective character and whose object is something in the scene. For example,

- (14) ^{14.1}She [Yalith] was not sure why she was hesitant. ^{14.2}She breathed in the strange odor of his wings, smelling of stone, of the cold, dark winds which came during the few brief weeks of winter. ^{14.3}Enveloped in Eblis's wings, she did not hear the rhythmic thud as a great lion galloped toward them across the desert, roaring as it neared them. ^{14.4}Then both Yalith and Eblis turned and saw the lion rising to its hind legs . . . [L'Engle, *Many Waters*, p. 47]

The subjective context in the first paragraph is ended by the negated perceptual verb in (14.3). Sentence (14.4) then establishes a new subjective context attributed to Yalith and Eblis. In a similar way, subjective contexts can be ended by negated factive verbs, as in "He did not realize that . . .".

6. BELIEF AND SUBJECTIVE CONTEXTS. Since subjective contexts portray thoughts and perceptions, the reader understands that the information they convey reflects the subjective character's beliefs (cf. Fillmore (1974), Banfield (1982), Uspensky (1973)). Whatever else the reader may infer that the characters believe, she has to attribute the information in subjective contexts to the subjective character. Brian Reiser (1981) showed that one of the effects of perspective on a reader's understanding is that it focuses processing. In particular, he showed that the reader primarily infers the goals and plans of the character whose perspective the narrative is taking. In a similar way, perspective focuses the reader's attribution of beliefs to the characters.

References in subjective contexts reflect just what the subjective character believes. (Cf. Clark and Marshall (1981), Cohen, Perrault, and Allen (1982), and Wilks and Bien (1983) for discussions of belief and reference in conversation.) The subjective character might be mistaken, or know less about the referent than the reader or the other characters know, or know more than the other characters. The remainder of this paper addresses the attribution of beliefs to characters in order to understand references in subjective contexts.

6.1. An Algorithm for Understanding References Using the CP. Our belief representation, described in Rapaport and Shapiro (1984), Rapaport (1986), and Wiebe and Rapaport (1986), is based on the notion of belief spaces. A *belief space* is accessed by a stack of individuals, and consists of what the bottom member of the stack believes that . . . the top member believes. The reader is always the bottom member of the stack, and the belief space corresponding to a stack consisting only of the reader contains the set of

propositions that the reader believes are true. All propositions in the knowledge base appear in at least one belief space, and a single proposition can appear in more than one belief space. This occurs, for example, if the reader believes a proposition and believes that a character believes it, too.

The CP determines the current belief space with respect to which references are understood. So far, our analysis extends only to non-anaphoric, specific references. The following is our algorithm for understanding a non-anaphoric, specific reference 'X' in third-person narrative (there may actually be more than one proposition found or built in order to understand 'X', for example, if 'X' is plural or a possessive):

(A2)

- If 'X' is an indefinite noun phrase of the form 'a Y',
 - ¹then create a new concept, N; build in the CP's belief space the proposition that N is a Y; return N
- else if 'X' is a definite noun phrase or proper name,
 - then if a proposition that N is X can be found in the CP's belief space,
 - ²then return N
 - else if a proposition that N is X can be found in a belief space other than the CP's,
 - ³then add the found proposition to the CP's belief space; return N
 - ⁴else create a new concept, N; build in the CP's belief space the proposition that N is X; return N.

6.1.1. Discussion of branch 1. Indefinite references introduce new individuals into the CP's belief space. We discuss this in Section 6.3, below.

6.1.2. Discussion of branches 2 and 3. The search for the referent of a non-anaphoric definite noun phrase or proper name starts in the CP's belief space. Branch 2 is taken if the referent can be found there.

If the test in branch 2 fails, then the rest of the knowledge base must be searched. To see why this is so, suppose that the reference is 'Ellen' and that it occurs in a subjective context. It is possible that Ellen has been referred to previously in the narrative but not under any circumstances that would have required the reader to explicitly attribute the belief that she is named 'Ellen' to the subjective character. Perhaps she has only been referred to in objective contexts, for example. So, to find the referent, other belief spaces than the CP's must be searched.

Branch 3 is taken if the search is successful, and, before the referent is returned, the proposition that the referent is X is added to the CP's belief space. In the case just discussed, the fact that the reference occurs in a subjective context indicates that the belief that the referent is named 'Ellen' should now be attributed to the subjective character.

6.1.3. Discussion of branch 4. Branch 4 is taken in order to understand definite noun phrases and proper names that refer to individuals who have not been previously introduced into the narrative. For example, in (3.2), above, neither the mirror, the gifts, the baby, nor Ellen have been mentioned before in the novel. A new referent for each is introduced into the CP's belief space; that is, by virtue of understanding the references in (3.2), 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'.

6.1.4. An example. We now illustrate our algorithms on a passage that reflects a character's mistaken belief. The passage is from a novel in which a character, Dwayne, mistakenly believes that another character, Casey, is a boy:

- (15) ^{15.1}His [Dwayne's] brain worked slowly through what he knew about this person [Casey]. ^{15.2}David's kid. ^{15.3}The name stumbled into place. ^{15.4}This was David's boy. ^{15.5}David was in the war, and here was his kid in the arcade scared of something. [Bridgers, *All Together Now*, p. 91]

Note that (15.1) and (15.3) are psychological reports that employ metaphor, rather than psychological verbs, to report the character's psychological experience ((9.1), above, employs metaphor in a similar way). Metaphor is beyond the scope of this work, so, before applying our algorithm to this passage, we paraphrase it as follows:

- (15a) ^{15a.1}He [Dwayne] thought of what he knew about this person [Casey]. ^{15a.2}David's kid. ^{15a.3}He remembered the name. ^{15a.4}This was David's boy. ^{15a.5}David was in the war, and here was his kid in the arcade scared of something.

First, consider the operation of our discourse process. Sentence (15a.1) is a psychological report, and so Dwayne, its subject, is pushed onto the CP; this establishes a subjective context, attributed to Dwayne, which is continued throughout the passage. Note that when the sentence fragment (which is a subjective element) is encountered, no change is made to the CP because there is already a character on the top of it. Similarly, no change is made to the CP when (15a.3), a psychological report, and the second conjunct of (15a.5), a predicate-adjective sentence with a psychological adjective, are encountered, since there is already a character on the top of the CP.

Now, consider the reference to David in (15a.2). The reader knows that David is Casey's father. If, before reading (15a.2), the reader didn't explicitly believe that Dwayne knew about David too, then branch 3 of algorithm (A2) would be taken to understand this reference; the result is that the reader now explicitly believes that Dwayne knows about David.

'David's boy' in (15a.4) reflects Dwayne's mistaken belief about Casey, and branch 2 of algorithm (A2) is taken in order to understand it. To illustrate that information in subjective contexts is attributed to the subjective character, suppose that (15a.4) were "This was David's girl"; in that case, the reader would have to infer that Dwayne had somehow found out that Casey is a girl.

6.1.5. Further discussion of algorithm (A2). Note that if a reference is a subjective element, such as 'the bastard', it may be a *non-classificatory noun* (Banfield 1982); that is, it cannot be understood entirely propositionally, since it expresses subjectivity. How it should be understood depends on the particular subjective element. Thus, specific algorithms for nouns that are subjective elements must supersede algorithm (A2).

As mentioned above, algorithm (A2) is unable to understand anaphoric references. However, anaphor comprehension can be affected by perspective. Consider the following passage:

- (16) ^{16.1}The man had turned. ^{16.2}He started to walk away quickly in the direction of the public library. ^{16.3}"O.K.," said Joe, "get Rosie."
^{16.4}Zoe crept back to the blanket. ^{16.5}She felt hollow in her stomach. ^{16.6}She'd never really expected to see the Enemy again. [Oneal, *War Work*, p. 64]

In (16.6), 'the Enemy' is an anaphoric reference that occurs in a subjective context (established by (16.5), which is a psychological report); it co-specifies 'the man' in (16.1) and 'He' in (16.2). It reflects Zoe's belief that the man is an enemy spy, although it is not at all clear to the reader, at this point, that he is.

Personal pronouns can also reflect the beliefs of a character. The following passage is a continuation of passage (15) (italics ours):

- (17) He [Dwayne] wasn't sure of what. What in the arcade could scare a boy like that? He rubbed his head under his baseball cap. He could see tears in Casey'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 Casey was in a new place, too, feeling maybe the same as him. If he just knew what to do about it.*
 "Let's don't play that game anymore," he said. "I don't like that one."
 Casey wiped *her* face on her sleeve . . . [Bridgers, *All Together Now*, p. 92]

Both italicized pronouns refer to Casey; the first occurs in a subjective context attributed to Dwayne, and the second occurs after the subjective context has ended (in this passage, the subjective context is ended by direct speech).

6.2. Assertive Indefinite Pronouns. Assertive indefinite pronouns—e.g., 'someone', 'something', 'somebody'—are specific, though unspecified (Quirk et al. 1985); that is, they generally refer to particular people, things, etc., without identifying them. When referring to a particular referent, a speaker typically uses an assertive indefinite pronoun if (1) she doesn't know the identity of the referent, (2) she doesn't want the addressee to know the identity of the referent, or (3) she doesn't believe that the identity of the referent is relevant to the conversation. A character's thoughts and perceptions are not directed toward an addressee, and so the first of these uses is the predominant one in subjective contexts. Used in this way, they express a lack of knowledge, and so are subjective elements. When one of them appears in a subjective context, the reader understands that the subjective character does not know who or what the referent is. Often, the pronoun is the only source of this information. Consider the following example:

- (18) ^{18.1}Suddenly she [Zoe] gasped. ^{18.2}She had touched somebody! [Oneal, *War Work*, p. 129]

There is no explicit statement in the novel that Zoe does not know whom she touched; this has to be inferred from the use of 'somebody'. Sentences (6) and (15.5) provide further examples.

6.3. Indefinite References. In conversation, definite references are used only if the speaker believes that the addressee has enough information to interpret them. As mentioned above, thoughts and perceptions are not directed toward an addressee, and so the use of definite references in subjective contexts is not subject to this constraint; as illustrated by (3.2), they are used to refer to referents familiar to the subjective character, whether or not the reader has been told about them before. So, when a specific indefinite reference appears in a subjective context, the reader understands that the referent is unfamiliar to the subjective character; otherwise, a definite reference would have appeared (Fillmore 1974).

However, the referent may not be unknown to the reader or to the other characters. For example,

- (19) 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 reader knows that 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. When the King and his men are on the top of the CP (after 'saw' and continued by 'were aware of'), the hobbits are not referred to by name, but as 'two small figures'. Branch 1 of algorithm (A2) creates new referents and, in the belief space of the King and his men,

builds propositions that they are small figures. The new referents can be asserted to be co-extensional with the concepts who the reader and other characters believe are named 'Merry' and 'Pippin'.

Indefinite references can sometimes indicate that the subjective character doesn't even know what the referent is. This occurs when the head noun is a superordinate, rather than a basic-level, term (Rosch and Lloyd 1978). The basic level is the preferred level at which people identify things. If a superordinate, rather than a basic-level, term appears in an indefinite reference in a subjective context, the reader understands that the subjective character can't even identify the referent at the basic level. In example (19), the hobbits are referred to as 'two small figures', because the King and his men have never seen hobbits before. Here is an example that is not from a fantasy novel:

- (20) 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 only identify them as 'small dots of purple'. Another occurs in (2.2): The fact that the reference 'several large cylindrical objects' includes the superordinate term 'objects' indicates that Muhammad doesn't know what the referents are. Another example is (21):

- (21) He felt firm restraints of some sort holding him in place. [Wu, *Cyborg*, p. 141]

Peters and Shapiro (1987ab) describe a SNePS representation for natural category systems in which superordinate categories can be distinguished from basic-level and subordinate categories. After an indefinite reference with a superordinate term in a subjective context has been parsed, the fact that the subjective character was able to identify the referent only at a superordinate level is represented in the knowledge base by using their representation.

7. CONCLUSIONS AND FUTURE RESEARCH. Many problems remain to be solved. Our discourse process cannot recognize subjective contexts that are not established by the linguistic signals it relies on, and general principles are needed to explain how readers recognize the ending boundaries of subjective contexts. We are investigating how tense, deictic terms (cf. Bruder et al. (1986), Banfield (1982)), the characters' goals (cf. Wilensky (1983)), and the argument structure (cf. Cohen (1987)) often exhibited by thoughts might be used to recognize the boundaries of subjective contexts. Branches 2 and 3 of algorithm (A1) need to be expanded to determine who the subjective character is if the actor focus isn't a reasonable candidate and no parenthetical appears. We are investigating how focus of attention (cf. Grosz (1981), Sidner (1983)) can be incorporated into algorithm (A2) in such a way that anaphoric

references reflecting the beliefs of a character can be understood. Finally, there is the general problem of revision. Our algorithms assume that signals occur at the beginning of subjective contexts. However, there are cases when a subjective context cannot be recognized until some of it has already been parsed. A difficult case is illustrated by the following: "C₁Jody was rich and famous. C₂Why wasn't she happy? Bill wondered." Only after reading (C2) can the reader recognize that (C1) is a represented thought.

We have argued that a discourse-level approach must be taken to the problem of recognizing character's thoughts and perceptions in third-person narrative. Our discourse process, which is implemented in an ATN grammar interfaced to SNePS, recognizes subjective contexts that are linguistically signaled in ways frequently employed in naturally-occurring narratives. By using the results of the discourse process to determine the belief context needed to understand references, our reference algorithm demonstrates how perspective affects reference in third-person narrative.

8. ACKNOWLEDGMENTS. We are indebted to Mary Galbraith, David Zubin, Sandra Peters, Stuart Shapiro, and the other members of the SUNY Buffalo Graduate Group in Cognitive Science and the SNePS Research Group for many discussions and ideas. This research was supported in part by NSF grants IST-8504713, IRI-8610517.

REFERENCES

- Banfield, Ann (1982), *Unspeakable Sentences: Narration and Representation in the Language of Fiction* (Boston: Routledge & Kegan Paul).
- Bruder, Gail A.; Duchan, Judy F.; Rapaport, William J.; Segal, Erwin M.; Shapiro, Stuart C.; and Zubin, David A. (1986), "Deictic Centers in Narrative: An Interdisciplinary Cognitive-Science Project," *Technical Report 86-20* (Buffalo: SUNY Buffalo Dept. of Computer Science).
- Chatman, Seymour (1978), *Story and Discourse: Narrative Structure in Fiction and Film* (Ithaca, NY: Cornell University Press).
- Clark, Herbert H. and Marshall, Catharine R. (1981), "Definite Reference and Mutual Knowledge," in A. Joshi, B. Webber, and I. Sag (eds.), *Elements of Discourse Understanding* (Cambridge: Cambridge University Press): 10-63.
- Cohen, Philip R.; Perrault, C. Raymond; and Allen, James F. (1982), "Beyond Question Answering," in W. Lehnert and M. Ringle (eds.), *Strategies for Natural Language Processing* (Hillsdale, NJ: Lawrence Erlbaum): 245-274.
- Cohen, Robin (1987), "Analyzing the Structure of Argumentative Discourse," *Computational Linguistics* 13: 11-24.
- Cohn, Dorrit (1978), *Transparent Minds: Narrative Modes for Representing Consciousness in Fiction* (Princeton: Princeton University Press).

- Doležal, Lubomir (1973), *Narrative Modes in Czech Literature* (Toronto: University of Toronto Press).
- Fillmore, Charles (1974), "Pragmatics and the Description of Discourse," in C. Fillmore, G. Lakoff, and R. Lakoff (eds.), *Berkeley Studies in Syntax and Semantics I* (Berkeley: University of California Dept. of Linguistics and Institute of Human Learning): V1-V21.
- Grosz, Barbara J. (1981), "Focusing and Description in Natural Language Dialogues," in A. Joshi, B. Webber, and I. Sag (eds.), *Elements of Discourse Understanding* (Cambridge: Cambridge University Press): 84-105.
- Grosz, Barbara J. and Sidner, Candace L. (1986), "Attention, Intentions, and the Structure of Discourse," *Computational Linguistics*, 12: 175-204.
- Lakoff, Robin (1974), "Remarks on 'this' and 'that,'" *Papers from the Tenth Regional Meeting of the Chicago Linguistic Society* (Chicago: Chicago Linguistic Society): 345-356.
- Li, Naicong (1986), "Pronoun Resolution in SNePS," *SNeRG Technical Note 18* (Buffalo: SUNY Buffalo Dept. of Computer Science).
- Maida, Anthony S. and Shapiro, Stuart C. (1982), "Intensional Concepts in Propositional Semantic Networks," *Cognitive Science* 6: 291-330.
- Quirk, Randolph; Greenbaum, Sidney; Leech, Geoffrey; and Svartvik, Jan (1985), *A Comprehensive Grammar of the English Language* (New York: Longman).
- Peters, Sandra L. and Shapiro, Stuart C. (1987a), "A Representation for Natural Category Systems," *Proceedings of the 9th Annual Conference of the Cognitive Science Society (Seattle)* (Hillsdale, NJ: Lawrence Erlbaum): 379-390.
- Peters, Sandra L. and Shapiro, Stuart C. (1987b), "A Representation for Natural Category Systems," *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87; Milan)* (Los Altos, CA: Morgan Kaufmann): 140-146.
- Rapaport, William J. and Shapiro, Stuart C. (1984), "Quasi-Indexical Reference in Propositional Semantic Networks," *Proceedings of the 10th International Conference on Computational Linguistics (COLING-84; Stanford Univ.)* (Morristown, NJ: Assoc. for Computational Linguistics): 65-70.
- Rapaport, William J. (1986), "Logical Foundations for Belief Representation," *Cognitive Science* 10: 371-422.
- Reiser, Brian J. (1981), "Character Tracking and the Understanding of Narrative," *Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI-81; Vancouver)* (Los Altos, CA: Morgan Kaufmann): 209-211.
- Rosch, Eleanor and Lloyd, B.B. (1978), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates).
- Shapiro, Stuart C. (1979), "The SNePS Semantic Network Processing System," in N.V. Findler (ed.), *Associative Networks* (New York: Academic): 179-203.
- Shapiro, Stuart C. and Rapaport, William J. (1987), "SNePS Considered as a Fully Intensional Propositional Semantic Network," in N. Cercone and G. McCalla (eds.), *The Knowledge Frontier* (New York: Springer-Verlag): 262-315.
- Sidner, Candace L. (1983), "Focusing in the Comprehension of Definite Anaphora," in M. Brady and R. Berwick (eds.), *Computational Models of Discourse* (Cambridge, MA: The MIT Press): 267-330.
- Uspensky, Boris (1973), *A Poetics of Composition* (Berkeley: University of California Press).
- Wiebe, Janyce M. and Rapaport, William J. (1986), "Representing *De Re* and *De Dicto* Belief Reports in Discourse and Narrative," *Proc. IEEE* 74: 1405-1413.
- Wilensky, Robert (1983), *Planning and Understanding* (Reading, MA: Addison-Wesley).
- Wilks, Yorick and Bien, Janusz (1983), "Beliefs, Points of View, and Multiple Environments," *Cognitive Science* 7: 95-119.

CITED TEXTS

- Allegro, John (1977), *The Dead Sea Scrolls* (Harmondsworth, Eng.: Penguin).
- Bridgers, Sue Ellen (1979), *All Together Now* (New York: Knopf).
- Clancy, Tom (1986), *Red Storm Rising* (New York: G. P. Putnam's Sons).
- Franchere, Ruth (1964), *Hannah Herself* (New York: Thomas Y. Crowell).
- Gage, Wilson (1963), *Miss Osborne-the-Mop* (Cleveland: World Publishing).
- L'Engle, Madeleine (1986), *Many Waters* (New York, Dell Publishing).
- Oneal, Zibby (1971), *War Work* (New York, Viking Press).
- Tolkien, J.R.R. (1965), *The Two Towers* (New York: Ballantine Books).
- Wu, William (1987), *Cyborg* (New York: Ace Books).