

Annotating Attributions and Private States

Theresa Wilson
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
twilson@cs.pitt.edu

Janyce Wiebe
Department of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260
wiebe@cs.pitt.edu

Abstract

This paper describes extensions to a corpus annotation scheme for the manual annotation of attributions, as well as opinions, emotions, sentiments, speculations, evaluations and other *private states* in language. It discusses the scheme with respect to the “Pie in the Sky” *Check List of Desirable Semantic Information for Annotation*. We believe that the scheme is a good foundation for adding private state annotations to other layers of semantic meaning.

1 Introduction

This paper describes a fine-grained annotation scheme for key components and properties of opinions, emotions, sentiments, speculations, evaluations, and other *private states* in text. We first give an overview of the core scheme. We then describe recent extensions to the scheme, namely refined annotations of *attitudes* and *targets*, or objects, of private states. Finally, we discuss related items from the “Pie in the Sky” *Check List of Desirable Semantic Information for Annotation*, and related work. We believe our scheme would provide a foundation for adding private state annotations to other layers of semantic and pragmatic meaning.

2 The Core Scheme

This section overviews the core of the annotation scheme. Further details may be found in (Wilson and Wiebe, 2003; Wiebe et al., 2005).

2.1 Means of Expressing Private States

The goals of the annotation scheme are to represent internal mental and emotional states, and to distinguish subjective information from material presented as fact. As a result, the annotation scheme is centered on the notion of *private state*, a general term that covers opinions,

beliefs, thoughts, feelings, emotions, goals, evaluations, and judgments. As Quirk et al. (1985) define it, a *private state* is a state that is not open to objective observation or verification: “a person may be observed to *assert that God exists*, but not to *believe that God exists*. Belief is in this sense ‘private.’” (p. 1181) Following literary theorists such as Banfield (1982), we use the term *subjectivity* for linguistic expressions of private states in the contexts of texts and conversations.

We can further view private states in terms of their functional components — as states of *experiencers* holding *attitudes*, optionally toward *targets*. For example, for the private state in the sentence *John hates Mary*, the experiencer is “John,” the attitude is “hate,” and the target is “Mary.”

We create private state frames for three main types of private state expressions in text:

- explicit mentions of private states
- speech events expressing private states
- expressive subjective elements

An example of an explicit mention of a private state is “fears” in (1):

(1) “The U.S. **fears** a spill-over,” said Xirao-Nima.

An example of a *speech event* expressing a private state is “said” in (2):

(2) “The report is **full of absurdities**,” Xirao-Nima **said**.

Note that we use the term *speech event* to refer to both speaking and writing events.

The phrase “full of absurdities” in (2) above is an *expressive subjective element* (Banfield, 1982). Other examples can be found in (3):

(3) **The time has come, gentlemen**, for Sharon, **the assassin**, to realize that **injustice cannot last long**.

The private states in this sentence are expressed entirely by the words and the style of language that is used. In (3), although the writer does not explicitly say that he hates Sharon, his choice of words clearly demonstrates a negative attitude toward him. As used in these sentences, the phrases “The time has come,” “gentlemen,” “the assassin,” and “injustice cannot last long” are all expressive subjective elements. Expressive subjective elements are used by people to express their frustration, anger, wonder, positive sentiment, etc., without explicitly stating that they are frustrated, angry, etc. Sarcasm and irony often involve expressive subjective elements.

2.2 Private State Frames

We propose two types of private state frames: *expressive subjective element frames* will be used to represent expressive subjective elements; and *direct subjective frames* will be used to represent both subjective speech events (i.e., speech events expressing private states) and explicitly mentioned private states. The frames have the following attributes:

Direct subjective (subjective speech event or explicit private state) frame:

- **text anchor:** a pointer to the span of text that represents the speech event or explicit mention of a private state.
- **source:** the person or entity that expresses or experiences the private state, possibly the writer.
- **target:** the target or topic of the private state, i.e., what the speech event or private state is about.
- **properties:**
 - **intensity:** the intensity of the private state (*low, medium, high, or extreme*).
 - **expression intensity:** the contribution of the speech event or private state expression itself to the overall intensity of the private state. For example, “say” is often neutral, even if what is uttered is not neutral, while “excoriate” itself implies a very strong private state.
 - **insubstantial:** true, if the private state is not substantial in the discourse. For example, a private state in the context of a conditional often has the value *true* for attribute *insubstantial*.
 - **attitude type:** the type of attitude(s) composing the private state.

Expressive subjective element frame:

- **text anchor:** a pointer to the span of text that denotes the subjective or expressive phrase.

- **source:** the person or entity that is expressing the private state, possibly the writer.
- **properties:**
 - **intensity:** the intensity of the private state.
 - **attitude type**

2.3 Objective Speech Event Frames

To distinguish opinion-oriented material from material presented as factual, we also define *objective speech event frames*. These are used to represent material that is attributed to some source, but is presented as objective fact. They include a subset of the slots in private state frames:

Objective speech event frame:

- **text anchor:** a pointer to the span of text that denotes the speech event.
- **source:** the speaker or writer.
- **target:** the target or topic of the speech event, i.e., the content of what is said.

For example, an objective speech event frame is created for “said” in the following sentence (assuming no undue influence from the context):

(4) Sargeant O’Leary said the incident took place at 2:00pm.

That the incident took place at 2:00pm is presented as a fact with Sargeant O’Leary as the source of information.

2.4 Agent Frames

The annotation scheme includes an *agent frame* for noun phrases that refer to sources of private states and speech events, i.e., for all noun phrases that act as the experiencer of a private state, or the speaker/writer of a speech event. Each agent frame generally has two slots. The *text anchor* slot includes a pointer to the span of text that denotes the noun phrase source. The *source* slot contains a unique alpha-numeric ID that is used to denote this source throughout the document. The agent frame associated with the first informative (e.g., non-pronominal) reference to this source in the document includes an *id* slot to set up the document-specific source-id mapping.

2.5 Nested Sources

The source of a speech event is the speaker or writer. The source of a private state is the experiencer of the private state, i.e., the person whose opinion or emotion is being expressed. The writer of an article is always a source, because he or she wrote the sentences of the article, but the writer may also write about other people’s private states

and speech events, leading to multiple sources in a single sentence. For example, each of the following sentences has two sources: the writer (because he or she wrote the sentences), and Sue (because she is the source of a speech event in (5) and of private states in (6) and (7)).

- (5) Sue said, “The election was fair.”
- (6) Sue thinks that the election was fair.
- (7) Sue is afraid to go outside.

Note, however, that we don’t really know what Sue says, thinks or feels. All we know is what the writer tells us. For example, Sentence (5) does not directly present Sue’s speech event but rather Sue’s speech event according to the writer. Thus, we have a natural *nesting of sources* in a sentence.

In particular, private states are often filtered through the “eyes” of another source, and private states are often directed toward the private states of others. Consider sentence (1) above and (8) following:

- (8) China criticized the U.S. report’s criticism of China’s human rights record.

In sentence (1), the U.S. does not directly state its fear. Rather, according to the writer, according to Xirao-Nima, the U.S. fears a spill-over. The source of the private state expressed by “fears” is thus the *nested source* (*writer, Xirao-Nima, U.S.*). In sentence (8), the U.S. report’s criticism is the target of China’s criticism. Thus, the nested source for “criticism” is (*writer, China, U.S. report*).

Note that the shallowest (left-most) agent of all nested sources is the writer, since he or she wrote the sentence. In addition, nested source annotations are composed of the IDs associated with each source, as described in the previous subsection. Thus, for example, the nested source (*writer, China, U.S. report*) would be represented using the IDs associated with the writer, China, and the report being referred to, respectively.

2.6 Examples

We end this section with examples of direct subjective, expressive subjective element, and objective speech event frames (sans target and attitude type attributes, which are discussed in the next section).

First, we show the frames that would be associated with sentence (9), assuming that the relevant source ID’s have already been defined:

- (9) “The US fears a spill-over,” said Xirao-Nima.

```
Objective speech event:
  Text anchor: the entire sentence
  Source: <writer>
  Implicit: true
```

```
Objective speech event:
  Text anchor: said
  Source: <writer,Xirao-Nima>
Direct subjective:
  Text anchor: fears
  Source: <writer,Xirao-Nima,U.S.>
  Intensity: medium
  Expression intensity: medium
```

The first objective speech event frame represents that, according to the writer, it is true that Xirao-Nima uttered the quote and is a professor at the university referred to. The *implicit* attribute is included because the writer’s speech event is not explicitly mentioned in the sentence (i.e., there is no explicit phrase such as “I write”).

The second objective speech event frame represents that, according to the writer, according to Xirao-Nima, it is true that the US fears a spillover. Finally, when we drill down to the subordinate clause we find a private state: the US fear of a spillover. Such detailed analyses, encoded as annotations on the input text, would enable a person or an automated system to pinpoint the subjectivity in a sentence, and attribute it appropriately.

Now, consider sentence (10):

- (10) “The report is full of absurdities,” Xirao-Nima said.

```
Objective speech event:
  Text anchor: the entire sentence
  Source: <writer>
  Implicit: true
Direct subjective:
  Text anchor: said
  Source: <writer,Xirao-Nima>
  Intensity: high
  Expression intensity: neutral
Expressive subjective element:
  Text anchor: full of absurdities
  Source: <writer,Xirao-Nima>
  Intensity: high
```

The objective frame represents that, according to the writer, it is true that Xirao-Nima uttered the quoted string. The second frame is created for “said” because it is a subjective speech event: private states are conveyed in what is uttered. Note that *intensity* is *high* but *expression intensity* is *neutral*: the private state being expressed is strong, but the specific speech event phrase “said” does not itself contribute to the intensity of the private state. The third frame is for the expressive subjective element “full of absurdities.”

3 Annotation Process

To date, over 11,000 sentences in 550 documents have been annotated according to the annotation scheme described above. The documents are English-language versions of news documents from the world press. The documents are from 187 different news sources in a variety

of countries. The original documents and their annotations are available at <http://nrrc.mitre.org/NRRC/publications.htm>.

The annotation process and inter-annotator agreement studies are described in (Wiebe et al., 2005). Here, we want to highlight two themes of the annotation instructions:

1. There are no fixed rules about how particular words should be annotated. The instructions describe the annotations of specific examples, but do not state that specific words should always be annotated a certain way.
2. Sentences should be interpreted with respect to the contexts in which they appear. The annotators should not take sentences out of context and think what they *could* mean, but rather should judge them as they are being used in that particular sentence and document.

We believe that these general strategies for annotation support the creation of corpora that will be useful for studying expressions of subjectivity in context.

4 Extensions: Attitude and Target Annotations

Before we describe the new attitude and target annotations, consider the following sentence.

(11) “I think people are happy because Chavez has fallen.”

This sentence contains two private states, represented by direct subjective annotations anchored on “think” and “happy,” respectively.

The word “think” is used to express an opinion about what is true according to its source (a *positive arguing* attitude type; see Section 4.1). The target of “think” is “people are happy because Chavez has fallen.”

The word “happy” clearly expresses a positive attitude, with target “Chavez has fallen.” However, looking more closely at the private state for “happy,” we see that we can also infer a negative attitude toward Chavez, from the phrase “happy because Chavez has fallen.”

Sentence (11) illustrates some of the things we need to consider when representing attitudes and targets. First, we see that more than one type of attitude may be involved when a private state is expressed. In (11), there are three (a positive attitude, a negative attitude, and a positive arguing attitude). Second, more than one target may be associated with a private state. Consider “happy” in (11). The target of the positive attitude is “Chavez has fallen,” while the target of the inferred negative attitude is “Chavez.”

Positive Attitudes	Positive Arguing
Negative Attitudes	Negative Arguing
Positive Intentions	Speculation
Negative Intentions	Other Attitudes

Table 1: Attitude Types

The representation also must support multiple targets for a single attitude, as illustrated by Sentence (12):

(12) Tsvangirai said the election result was a clear case of highway robbery by Mugabe, his government and his party, Zanu-PF.

In (12), the phrase “a clear case of highway robbery” expresses a negative attitude of Tsvangirai. This negative attitude has two targets: “the election results” and “Mugabe, his government and his party, Zanu-PF.”

To capture the kind of detailed attitude and target information that we described above, we propose two new types of annotations: *attitude frames* and *target frames*. We describe these new annotations in Sections 4.2 and 4.3, but first we introduce the set of attitude types that we developed for the annotation scheme.

4.1 Types of Attitudes

One of our goals in extending the annotation scheme for private states was to develop a set of attitude types that would be useful for NLP applications. It is also important that the set of attitude types provide good coverage for the range of possible private states. Working with our annotators and looking at the private states already annotated, we developed the set of attitude types listed in Table 1.

Below we give a brief description of each attitude type, followed by an example. In each example, the span of text that expresses the attitude type is in bold, and the span of text that refers to the target of the attitude type (if a target is given) is in angle brackets.

Positive Attitudes: positive emotions, evaluations, judgments and stances.

(13) The Namibians went as far as to say <Zimbabwe’s election system> was “**water tight, without room for rigging**”.

Negative Attitudes: negative emotions, evaluations, judgments and stances.

(14) His disenfranchised supporters **were seething**.

Positive Arguing: arguing for something, arguing that something is true or so, arguing that something did happen or will happen, etc.

(15) Iran **insists** ⟨its nuclear program is purely for peaceful purposes⟩.

Negative Arguing: arguing against something, arguing that something is not true or not so, arguing that something did not happen or will not happen, etc.

(16) Officials in Panama **denied** that ⟨Mr. Chavez or any of his family members had asked for asylum⟩.

Positive Intentions: aims, goals, plans, and other overtly expressed intentions.

(17) The Republic of China government believes in the US **commitment** ⟨to separating its anti-terrorism campaign from the Taiwan Strait issue⟩, an official said Thursday.

Negative Intentions: expressing that something is not an aim, not a goal, not an intention, etc.

(18) The Bush administration **has no plans** ⟨to ease sanctions against mainland China⟩.

Speculation: speculation or uncertainty about what may or may not be true, what may or may not happen, etc.

(19) ⟨The president is **likely** to endorse the bill⟩.

Other Attitudes: other types of attitudes that do not fall into one of the above categories.

(20) To the **surprise** of many, ⟨the dollar hit only 2.4 pesos and closed at 2.1⟩.

4.2 Attitude Frames

With the introduction of the attitude frames, two issues arise. First, which spans of text should the new attitudes be anchored to? Second, how do we tie the attitude frames back to the private states that they are part of?

The following sentence illustrates the first issue.

(21) The MDC leader said systematic cheating, spoiling tactics, rigid new laws, and shear obstruction - as well as political violence and intimidation - were just some of the irregularities practised by the authorities in the run-up to, and during the poll.

In (21), there are 5 private state frames attributed to the MDC leader: a direct subjective frame anchored to “said,” and four expressive subjective element frames anchored respectively to “systematic cheating . . . obstruction,” “as well as,” “violence and intimidation,” and “just some of the irregularities.” We could create an attitude frame for each of these private state frames,

but we believe the following is a better solution. For each direct subjective frame, the annotator is asked to consider the direct subjective annotation and everything within the scope of the annotation when deciding what attitude types are being expressed by the source of the direct subjective frame. Then, for each attitude type identified, the annotator creates an attitude frame and anchors the frame to whatever span of text completely captures the attitude type. In to sentence (21), this results in just one attitude frame being created to represent the negative attitude of the MDC leader. The anchor for this attitude frame begins with “systematic cheating” and ends with “irregularities.”

Turning to the second issue, tying attitude frames to their private states, we do two things. First, we create a unique ID for the attitude frame. Then, we change the attitude type attribute on the direct subjective annotation into a new attribute called an *attitude link*. We place the attitude frame ID into the attitude link slot. The attitude link slot can hold more than one attitude frame ID, allowing us to represent a private state composed of more than one type of attitude.

Because we expect the attitude annotations to overlap with most of the expressive subjective element annotations, we chose not to link attitude frames to expressive subjective element frames. However, this would be possible to do should it become necessary.

The attitude frame has the following attributes:

Attitude frame:

- **id:** a unique alphanumeric ID for identifying the attitude annotation. The ID is used to link the attitude annotation to the private state it is part of.
- **text anchor:** a pointer to the span of text that captures the attitude being expressed.
- **attitude type:** one of the attitude types listed in Table 1.
- **target link:** one or more target annotation IDs (see Section 4.3).
- **intensity:** the intensity of the attitude.
- **properties:**
 - **inferred:** true, if the attitude is inferred.
 - **sarcastic:** true, if the attitude is realized through sarcasm.
 - **repetition:** true, if the attitude is realized through the repetition of words, phrases, or syntax.
 - **contrast:** true, if the attitude is realized only through contrast with another attitude.

Of the four attitude-frame properties, *inferred* was already discussed. The property *sarcastic* marks attitudes expressed using sarcasm. In general, we think this property will be of interest for NLP applications working with opinions. Detecting sarcasm may also help a system learn to distinguish between positive and negative attitudes. The sarcasm in Sentence (22), below, makes the word “Great” an expression of a negative rather than a positive attitude.

(22) “Great, keep on buying dollars so there’ll be more and more poor people in the country,” shouted one.

The *repetition* and *contrast* properties are also for marking different ways in which an attitude might be realized. We feel these properties will be useful for developing an automatic system for recognizing different types of attitudes.

4.3 Target Frames

The target frame is used to mark the target of each attitude. A target frame has two slots, the *id* slot and the *text anchor* slot. The *id* slot contains a unique alpha-numeric ID for identifying the target annotation. We use the target frame ID to link the target back to the attitude frame. The attitude frame has a *target-link* slot that can hold one or more target frame IDs. This allows us to represent when a single attitude is directed at more than one target.

The text anchor slot has a pointer to the span of text that denotes the target. If there is more than one reference to the target in the sentence, the most *syntactically relevant* reference is chosen.

To illustrate what we mean by syntactically relevant, consider the following sentence.

(23) African observers **generally approved** of ⟨his victory⟩ while Western governments **denounced** ⟨it⟩.

The target of the two attitudes (in bold) in the above sentence is the same entity in the discourse. However, although we anchor the target for the first attitude to “his victory,” the anchor for the target of the second attitude is the pronoun “it.” As the direct object of the span that denotes the attitude “denounced,” “it” is more syntactically relevant than “his victory.”

4.4 Illustrative Examples

Figures 4.4 and 4.4 give graphical representations for the annotations in sentences (11) and (12). With attitude frame and target frame extensions, we are able to capture more detail about the private states being expressed in the text than the original core scheme presented in (Wiebe et al., 2005).

5 Pie in the Sky Annotation

Among the items on the “Pie in the Sky” *Check List of Desirable Semantic Information for Annotation*,¹ the most closely related are *epistemic values* (“attitude?”), *epistemic*, *deontic*, and *personal attitudes*. These all fundamentally involve a *self* (Banfield, 1982), a subject of consciousness who is the source of knowledge assessments, judgments of certainty, judgments of obligation/permission, personal attitudes, and so on. Any explicit epistemic, deontic, or personal attitude expressions are represented by us as private state frames, either direct subjective frames (e.g., for verbs such as “know” referring to an epistemic state) or expressive subjective element frames (e.g., for modals such as “must” or “ought to”). Importantly, many deontic, epistemic, and personal attitude expressions do not directly express the speaker or writer’s subjectivity, but are attributed by the speaker or writer to agents mentioned in the text (consider, e.g., “John believes that Mary should quit her job”). Our frame and nested-source representations were designed to support attributing subjectivity to appropriate sources. In future work, additional attributes could be added to private state frames to distinguish between, for example, deontic and epistemic usages of “must” and to represent different epistemic values.

Other phenomena on the list overlap with subjectivity, such as *modality* and *social style/register*. As mentioned above, some modal expressions are subjective, such as those expressing deontic or epistemic judgments. However, hypotheticals and future expressions need not be subjective. For example, “The company announced that if its profits decrease in the next quarter, it will lay off some employees” may easily be interpreted as presenting objective fact. As for style, some are subjective by their nature. One is the literary style *represented thought*, used to present consciousness in fiction (Cohn, 1978; Banfield, 1982). Others are sarcastic or dismissive styles of speaking or writing. In our annotation scheme, sentences perceived to represent a character’s consciousness are represented with private-state frames, as are expressions perceived to be sarcastic or dismissive. On the other hand, some style distinctions, such as degree of formality, are often realized in other ways than with explicit subjective expressions (e.g., “can’t” versus “cannot”).

Polarity, another item on the checklist, also overlaps with subjective positive and negative attitude types. Although many negative and positive polarity words are seldom used outside subjective expressions (such as “hate” and “love”), others often are. For example, words such as “addicted” and “abandoned” are included as negative polarity terms in the General Inquirer lexicon (General-Inquirer, 2000), but they can easily appear in objective

¹Available at: <http://nlp.cs.nyu.edu/meyers/frontiers/2005.html>

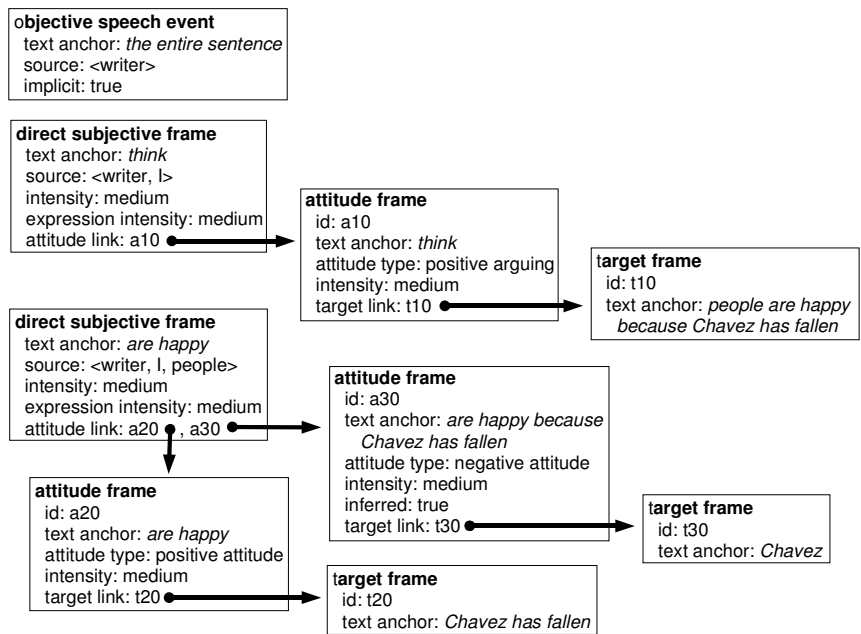


Figure 1: Graphical representation of annotations for Sentence (11)

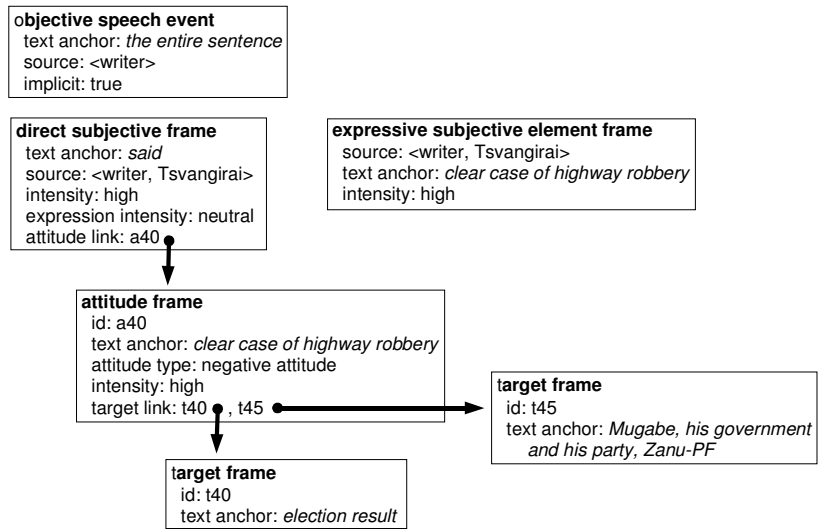


Figure 2: Graphical representation of annotations for Sentence (12)

sentences (e.g., “Thomas De Quincy was addicted to opium and lived in an abandoned shack”).

Integrating subjectivity with other layers of annotation proposed in the “Pie in the Sky” project would afford the opportunity to investigate how they interact. It would also enrich our subjectivity representations. While our scheme promises to be a good base, much remains to be added. For example, annotations of thematic roles and co-reference would add needed structure to the target annotations, which are now only spans of text. In addition, temporal and modal annotations would flesh out the *insubstantial* attribute, which is currently only a binary marker. Furthermore, individual private state expressions must be integrated with respect to the discourse context. For example, which expressions of opinions oppose versus support one another? Which sentences presented as objective fact are included to support a subjective opinion? A challenging dimension to add to the “Pie in the Sky” project would be the *deictic center* as conceived of in (Duchan et al., 1995), which consists of *here*, *now*, and *I* reference points updated as the text or conversation unfolds. Our annotation scheme was developed with this framework in mind.

6 Related Work

The work most similar to ours is Appraisal Theory (Martin, 2000; White, 2002) from systemic functional linguistics (see Halliday (1985/1994)). Both Appraisal Theory and our annotation scheme are concerned with identifying and characterizing expressions of opinions and emotions in context. The two schemes, however, make different distinctions. Appraisal Theory distinguishes different types of positive and negative attitudes and also various types of “intersubjective positioning” such as attribution and expectation. Appraisal Theory does not distinguish, as we do, the different ways that private states may be expressed (i.e., directly, or indirectly using expressive subjective elements). It also does not include a representation for nested levels of attribution.

In addition to Appraisal Theory, subjectivity annotation of text in context has also been performed in Yu and Hatzivassiloglou (2003), Bruce and Wiebe (1999), and Wiebe et al. (2004). The annotations in Yu and Hatzivassiloglou (2003) are sentence-level subjective vs. objective and polarity judgments. The annotation schemes used in Bruce and Wiebe (1999) and Wiebe et al. (2004) are earlier, much less detailed versions of the annotation scheme presented in this paper.

7 Conclusion

We have described extensions to an annotation scheme for private states and objective speech events in language. We look forward to integrating and elaborating

this scheme with other layers of semantic meaning in the future.

8 Acknowledgments

This work was supported in part by the National Science Foundation under grant IIS-0208798 and by the Advanced Research and Development Activity (ARDA).

References

- A. Banfield. 1982. *Unspeakable Sentences*. Routledge and Kegan Paul, Boston.
- R. Bruce and J. Wiebe. 1999. Recognizing subjectivity: A case study of manual tagging. *Natural Language Engineering*, 5(2):187–205.
- D. Cohn. 1978. *Transparent Minds: Narrative Modes for Representing Consciousness in Fiction*. Princeton University Press, Princeton, NJ.
- J. Duchan, G. Bruder, and L. Hewitt, editors. 1995. *Deixis in Narrative: A Cognitive Science Perspective*. Lawrence Erlbaum Associates.
- The General-Inquirer. 2000. http://www.wjh.harvard.edu/~inquirer/spreadsheet_guide.htm.
- M.A.K. Halliday. 1985/1994. *An Introduction to Functional Grammar*. London: Edward Arnold.
- J.R. Martin. 2000. Beyond exchange: APPRAISAL systems in English. In Susan Hunston and Geoff Thompson, editors, *Evaluation in Text: Authorial stance and the construction of discourse*, pages 142–175. Oxford: Oxford University Press.
- R. Quirk, S. Greenbaum, G. Leech, and J. Svartvik. 1985. *A Comprehensive Grammar of the English Language*. Longman, New York.
- P.R.R. White. 2002. Appraisal: The language of attitudinal evaluation and intersubjective stance. In Verschueren, Ostman, blommaert, and Bulcaen, editors, *The Handbook of Pragmatics*, pages 1–27. Amsterdam/Philadelphia: John Benjamins Publishing Company.
- J. Wiebe, T. Wilson, R. Bruce, M. Bell, and M. Martin. 2004. Learning subjective language. *Computational Linguistics*, 30(3):277–308.
- J. Wiebe, T. Wilson, and C. Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation (formerly Computers and the Humanities)*, 1(2).
- T. Wilson and J. Wiebe. 2003. Annotating opinions in the world press. In *SIGdial-03*.
- H. Yu and V. Hatzivassiloglou. 2003. Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. In *EMNLP-2003*.