

Annotating Expressions of Opinions and Emotions in Language

To appear in *Language Resources and Evaluation 1(2)*

*This is not the final version**

Janyce Wiebe (wiebe@cs.pitt.edu)

Department of Computer Science, University of Pittsburgh

Theresa Wilson (twilson@cs.pitt.edu)

Intelligent Systems Program, University of Pittsburgh

Claire Cardie (cardie@cs.cornell.edu)

Department of Computer Science, Cornell University

June 2004

Abstract. This paper describes a corpus annotation project to study issues in the manual annotation of opinions, emotions, sentiments, speculations, evaluations and other private states in language. We present the resulting corpus annotation scheme as well as examples of its use. In addition, we describe the manual annotation process and the results of an interannotator agreement study on a 10,000-sentence corpus of articles drawn from the world press.

Keywords: affect, attitudes, corpus annotation, emotion, natural language processing, opinions, sentiment, subjectivity

1. Introduction

There has been a recent swell of interest in the automatic identification and extraction of opinions, emotions, and sentiments in text. Motivation for this task comes from the desire to provide tools for information analysts in government, commercial, and political domains, who want to automatically track attitudes and feelings in the news and on-line forums. How do people feel about recent events in the Middle East? Is the rhetoric from a particular opposition group intensifying? What is the range of opinions being expressed in the world press about the best

* This work was supported in part by the National Science Foundation under grants IIS-0208798 and IIS-0208028, the Advanced Research and Development Activity (ARDA)'s Advanced Question Answering for Intelligence (AQUAINT) Program, and by the Northeast Regional Research Center (NRRC) which is sponsored by the Advanced Research and Development Activity (ARDA), a U.S. Government entity which sponsors and promotes research of import to the Intelligence Community which includes but is not limited to the CIA, DIA, NSA, NIMA, and NRO.



course of action in Iraq? A system that could automatically identify opinions and emotions from text would be an enormous help to someone trying to answer these kinds of questions.

Researchers from many subareas of Artificial Intelligence and Natural Language Processing have been working on the automatic identification of opinions and related tasks (e.g., Pang et al. (2002), Dave et al. (2003), Gordon et al. (2003), Riloff and Wiebe (2003), Riloff et al. (2003), Turney and Littman (2003), Yi et al. (2003), and Yu and Hatzivassiloglou (2003)). To date, most such work has focused on sentiment or subjectivity classification at the document or sentence level. Document classification tasks include, for example, distinguishing editorials from news articles and classifying reviews as positive or negative (Wiebe et al., 2001b; Pang et al., 2002; Yu and Hatzivassiloglou, 2003). A common sentence-level task is to classify sentences as subjective or objective (Yu and Hatzivassiloglou, 2003; Riloff et al., 2003).

However, for many applications, just identifying opinionated documents or sentences may not be sufficient. In the news, it is not uncommon to find two or more opinions in a single sentence, or to find a sentence containing opinions as well as factual information. *Information extraction (IE)* systems are natural language processing (NLP) systems that extract from text any information relevant to a pre-specified topic. An IE system trying to distinguish between factual information (which should be extracted) and non-factual information (which should be discarded or labeled uncertain) would find it helpful to be able to pinpoint the particular clauses that contain opinions. This ability would also be important for multi-perspective question answering systems, which aim to present to the user multiple answers to non-factual questions based on opinions derived from different sources; and for multi-document summarization systems, which need to summarize differing opinions and perspectives.

Many applications would benefit from being able to determine not just whether a document or text snippet is opinionated but also the *strength* of the opinion. Flame detection systems, for example, want to identify strong rants and emotional tirades, while letting milder opinions pass through (Spertus, 1997; Kaufer, 2000). In addition, information analysts need to recognize changes over time in the degree of virulence expressed by persons or groups of interest, and to detect when rhetoric is heating up, or cooling down. Furthermore, knowing the *types of attitude* being expressed (e.g., positive versus negative evaluations) would enable a natural language processing (NLP) application to target particular types of opinions.

Very generally then, we assume that the existence of corpora annotated with rich information about opinions and emotions would support

the development and evaluation of NLP systems that exploit such information. In particular, statistical and machine learning approaches have become the method of choice for constructing a wide variety of practical NLP applications. These methods, however, typically require training and test corpora that have been manually annotated with respect to each language-processing task to be acquired.

The high-level goal of this paper, therefore, is to investigate the use of opinion and emotion in language through a corpus annotation study. In particular, we propose a detailed annotation scheme that identifies key components and properties of opinions, emotions, sentiments, speculations, evaluations, and other *private states* (Quirk et al., 1985), i.e., internal states that cannot be directly observed by others. We argue, through the presentation of numerous examples, that this annotation scheme covers a broad and useful subset of the range of linguistic expressions and phenomena employed in naturally occurring text to express opinion and emotion.

We propose a relatively fine-grained annotation scheme, annotating text at the word- and phrase-level rather than at the level of the document or sentence. For every expression of a private state in each sentence, a *private state frame* is defined. A private state frame includes the source of the private state (i.e., whose private state is being expressed), the target (i.e., what the private state is about), and various properties involving strength, significance, and type of attitude. An important property of sources in the annotation scheme is that they are *nested*, reflecting the fact that private states and speech events are often embedded in one another. The representation scheme also includes frames representing material that is attributed to a source, but presented objectively, without evaluation, speculation, or other type of private state by that source.

The annotation scheme has been employed in the manual annotation of a 10,000-sentence corpus of articles from the world press.¹ We describe the annotation procedure in this paper, and present the results of an interannotator agreement study.

A focus of this work is identifying private state expressions in context, rather than judging words and phrases themselves, out of context. That is, the annotators are not presented with word- or phrase-lists to judge (as in, e.g., Osgood et al. (1957), Heise (1965; 2001), and Subasic and Huettner (2001)). Further, the annotation instructions do not specify how specific words should be annotated, and the annotators were not limited to marking any particular words, parts of

¹ The corpus is freely available at: <http://nrrc.mitre.org/NRRC/publications.htm>. To date, *target* annotations are included for only a subset of the sentences, as specified in Section 2.2.

speech, or grammatical categories. Consequently, a tremendous range of words and constituents were marked by the annotators, not only adjectives, modals, and adverbs, but also verbs, nouns, and various types of constituents. The contextual nature of the annotations makes the annotated data valuable for studying ambiguities that arise with subjective language. Such ambiguities range from word-sense ambiguity (e.g., objective senses of *interest* as in *interest rate* versus subjective senses as in *take an interest in*), to ambiguity in idiomatic versus non-idiomatic usages (e.g., *bombed* in *The comedian really bombed last night* versus *The troops bombed the building*), to various pragmatic ambiguities involving irony, sarcasm, and metaphor.

To date, the annotated data has served as training and testing data in opinion extraction experiments classifying sentences as subjective or objective (Riloff et al., 2003; Riloff and Wiebe, 2003) and in experiments classifying the strengths of private states in individual clauses (Wilson et al., 2004). However, these experiments abstracted away from the details in the annotation scheme, so there is much room for additional experimentation in the automatic extraction of private states, and in exploiting the information in NLP applications.

The remainder of this paper is organized as follows. Section 2 gives an overview of the annotation scheme, ending with a short example. Section 3 elaborates on the various aspects of the annotation scheme, providing motivations, examples, and further clarifications; this section ends with an extended example, which illustrates the various components of the annotation scheme and the interactions among them. Section 4 presents observations about the annotated data. Section 5 describes the corpus and 6 presents the results of an interannotator agreement study. Section 7 discusses related work, Section 8 discusses future work, and Section 9 presents conclusions.

2. Overview of the Annotation Scheme

2.1. MEANS OF EXPRESSING PRIVATE STATES.

The goal of the annotation scheme is to represent internal mental and emotional states. 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.”

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 expressed 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**.

In this work, the term *speech event* is used to refer to any speaking or writing event. A speech event has a writer or speaker as well as a target, which is whatever is written or said.

The phrase “full of absurdities” in (2) above is an *expressive subjective element* (Banfield, 1982). There are a number of additional examples of expressive subjective elements in sentences (3) and (4):

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

(4) “We foresaw electoral **fraud** but not **daylight robbery**,” Tsvangirai said.

The private states in these sentences 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. In sentence (4), describing the election as “daylight robbery” clearly reflects the anger being experienced by the speaker, Tsvangirai. As used in these sentences, the phrases “The time has come,” “gentlemen,” “the assassin,” “injustice cannot last long,” “fraud,” and “daylight robbery” are all expressive subjective elements. Expressive subjective elements are used by people to express their frustration, anger, wonder, positive sentiment, mirth, etc., without explicitly stating that they are frustrated, angry, etc. Sarcasm and irony often involve expressive subjective elements.

As mentioned above, “full of absurdities” in (2) is an expressive subjective element. In fact, two private state frames are created for

sentence (2): one for the speech event and one for the expressive subjective element. The first represents the more general fact that private states are expressed in what was said; the second pinpoints a specific expression used to express Tsvangairai’s negative evaluation.

In the subsections below, we describe how private states, speech events, and expressive subjective elements are explicitly mapped onto components of the annotation scheme.

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. Direct subjective expressions are typically more explicit than expressive subjective element expressions, which is reflected in the fact that direct subjective frames contain more attributes than expressive subjective element frames. Specifically, the frames have the following attributes:

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

- **text span:** a pointer to the span of text that represents the speech event or explicit mention of a private state. (Text spans are described more fully in Section 3.1.)
- **source:** the person or entity that is expressing the private state, possibly the writer. (See Sections 2.5 and 3.2 for more information on sources.)
- **target:** the target or topic of the private state, i.e., what the speech event or private state is about. To date, our corpus includes only targets that are agents (see Section 2.4) and that are targets of negative or positive private states (see Section 4.4).
- **properties:**
 - **strength:** the strength of the private state (*low*, *medium*, *high*, or *extreme*). (The *strength* attribute is described further in Section 3.4.)
 - **expression strength:** the contribution of the speech event or private state expression itself to the overall strength of the private state (*neutral*, *low*, *medium*, *high*, or *extreme*.) For example, *say* is often neutral, even if what is uttered is not

neutral, while *excoriate* itself implies a very strong private state. (The *expression-strength* property will be described in more detail in Section 3.4.)

- **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*. (This attribute is described in more detail in Section 3.6.)
- **attitude type:** This attribute currently represents the *polarity* of the private state. The possible values are *positive*, *negative*, *other*, or *none*. In ongoing work, we are developing a richer set of attitude types to make more fine-grained distinctions (see Section 8).

Expressive subjective element frame:

- **text span:** 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:**
 - **strength:** the strength of the private state (*low*, *medium*, *high*, or *extreme*)
 - **attitude type:** This attribute represents the *polarity* of the private state. The possible values are *positive*, *negative*, *other*, or *none*.

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, namely the *span*, *source*, and *target* slots.

Objective speech event frame:

- **text span:** 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. To date, targets of objective speech event frames are not yet annotated in our corpus.

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

(5) 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 *span* 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.

For example, suppose that *nima* is the ID created for Xirao-Nima in a document that quotes him. Consider the following consecutive sentences from that document:

(6) “I have been to Tibet many times. I have seen the truth there, which is very different from what some US politicians with ulterior motives have described,” said Xirao-Nima, who is a Tibetan.

(7) Some Westerners who have been there have also seen the ever-improving human rights in the Tibet Autonomous Region, he added.

The following agent frames are created for the references to Xirao-Nima in these sentences:

Agent:

Span: Xirao-Nima in (6)

Source: nima

Agent:

Span: he in (7)

Source: nima

The connection between agent frames and the *source* slots of the private state and objective speech event frames will be explained in the following subsection.

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. Obviously, the writer of an article is a source, because he or she wrote the sentences composing 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 (8) and of private states in (9) and (10)).

(8) Sue said, "The election was fair."

(9) Sue thinks that the election was fair.

(10) 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. Sentence (8), for example, 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 the following sentences (the first is sentence (1), reprinted here):

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

(11) 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* $\langle \textit{writer}, \textit{Xirao-Nima}, \textit{U.S.} \rangle$. In sentence (11), the U.S. report's criticism is the target of China's criticism. Thus, the nested source for "criticism" is $\langle \textit{writer}, \textit{China}, \textit{U.S. report} \rangle$.

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 $\langle \textit{writer}, \textit{China}, \textit{U.S. report} \rangle$ 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. Throughout this paper, targets are indicated only in cases where the targets are agents and are the targets of positive or negative private states, as those are the targets labeled in our annotated corpus.

First, we show the frames that would be associated with sentence (12), assuming that the relevant source ID's have already been defined:

(12) “The US fears a spill-over,” said Xirao-Nima, a professor of foreign affairs at the Central University for Nationalities.

Objective speech event:

Span: the entire sentence
Source: <writer>
Implicit: true

Objective speech event:

Span: said
Source: <writer,Xirao-Nima>

Direct subjective:

Span: fears
Source: <writer,Xirao-Nima,U.S.>
Strength: medium
Expression strength: medium
Attitude type: negative

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 (13):

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

Objective speech event:

Span: the entire sentence
 Source: <writer>
 Implicit: true

Direct subjective:

Span: said
 Source: <writer,Xirao-Nima>
 Strength: high
 Expression strength: neutral
 Target: report
 Attitude type: negative

Expressive subjective element:

Span: full of absurdities
 Source: <writer, Xirao-Nima>
 Strength: high
 Attitude type: negative

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 *strength* is *high* but *expression strength* is *neutral*: the private state being expressed is strong, but the specific speech event phrase “said” does not itself contribute to the strength of the private state. The third frame is for the expressive subjective element “full of absurdities.”

3. Elaborations and Illustrations

3.1. TEXT SPANS IN DIRECT SUBJECTIVE AND OBJECTIVE SPEECH EVENT FRAMES

All frames in the private state annotation scheme are directly encoded as XML annotations on the underlying text. In particular, each XML annotation frame is anchored to a particular location in the underlying text via its associated text span slot. This section elaborates on the appropriate text spans to include in direct subjective and objective speech event frames and further explains the notion of an “implicit” speech event.

Consider a sentence that explicitly presents a private state or speech event. For the discussion in this subsection, it will be useful to distinguish between the following:

The private state or speech event phrase: For private states, this is the text span that designates the attitude (or attitudes) being expressed. For speech event phrases, this is the text span that refers to the speaking or writing event.

The subordinated constituents: The constituents of the sentence that are inside the scope of the private state or speech event phrase. This is the text span that designates the target.

Consider sentence (14):

(14) “It is heresy,” said Cao, “the ‘Shouters’ claim they are bigger than Jesus.”

First consider the writer’s top-level speech event (i.e., the writing of the sentence itself). The source and speech event phrase are implicit; that is, we understand the sentence as implicitly in the scope of “I write that ...” or “According to me ...”. Thus, the entire sentence is subordinated to the (implicit) speech event phrase.

Now consider Cao’s speech event:

- **Source:** $\langle \textit{writer}, \textit{Cao} \rangle$
- **private state or speech event phrase:** “said”
- **Subordinated constituents:** “It is heresy”; “the ‘Shouters’ claim they are bigger than Jesus.”

Finally, we have the Shouters’ claim:

- **Source:** $\langle \textit{writer}, \textit{Cao}, \textit{Shouters} \rangle$
- **private state or speech event phrase:** “claim”
- **Subordinated constituents:** “they are bigger than Jesus”

For sentences that explicitly present a private state or speech event, the *span* slot is filled with the private state or speech event phrase. Moreover, in the underlying text-based representation, the XML annotation for the private state is anchored on the private state or speech event phrase.

It is less clear what *span* to associate when the private state or speech event phrase is implicit, as was the case for the writer’s top-level speech

event in sentence (14). Since the phrase is implicit, it cannot serve as the anchor in the underlying representation. A similar situation arises when direct quotes are not accompanied by discourse parentheticals (such as “, she said”). An example is the second sentence in the following passage:

(15) “We think this is an example of the United States using human rights as a pretext to interfere in other countries’ internal affairs,” Kong said. “We have repeatedly stressed that no double standard should be employed in the fight against terrorism.”

In these cases, we opted to make the entire sentence or quoted string the *span* for the frame (and to anchor the annotation on the sentence or quoted string, in the text-based XML representation).²

Currently, the subordinated constituents are not explicitly encoded in the annotation scheme.

3.2. NESTED SOURCES

Although the nested source examples in Section 2 were fairly simple in nature, the nesting of sources may be quite deep and complex in practice. For example, consider sentence (16):

(16) The Foreign Ministry said Thursday that it was “surprised, to put it mildly” by the U.S. State Department’s criticism of Russia’s human rights record and objected in particular to the “odious” section on Chechnya.

There are three sources in this sentence: the writer, the Foreign Ministry, and the U.S. State Department. The writer is the source of the overall sentence. The remaining explicitly mentioned private states and speech events in (16) have the following nested sources:

Speech event “said”:

- **Source:** \langle writer, Foreign Ministry \rangle
The relevant part of the sentence for identifying the source is “**The foreign ministry** said ...”

Private state “surprised, to put it mildly”:

- **Source:** \langle writer, Foreign Ministry, Foreign Ministry \rangle
The relevant part of the sentence for identifying the source is “**The foreign ministry** said **it** was surprised, to put it mildly ...”

² In addition, an *implicit* attribute is added to the frame, to record the fact that the speech event phrase was implicit. Sources that are implicit are also marked as *implicit*.

The *Foreign Ministry* appears twice because its “surprised” private state is nested in its “said” speech event. Note that the entire string “surprised, to put it mildly” is the private state phrase, rather than only “surprised”, because “to put it mildly” intensifies the private state. The Foreign Ministry is not only surprised, it is very surprised. As shown below, “to put it mildly” is also an expressive subjective element.

Private state “criticism”:

- **Source:** $\langle \textit{writer}, \textit{Foreign Ministry}, \textit{Foreign Ministry}, \textit{U.S. State Department} \rangle$
The relevant part of the sentence for identifying the source is “**The foreign ministry** said it was surprised, to put it mildly by the **U.S. State Department’s** criticism . . .”

Private state/speech event “objected”:

- **Source:** $\langle \textit{writer}, \textit{Foreign Ministry} \rangle$
The relevant part of the sentence for identifying the source is “**The foreign ministry** . . . objected . . .” To see that the source contains only *writer* and *Foreign Ministry*, note that the sentence is a compound sentence, and that “objected” is not in the scope of “said” or “surprised.”

Expressive subjective elements also have nested sources. The expressive subjective elements in (16) have the following sources:

Expressive subjective element “to put it mildly”:

- **Source:** $\langle \textit{writer}, \textit{Foreign Ministry} \rangle$
The Foreign Ministry uses a subjective intensifier, “to put it mildly”, to express sarcasm while describing its surprise. This subjectivity is at the level of the Foreign Ministry’s speech, so the source is $\langle \textit{writer}, \textit{Foreign Ministry} \rangle$ rather than $\langle \textit{writer}, \textit{Foreign Ministry}, \textit{Foreign Ministry} \rangle$.

Expressive subjective element “odious”:

- **Source:** $\langle \textit{writer}, \textit{Foreign Ministry} \rangle$
The word “odious” is not within the scope of the “surprise” private state, but rather attaches to the “objected” private state/speech event. Thus, as for “to put it mildly,” the source is nested two levels, not three.

As we can see in the frames above, the expressive subjective elements in (16) have the same nested sources as their immediately dominating

private state or speech terms (i.e., “to put it mildly” and “said” have the same nested source; and “odious” and “objected” have the same nested source). However, expressive subjective elements might attach to higher-level speech events or private states.³ For example, consider “bigger than Jesus” in the following sentence from a Chinese news article:

(14) “It is heresy,” said Cao, “the ‘Shouters’ claim they are bigger than Jesus.”

The nested source of the subjectivity expressed by “bigger than Jesus” is $\langle \text{writer}, \text{Cao} \rangle$, while the nested source of “claim” is $\langle \text{writer}, \text{Cao}, \text{Shouters} \rangle$. In particular, the Shouters aren’t really making this claim in the text; instead, it seems clear from the sentence that it’s Cao’s interpretation of the situation that comprises the “claim.”

3.3. SPEECH EVENTS

This section focuses on the distinction between objective speech events, and subjective speech events (which, recall, are represented by direct subjective frames). To help the reader understand the distinction being made, we first give examples of subjective versus objective speech events, including explicit speech events as well as implicit speech events attributed to the writer. Next, the distinction is more formally specified. Finally, we discuss an interesting context-dependent aspect of the subjective versus objective distinction.

The following two sentences illustrate the distinction between subjective and objective speech events when the speech event term is explicit. Note that, in both sentences, the speech event term is “said,” which itself is neutral.

(4) “We foresaw electoral fraud but not daylight robbery,” Tsvangirai said.

(17) Medical Department head Dr Hamid Saeed said the patient’s blood had been sent to the Institute for Virology in Johannesburg for analysis.

In both cases, the writer’s top-level speech event is represented with an objective speech event frame (that someone said something is presented as objectively true). Of interest to us here are the explicit speech events referred to with “said.” The one in sentence (4) is opinionated. Its representation is a direct subjective frame with an *expression strength* rating of *neutral*, but a *strength* rating of *high*, reflecting the strong negative

³ As discussed in Wiebe (1991), this mirrors the *de re/de dicto* ambiguity of references in opaque contexts (Castaneda, ; Quine, 1976; Fodor, 1979).

evaluation expressed by *Tsvangirai*. In contrast, the information in (17) is simply presented as fact, and the speech event referred to by “said” is represented with an *objective speech event* frame (which contains no strength ratings).

The following two sentences illustrate the distinction between *implicit* subjective and objective speech events attributed to the writer:

(18) The report is flawed and inaccurate.

(19) Bell Industries Inc. increased its quarterly to 10 cents from 7 cents a share.

Consider the frames created for the writer’s top-level speech events in (18) and (19). The frame for (18) is a direct subjective frame, reflecting the writer’s negative evaluations of the report. In contrast, the frame for (19) is an objective speech event frame, because the sentence describes an event presented by the writer as true (assuming nothing in context suggests otherwise).

When the speech event term is neutral, as in (4) and (17), or if there isn’t an explicit speech event term, as in (18) and (19), whether the speech event is subjective or objective depends entirely on the context and the presence or absence of expressive subjective elements.

Let us consider more formally the distinction between subjective and objective speech events. Suppose that the annotator has identified a speech event S with nested source $\langle X_1, X_2, X_3 \rangle$. The critical question is, according to X_1 , according to X_2 , does S express X_3 ’s private state? If yes, the speech event is subjective (and a direct subjective frame is used). Otherwise, it is objective (and an objective speech event frame is used). Note that the frames for a given sentence may be mixtures of subjective and objective speech events. For example, the frames for sentence (2) given in Section 2.6 above include an objective speech event frame for the writer’s top-level speech event (the writer presents it as true that Xirao-Nima uttered the quoted string), as well as a direct subjective frame for Xirao-Nima’s speech event (Xirao-Nima expresses negative evaluation — that the report is full of absurdities — in his utterance). Note also that, even if a speech event is subjective, it may still express something the immediate source believes is true. Consider the sentence “John criticized Mary for smoking.” According to the writer, John expresses a private state (his negative evaluation of Mary’s smoking). However, this does not mean that, according to the writer, John does not believe that Mary smokes.

We complete this subsection with a discussion of an interesting class of subjective speech events, namely those expressing facts and claims that are disputed in the context of the article. Consider the statement “Smoking causes cancer.” In some articles, this speech event would

be objective, while in others, it would be subjective. The annotation frames selected to represent *smoking causes cancer* should reflect the status of the proposition in the article. In a modern scientific article, it is likely to be treated as an undisputed fact. However, in an older article, in which views of scientists and tobacco executives are given, it may be a fact under dispute. When the proposition is disputed, the speech event is represented as subjective. Even if only the views of the scientists or only those of the tobacco executives are explicitly given in the article, a subjective representation might still be the appropriate one. It would be the appropriate representation if, for example, the scientists are arguing against the idea that smoking does not cause cancer. The scientists would be going beyond simply presenting something they believe is a fact; they would be arguing against an alternative view, and for the truthfulness of their own view.

3.4. STRENGTH RATINGS

Strength ratings are included in the annotation scheme to indicate the strengths of the private states expressed in subjective sentences. This is an informative feature in itself; for example, strength would be informative for distinguishing inflammatory messages from reasoned arguments, and for recognizing when rhetoric is ratcheting up or toning down in a particular forum. In addition, strength ratings help in distinguishing borderline cases from clear cases of subjectivity and objectivity: the difference between no subjectivity and a low-strength private state might be highly debatable, but the difference between no subjectivity and a medium or high-strength private state is often much clearer. The annotation study presented below in Section 6 provides evidence that annotator agreement is quite high concerning which are the clear cases of subjective and objective sentences.

As described in Section 2.2, all subjective frames (both expressive subjective element and direct subjective frames) include a strength rating reflecting the overall strength of the private state represented by the frame. The values are *low*, *medium*, *high*, and *extreme*.

For direct subjective frames, there is an additional strength rating, namely the *expression strength*, which deserves additional explanation. The *expression strength* attribute represents the contribution to strength made specifically by the private state or speech event phrase. For example, the *expression strength* of *said*, *added*, *told*, *announce*, and *report* is typically neutral, the *expression strength* of *criticize* is typically medium, and the *expression strength* of *vehemently denied* is typically high or extreme.

3.5. MIXTURES OF PRIVATE STATES AND SPEECH EVENTS.

This section notes something the reader may have noticed earlier: many speech event terms imply mixtures of private states and speech. Examples are *berate*, *object*, *praise*, and *criticize*. This had two effects on the development of our annotation scheme. First, it motivated the decision to use a single frame type, direct subjective, for both subjective speech events and explicit private states. With a single frame type, there is no need to classify an expression as either a speech event or a private state.

Second, it motivated, in part, our inclusion of the *expression strength* attribute described in the previous subsection. Purely speech terms are typically assigned *expression strength* of *neutral*, while mixtures of private states and speech events, such as *criticize* and *praise*, are typically assigned a rating between *low* and *extreme*.

3.6. INSUBSTANTIAL PRIVATE STATES AND SPEECH EVENTS

Recall that direct subjective frames can include the *insubstantial* attribute. This section provides additional discussion regarding the use of this attribute and gives examples illustrating situations in which it is included.

The motivation for including the *insubstantial* attribute is that some NLP applications might need to identify all private state and speech event expressions in a document (for example, systems performing lexical acquisition to populate a dictionary of subjective language), while others might want to find only those opinions and other private states that are substantial in the discourse (for example, summarization and question answering systems). The *insubstantial* attribute allows applications to choose which they want: all private states, or just those whose frames have the value *false* for the *insubstantial* attribute.

There are two cases of *insubstantial* frames, corresponding to the following two dictionary meanings of *insubstantial*:

- (1) Lacking substance or reality and (2) Negligible in size or amount.
[*The American Heritage Dictionary of the English Language*, Fourth Edition, Copyright 2000 by Houghton Mifflin Company]

Thus, the *insubstantial* attribute is true in direct subjective frames whose private states are either (1) not real or (2) not significant.

Let us first consider private states that are *insubstantial* because they are not “real.” A “real” speech event or private state is presented as an existing event within the domain of discourse, e.g., it is not hypothetical. For speech events and private states that are not real, the

presupposition that the event occurred or the state exists is removed via the context (or, the event or state is explicitly asserted not to exist).

The following sentences all contain one or more private states or speech events that are not real under our criterion (highlighted in bold).

(20) If the Europeans **wish** to influence Israel in the political arena...
(in a conditional, so not real)

(21) “And we are seeking **a declaration** that the British government **demand**s that Abbasi should not face trial in a military tribunal with the death penalty.” (not real, i.e., the declaration of the demand is just being sought)

(22) No one who has ever studied realist political science **will find this surprising**. (not real since a specific “surprise” state is not referred to; note that the subject noun phrase is attributive rather than referential (Donnellan, 1966))

Of course, the criterion refers not to actual reality, but rather reality within the domain of discourse. Consider the following sentence from a novel about an imaginary world:

(23) “It’s wonderful!” said Dorothy. [*Dorothy and the Wizard in Oz*, L. Frank Baum, Chapter 2].

Even though Dorothy and the world of Oz don’t exist, Dorothy does utter “It’s wonderful” in that world, which expresses her private state. Thus, the *insubstantial* attribute in the frame for “said” in (23) would be false.

We now turn to private states that are *insubstantial* because they are “not significant.” By “not significant” we mean that the sentence in which the private state is marked does not contain a significant portion of the contents (target) of the private state or speech event. Consider the following sentence:

(24) **Such wishful thinking** risks making the US an accomplice in the destruction of human rights. (not significant)

There are two private state frames created for “such wishful thinking” in (24):

Expressive subjective element:

Span: Such wishful thinking

Source: <writer>

Strength: medium

Attitude type: negative

Direct subjective:

Span: Such wishful thinking
 Source: <writer,US>
 Insubstantial: true
 Strength: medium
 Expression strength: medium

The first frame represents the writer’s negative subjectivity in describing the US’s view as “such wishful thinking” (note that the source is simply $\langle writer \rangle$). The second frame is the one of interest in this subsection: it represents the US’s “thinking” private state (attributed to it by the writer, hence the nested source $\langle writer, US \rangle$). The *insubstantial* attribute for the frame is *true* because the sentence does not present the contents of the private state; it does not identify the US view which the writer thinks is merely “wishful thinking.” The presence of this attribute serves as a signal to NLP systems that this sentence is not informative with respect to the contents of the US’s “thinking” private state.

3.7. PRIVATE STATE ACTIONS

Thus far, we have seen private states expressed in text via a speech event or by expressions denoting subjectivity, emotion, etc. Occasionally, however, private states are expressed by direct physical actions. Such actions are called *private state actions* (Wiebe, 1994). Examples include booing someone, sighing heavily, shaking ones fist angrily, waving ones hand dismissively, smirking, and frowning. “Applaud” in sentence (25) is an example of a positive-evaluative private state action.

(25) As the long line of would-be voters marched in, those near the front of the queue began to spontaneously applaud those who were far behind them.

Because private state actions are not common, we did not introduce a distinct type of frame into the annotation scheme for them. Instead, they are represented using direct subjective frames.

3.8. EXTENDED EXAMPLE

This section gives the speech event and private states frames for a passage from an article from the Beijing China Daily:

(Q1) As usual, the US State Department published its annual report on human rights practices in world countries last Monday.

(Q2) And as usual, the portion about China contains little truth and many absurdities, exaggerations and fabrications.

(Q3) Its aim of the 2001 report is to tarnish China's image and exert political pressure on the Chinese Government, human rights experts said at a seminar held by the China Society for Study of Human Rights (CSSHR) on Friday.

(Q4) "The United States was slandering China again," said Xirao-Nima, a professor of Tibetan history at the Central University for Nationalities.

[Beijing China Daily. (Internet Version-WWW) in English 11 Mar 02. Article by Xiao Xin from the "Opinion" page: "US human rights report defies truth".]

Sentence (Q1) is an objective sentence without speech events or private states (other than the writer's top-level speech event). Though a report is referred to, the sentence is about publishing the report, rather than what the report says.

Objective speech event:

Span: the entire sentence (Q1)
 Source: <writer>
 Implicit: true

Sentence (Q2) (reprinted here for convenience) expresses the writer's subjectivity:

(Q2) And as usual, the portion about China contains little truth and many absurdities, exaggerations and fabrications.

Thus, the top-level speech event is represented with a direct subjective frame:

Direct subjective:

Span: the entire sentence (Q2)
 Source: <writer>
 Implicit: true
 Strength: high
 Attitude type: negative
 Target: report

The frames for the individual subjective elements in (Q2) are the following:

Expressive subjective element:

Span: And as usual
 Source: <writer>
 Strength: low
 Attitude type: negative

Expressive subjective element:

Span: little truth
 Source: <writer>
 Strength: medium
 Attitude type: negative

Expressive subjective element:

Span: many absurdities, exaggerations and fabrications
 Source: <writer>
 Strength: high
 Attitude type: negative

The annotator who labeled this sentence identified three distinct subjective elements in the sentence. The first one, “And as usual”, is interesting because its subjectivity is highly contextual. The subjectivity is amplified by the fact that “as usual” is repeated from the sentence before. The third expressive subjective element, “many absurdities, exaggerations and fabrications,” could have been divided into multiple frames; annotators vary in the extent to which they identify long subjective elements or divide them into sequences of shorter ones. To date, we have not focused on achieving agreement with respect to the boundaries of expressive subjective elements. As described below in Section 6, the spans of subjective elements are considered to be matches if they overlap one another.

Sentence (Q3) is a mixture of private states and speech events at different levels:

(Q3) Its aim of the 2001 report is to tarnish China’s image and exert political pressure on the Chinese Government, human rights experts said at a seminar held by the China Society for Study of Human Rights (CSSHR) on Friday.

The entire sentence is attributed to the writer. The quoted content is attributed by the writer to the human rights experts, so the *source* for that speech event is *<writer, human rights experts>*. In addition, another level of nesting is introduced, with *source <writer, human rights experts, report>*, because a private state of the report is presented, namely

that the report has the aim to tarnish China's image and exert political pressure (according to the writer, according to the human rights experts). The specific frames created for the sentence are as follows.

The writer's speech event is represented with an objective speech event frame: the writer presents it as true, without emotion or other type of private state, that human rights experts said something at a particular location on a particular day.

Objective speech event:

Span: the entire sentence (Q3)
 Source: <writer>
 Implicit: true

Next we have the frame representing the human rights experts' speech:

Direct subjective:

Span: said
 Source: <writer,human rights experts>
 Strength: medium
 Expression strength: neutral
 Target: report
 Attitude type: negative

Note that a direct subjective rather than an objective speech event frame is used. The reason is that, in the context of the article, saying that the aim of the report is to tarnish China's image is argumentative. This is an example of a speech event being classified as subjective because the claim is controversial or disputed in the context of the article. In this article, people are arguing with what the report says and questioning its motives. The *expression strength* is neutral, because the *Span* is simply "said". The *strength*, however, is *medium*, reflecting the negative evaluation being expressed by the experts (according to the writer).

The subjectivity at this level is reflected in the expressive subjective element "tarnish"; The choice of the word "tarnish" reflects negative evaluation of the experts toward the motivations of the authors of the report (as presented by the writer):

Expressive subjective element:

Span: tarnish
 Source: <writer,human rights experts>
 Strength: medium
 Attitude type: negative

Finally, a direct subjective frame is introduced for the nested private state referred to by “aim”:

Direct subjective:

Span: aim in sentence (Q3)
 Source: <writer,human rights experts,report>
 Strength: medium
 Expression strength: low
 Attitude type: negative
 Target: China

According to the writer, according to the experts, the authors of the report have a negative intention toward China, namely to slander them.

Finally, sentence (Q4) is also a mixture of private states and speech events:

(Q4) “The United States was slandering China again,” said Xirao-Nima, a professor of Tibetan history at the Central University for Nationalities.

The writer’s speech event is objective (the writer objectively states that someone said something and provides information about the career of the speaker):

Objective speech event:

Span: the entire sentence (Q4)
 Source: <writer>
 Implicit: true

The frame representing Xirao-Nima’s speech is subjective, reflecting his negative evaluation:

Direct subjective:

Span: said
 Source: <writer,xirao-nima>
 Strength: high
 Expression strength: neutral
 Attitude type: negative
 Target: US

The subjectivity expressed by “slandering” in this sentence is multifaceted. When we consider the level of $\langle \textit{writer}, \textit{Xirao-Nima} \rangle$, the word “slanders” is a negative evaluation of the truthfulness of what the United States said. When we consider the level of $\langle \textit{writer}, \textit{Xirao-Nima}, \textit{United States} \rangle$, the word “slanders” communicates that, according to the writer, according to Xirao-Nima, the United States said something negative about China. Thus, two frames are created for the same text span:

Expressive subjective element:

Span: slandering
 Source: <writer,xirao-nima>
 Strength: high
 Attitude type: negative

Direct subjective:

Span: slandering
 Source: <writer, xirao-nima,US>
 Target: China
 Strength: high
 Expression strength: high

4. Observations

One might initially think that writers and speakers employ a relatively small set of linguistic expressions to describe private states. Our annotated corpus, however, indicates otherwise, and the goal of this section is to give the reader some sense of the complexity of the data. In particular, we provide here a sampling of corpus-based observations that attest to the variety and ambiguity of linguistic phenomena present in naturally occurring text.

The observations below are based on an examination of a subset of the full corpus (see Section 5), which was manually annotated according to our private state annotation scheme presented in this paper. More specifically, the observations are drawn from the subset of data that was used as training data in previously published papers (Riloff and Wiebe, 2003; Riloff et al., 2003), and consists of 66 documents, for a total of 1341 sentences.

4.1. WIDE VARIETY OF WORDS AND PARTS OF SPEECH

A striking feature of the data is the large variety of words that appear in subjective expressions.

First consider direct subjective expressions whose *expression strength* is not *neutral* and that are not *implicit*. There are 1046 such expressions (consisting of 2117 word tokens) in the data. Considering only content words, i.e., adjectives, verbs, nouns, and adverbs,⁴ and excluding a

⁴ The data was automatically tokenized and tagged for part-of-speech using the ANNIE tokenizer and tagger provided in the GATE NLP development environment (Cunningham et al., 2002).

small list of stop words (*be*, *have*, *not*, and *no*), there are 1438 word tokens. Among those, there are 638 distinct words (44%).

Considering expressive subjective elements, we also find a large variety of words. There are 1766 expressive subjective elements in the data, which contain 4684 word tokens. Considering only adjectives, verbs, nouns, and adverbs, and excluding the stop words listed above, there are 2844 word tokens. Among those, there are 1463 distinct words (51%).

Clearly, a small list of words would not suffice to cover the terms appearing in subjective expressions.

The prototypical direct subjective expressions are verbs such as *criticize* and *hope*. But there is more diversity in part-of-speech than one might think. Consider the same words as above, (i.e., adjectives, verbs, nouns and adverbs excluding the stop words *be*, *have*, *not*, and *no*), in the 1046 direct subjective expressions referred to above. While 54% of them are verbs, 6% are adverbs, 8% are adjectives, and 32% are nouns. Interestingly, 342 of the 1046 direct subjective expressions (33%) do not contain a verb other than *be* or *have*.

The prototypical expressive subjective elements are adjectives. Certainly much of the work on identifying subjective expressions in NLP has focused on learning adjectives (e.g., Hatzivassiloglou and McKeown (1997), Wiebe (2000), and Turney (2002)). Among the content words (as defined above) in expressive subjective elements, 14% are adverbs, 21% are verbs, 27% are adjectives, and 38% are nouns. Fully 1087 of the 1766 expressive subjective elements in the data (62%) do not contain adjectives.

4.2. AMBIGUITY OF INDIVIDUAL WORDS

We saw in the previous section that a small list of words will not suffice to cover subjective expressions. This section shows further that many words appear in both subjective and objective expressions.

Subjective expressions are defined in this section as expressive subjective elements whose expression strength is not *low*, and direct subjective expressions whose *expression strength* is not *neutral* or *low* and that are not *implicit*. The remainder constitute objective expressions. Note that expressions with strength *low* are included in the objective class. As discussed below in Section 6, the results of our inter-annotator agreement study suggest that expressions of strength *medium* or higher tend to be clear cases of subjective expressions; the borderline cases are most often *low*.

In this section, we consider how many words appear exclusively in subjective expressions, how many appear exclusively in objective

Table I. Word Occurrence in Subjective Expressions

Percentage in Subjective Expressions	Number of Words	Percentage of Total Words
≥ 0 and ≤ 10	1423	(58.5%)
> 10 and ≤ 20	175	(7.2%)
> 20 and ≤ 30	129	(5.3%)
> 30 and ≤ 40	154	(6.3%)
> 40 and ≤ 50	197	(8.1%)
> 50 and ≤ 60	25	(1.0%)
> 60 and ≤ 70	59	(2.4%)
> 70 and ≤ 80	42	(1.7%)
> 80 and ≤ 90	17	(0.7%)
> 90 and ≤ 100	213	(8.8%)

expressions, and how many appear in both. This gives us an idea of the degree of lexical (i.e., word-level) ambiguity with respect to subjectivity.

In the data, there are 2434 words that appear more than once (there is no reason to analyze those appearing just once, since there is no potential for them to appear in both subjective and objective expressions). For each of these words, we measure the percentage of its occurrences that appear in subjective expressions. Table I summarizes these results, showing the numbers of word types whose instances appear in subjective expressions to varying degrees. The first row, for example, represents words for which between 0 and 10% of its instances appear in subjective expressions. There are 1423 such words, 58.5% of the 2434 words being considered.

As Table I shows, a non-trivial proportion of the words, 33%, fall above the lowest decile and below the highest one, showing that many words appear in both subjective and objective expressions.

Table II shows the same analysis, but only for nouns, verbs, adjectives, and adverbs excluding the stop words (*be*, *have*, *not*, and *no*). Again, we only consider words appearing at least twice in the data. The degree of ambiguity is greater with this set: 38% of the words fall between the extreme deciles.

Although many approaches to subjectivity classification focus only on the presence of subjectivity cue words themselves, disregarding context (e.g., Hart (1984), Anderson and McMaster (1982), Hatzivassiloglou and McKeown (1997), Turney (2002), Gordon et al. (2003), Yi et al.

Table II. Content Word Occurrence in Subjective Expressions

Percentage in Subjective Expressions	Number of Words	Percentage of Total Words
≥ 0 and ≤ 10	986	(51.3%)
> 10 and ≤ 20	131	(6.9%)
> 20 and ≤ 30	112	(5.9%)
> 30 and ≤ 40	137	(7.3%)
> 40 and ≤ 50	192	(10.2%)
> 50 and ≤ 60	20	(1.1%)
> 60 and ≤ 70	58	(3.1%)
> 70 and ≤ 80	43	(2.3%)
> 80 and ≤ 90	18	(1.0%)
> 90 and ≤ 100	208	(11.0%)

(2003)). the observations in this section suggest that different usages of words, in context, need to be distinguished to understand subjectivity.

4.3. MANY SENTENCES ARE MIXTURES OF SUBJECTIVITY AND OBJECTIVITY

As we have seen in previous sections, a primary focus of our annotation scheme is identifying specific expressions of private states, rather than simply labeling entire sentences or documents as subjective or objective. In this section, we present corpus-based evidence of the need for this type of fine-grained analysis of opinion and emotion (i.e., below the level of the sentence). Specifically, we show that most sentences in the data set are mixtures of objectivity and subjectivity, and often contain subjective expressions of varying strengths.

This section does not consider specific words, as in the previous sections, but rather the private states evoked in the sentence. Thus, here we consider objective speech event frames and direct subjective frames. The expressive subjective element frames are not considered, because expressive subjective elements are always subordinated by direct subjective frames, and the *strength* ratings for direct subjective frames subsume the *strength* ratings of individual expressive subjective elements. We consider the *strength* rating rather than the *expression strength* rating, because the former is a rating of the private state being expressed, while the latter is a rating of the specific speech event or private state phrase being used.

Out of the 1341 sentences in the corpus subset under study, 556 (41.5%) contain no subjectivity at all or are mixtures of objectivity and direct subjective frames of strength only *low*. Practically speaking, we may consider these to be the objective sentences.

Fully 594 (44% over the total set of sentences) of the sentences are mixtures of two or more strength ratings, or are mixtures of objective and subjective frames. Of these, 210 are mixtures of three or more strength ratings, or are mixtures of objective frames and two or more strength ratings.

4.4. POLARITY AND STRENGTH

Recall that direct subjective frames include an attribute *attitude type* that represents the polarity of the private state. The possible values are *positive*, *negative*, *both*, and *neither*.⁵

One striking observation of the annotated data is that a significant number of the direct subjective frames have the *attitude type* value *neither*. The annotators were told to indicate *positive*, *negative*, or *both* only if they were comfortable with these values; otherwise, the value should be *neither*. Out of the 1689 direct subjective frames in the data, 69% were not assigned one of those polarity values. The large proportion of *neither* ratings replicates previous findings in a study involving different data and annotators (Wiebe et al., 2001a). It suggests that simple polarity is not a sufficient notion of attitude type,⁶ a main motivation for our new work expanding this attribute to include additional distinctions (see Section 8).

Of the 521 frames with non-*neither attitude type* values, 73% are *negative*, 26% are *positive*, and 1% are *both*. Thus, we see that the majority of polarity values that the annotators felt comfortable marking are *negative* values. Interestingly, *negative* ratings are positively correlated with higher *strength* ratings: stronger expressions of opinions and emotions tend to be more negative in this corpus. Specifically, 4.6% of the low-strength direct subjective frames are negative, 20% of the medium-strength direct subjective frames are negative, and 46% of the high or extreme strength direct subjective frames are negative. Positive polarity is middle-of-the-road: 67% of the *positive* frames are medium strength, while 15.8% are low-strength and 17.3% are high or extreme strength.

⁵ In the underlying representation, the *neither* value is not explicit. Instead, it corresponds to the lack of a polarity value for the private state represented by the frame.

⁶ Which is not surprising, given that many richer typologies of emotions and attitudes have been proposed in various fields; see Section 7.

In addition, the stronger the expression, the clearer the polarity. Fully 91% of the low-strength direct subjective frames have *attitude type neither* or *both*, while 69% of the medium-strength and only 49% of the high- or extreme-strength direct subjective frames have one of these values.

These observations lead us to believe that the strength of subjective expressions will be informative for recognizing polarity, and vice versa.

5. Data

To date, 10,657 sentences in 535 documents have been annotated according to the annotation scheme presented in this paper. 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. They date from June 2001 to May 2002.

The corpus was collected and annotated as part of the summer 2002 NRRC Workshop on Multi-Perspective Question Answering (MPQA) (Wiebe et al., 2003) sponsored by ARDA. The original documents and their annotations are available at

<http://nrcc.mitre.org/NRRC/publications.htm>.

Note that the terminology in this paper differs somewhat from terminology used in the current release of the corpus, though the two versions are equivalent. Later releases of the corpus will be updated to include the new terminology. In the meantime, to help readers map between the two, the appendix of this paper shows the same text annotated using both versions and describes the relationships between them.

6. Annotator Training and Intercoder Agreement Results

In this section, we describe the training process for annotators and the results of an intercoder agreement study.

6.1. CONCEPTUAL ANNOTATION INSTRUCTIONS

Annotators begin their training by reading a coding manual that presents the annotation scheme and examples of its application (Wiebe, 2002). Below is the introduction to the manual:

Picture an information analyst searching for opinions in the world press about a particular event. Our research goal is to help him or her find what they are looking for by automatically finding text segments expressing opinions, and organizing them in a useful way.

In order to develop a computer system to do this, we need people to annotate (mark up) texts with relevant properties, such as whether the language used is opinionated and whether someone expresses a negative attitude toward someone else.

Below are descriptions of the properties we want you to annotate. We will not give you formal criteria for identifying them. We don't know formal criteria for identifying them! We want you to use your human knowledge and intuition to identify the information. Our system will then look at your answers and try to figure out how it can make the same kinds of judgments itself.

This document presents the ideas behind the annotations. A separate document will explain exactly what to annotate and how. [*Details about accessing this document deleted.*]

When you annotate, please try to be as consistent as you can be. In addition, it is essential that you interpret sentences and words with respect to the context in which they appear. Don't take them out of context and think about what they *could* mean; judge them as they are being used in that particular sentence and document.

Three themes from this introduction are echoed throughout the 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. As stated in the quote above, 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.
3. The annotators should be as consistent as they can be with respect to their own annotations and the sample annotations given to them for training.

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

6.2. TRAINING

After reading the conceptual annotation instructions, annotator training proceeds in two stages. First, the annotator focuses on learning the annotation scheme. Then, the annotator learns how to perform the annotations in the annotation tool (<http://www.cs.pitt.edu/mpqa/opinion-annotations/gate-instructions>), which is implemented in GATE (Cunningham et al., 2002).

In the first stage of training, the annotator practices applying the annotation scheme to four to six training documents, using pencil and paper to mark the private state frames and objective speech frames and their attributes. The training documents are not trivial. Instead, they are news articles from the world press, drawn from the same corpus of documents that the annotator will be annotating. When the annotation scheme was first being developed, these documents were studied and discussed in detail, until consensus annotations were agreed upon that could be used as a gold standard. After annotating each training document, the annotator compares his or her annotations to the gold standard for the document. During this time, the annotator is encouraged to ask questions, to discuss where his or her tags disagree with the gold standard, and to reread any portion of the conceptual annotation scheme which may not yet be perfectly clear.

After the annotator has a firm grasp of the conceptual annotation scheme and can consistently apply the scheme on paper, the annotator learns to apply the scheme using the annotation tool. First, the annotator reads specific instructions and works through a tutorial on performing the annotations using GATE. The annotator then practices by annotating two or three new documents using the annotation tool.

The three annotators who participated in the agreement study were all trained as described above. One annotator was an undergraduate accounting major, one was a graduate student in computer science with previous annotation experience, and one was an archivist with a degree in Library Science. None of the annotators are authors of this paper. For an annotator with no prior annotation experience or exposure to the concepts in the annotation scheme, the basic training takes approximately 40 hours. At the time of the agreement study, each annotator had been annotating part-time (8–12 hours per week) for 3–6 months.

6.3. AGREEMENT STUDY

To measure agreement on various aspects of the annotation scheme, the three annotators (A, M, and S) independently annotated 13 documents with a total of 210 sentences. The articles are from a variety of topics

and were selected so that 1/3 of the sentences are from news articles reporting on objective topics, 1/3 of the sentences are from news articles reporting on opinionated topics (“hot-topic” articles), and 1/3 of the sentences are from editorials.⁷

In the instructions to the annotators, we asked them to rate the annotation difficulty of each article on a scale from 1 to 3, with 1 being the easiest and 3 being the most difficult. The annotators were not told which articles were about objective topics or which articles were editorials, only that they were being given a variety of different articles to annotate.

We hypothesized that the editorials would be the hardest to annotate and that the articles about objective topics would be the easiest. The ratings that the annotators assigned to the articles support this hypothesis. The annotators rated an average of 44% of the articles in the study as easy (rating 1) and 26% as difficult (rating 3). But, they rated an average of 73% of the objective-topic articles as easy, and 89% of the editorials as difficult.

It makes intuitive sense that “hot-topic” articles would be more difficult to annotate than articles about objective topics and that editorials would be more difficult still. Editorials and “hot-topic” articles contain many more expressions of private states, requiring an annotator to make more judgments than they would have to for articles about objective topics.

In the subsections that follow, we describe interrater agreement for various aspects of the annotation scheme.

6.3.1. *Measuring Agreement for Text Spans*

The first step in measuring agreement is to verify that annotators do indeed agree on which expressions should be marked. To illustrate this agreement problem, consider the words and phrases identified by annotators A and M in example (26). Text spans for direct subjective frames are in italics; text spans for expressive subjective elements are in bold.

(26)

A: We *applauded* this move because it was **not only just**, but it made us *begin to feel* that we, as Arabs, were an **integral** part of Israeli society.

M: We *applauded* this move **because** it was **not only just**, but it made us *begin to feel* that we, as Arabs, were an **integral part** of Israeli society.

⁷ The results presented in this section were first reported in the 2003 SIGdial workshop (Wilson and Wiebe, 2003).

In this sentence, the two annotators mostly agree on which expressions to annotate. Both annotators agree that “applauded” and “begin to feel” express private states and that “not only just” is an expressive subjective element. However, in addition to these text spans, annotator M also marked the words “because” and “but” as expressive subjective elements. The annotators also did not completely agree about the extent of the expressive subjective element beginning with “integral.”

The annotations from (26) illustrate two issues that need to be considered when measuring agreement for text spans. First, how should we define agreement for cases when annotators identify the same expression in the text, but differ in their marking of the expression boundaries? This occurred in (26) when A identified word “integral” and M identified the overlapping phrase “integral part.” The second question to address is which statistic is appropriate for measuring agreement between different sets of annotations, which occurs when annotators disagree w.r.t. the presence or absence of individual annotations.

Regarding the first issue, we did not attempt to define rules for boundary agreement in the annotation instructions, nor was boundary agreement stressed during training. For our purposes, we believed that it was most important that annotators identified the same general expression, and that boundary agreement was secondary. Thus, for this agreement study, we consider overlapping text spans, such as “integral” and “integral part” in (26), to be matches.

The second issue concerns the fact that, in this task, there is no guarantee that the annotators will identify the same set of expressions. In (26), the set of expressive subjective elements identified by A is {“not only just”, “integral”}. The set of expressive subjective elements identified by B is {“because”, “not only just”, “but”, “integral part”}. Thus, to measure agreement we want to consider how much intersection there is between the sets of expressions identified by the annotators. Contrast this annotation task with, for example, word sense annotation, where annotators are guaranteed to annotate exactly the same sets of objects (all instances of the words being sense tagged). Because the annotators will annotate different expressions, we use the *agr* metric rather than Kappa (κ) to measure agreement in identifying text spans.

Metric *agr* is defined as follows. Let A and B be the sets of spans annotated by annotators a and b . *agr* is a directional measure of agreement that measures what proportion of A marked by a were also marked by b . Specifically, we compute the agreement of b to a as:

$$agr(a||b) = \frac{|A \text{ matching } B|}{|A|}$$

Table III. Inter-annotator Agreement: Expressive subjective elements

a	b	$agr(a b)$	$agr(b a)$	average
A	M	0.76	0.72	
A	S	0.68	0.81	
M	S	0.59	0.74	
				0.72

The $agr(a||b)$ metric corresponds to the recall if a is the gold-standard and b the system, and to precision, if b is the gold standard and a the system.

6.3.2. Agreement for Expressive Subjective Element Text Spans

In the 210 sentences in the annotation study, the annotators A, M, and S respectively marked 311, 352 and 249 expressive subjective elements. Table III shows the pairwise agreement for these sets of annotations. For example, M agrees with 76% of the expressive subjective elements marked by A, and A agrees with 72% of the expressive subjective elements marked by M. The average agreement in Table III is the arithmetic mean of all six *agr*s.

We hypothesized that the stronger the expression of subjectivity, the more likely the annotators are to agree. To test this hypothesis, we measure agreement for the expressive subjective elements rated with a *strength* of medium or higher by at least one annotator. This excludes on average 29% of the expressive subjective elements. The average pairwise agreement rises to 0.80. When measuring agreement for the expressive subjective elements rated high or extreme, this excludes an average 65% of expressive subjective elements, and the average pairwise agreement increases to 0.88. Thus, annotators are more likely to agree when the expression of subjectivity is strong. Table IV gives a sample of expressive subjective elements marked with *strength* high or extreme by two or more annotators.

6.3.3. Agreement for Direct Subjective and Objective Speech Event Text Spans

This section measures agreement, collectively, for the text spans of objective speech event and direct subjective frames. For ease of reference, in this section we will refer to these frames collectively as *explicit*

Table IV. High and extreme strength expressive subjective elements

mother of terrorism
such a disadvantageous situation
will not be a game without risks
breeding terrorism
grown tremendously
menace
such animosity
throttling the voice
indulging in blood-shed and their lunaticism
ultimately the demon they have reared will eat up their own vitals
those digging graves for others, get engraved themselves
imperative for harmonious society
glorious
so exciting
disastrous consequences
could not have wished for a better situation
unconditionally and without delay
tainted with a significant degree of hypocrisy
in the lurch
floundering
the deeper truth
the Cold War stereotype
rare opportunity
would have been a joke

frames.⁸ For the agreement measured in this section, frame type is ignored. The next section measures agreement between annotators in distinguishing objective speech events from direct subjective frames.

As we did for expressive subjective elements above, we use the *agr* metric to measure agreement for the text spans of *explicit* frames. The three annotators, A, M, and S, respectively identified 338, 285, and 315

⁸ Note that implicit frames with source *<writer>* are excluded from this analysis. They are excluded because the text spans for the writer's implicit speech events are simply the entire sentence. The agreement for the text spans of these speech events is trivially 100%.

Table V. Inter-annotator Agreement: Explicitly-mentioned private states and speech events

a	b	$agr(a b)$	$agr(b a)$	average
A	M	0.75	0.91	
A	S	0.80	0.85	
M	S	0.86	0.75	
				0.82

explicit frames in the data. Table V shows the pairwise agreement for these sets of annotations.

The average pairwise agreement for the text spans of *explicit* frames is 0.82, which indicates that they are easier to annotate than expressive subjective elements.

6.3.4. Agreement Distinguishing between Objective Speech Event and Direct Subjective Frames

In this section, we focus on interrater agreement for judgments that reflect whether or not an opinion, emotion, or other private state is being expressed by a particular source. We measure agreement for these judgements by considering how well the annotators agree in distinguishing between objective speech event frames and direct subjective frames. We consider this distinction to be a key aspect of the annotation scheme—a higher-level judgement of subjectivity versus objectivity than what is typically made for individual expressive subjective elements.

For an example of the agreement we are measuring, consider sentence (27).

(27) “Those digging graves for others, get engraved themselves’, he [Abdullah] said while citing the example of Afghanistan.

Below are the objective speech event frames and direct subjective frames identified by annotators M and S in sentence (27)⁹.

⁹ To save space, some frame attributes have been omitted.

Annotator M	Annotator S
Objective speech event frame: Span: the entire sentence Source: <writer> Implicit: true	Objective speech event frame: Span: the entire sentence Source: <writer> Implicit: true
Direct subjective frame: Span: ‘‘said’’ Source: <writer,Abdullah> Strength: high Expression strength: neutral	Direct subjective frame: Span: ‘‘said’’ Source: <writer,Abdullah> Strength: high Expression strength: neutral
Direct subjective frame: Span: ‘‘citing’’ Source: <writer,Abdullah> Strength: low Expression strength: low	Direct subjective frame: Span: ‘‘citing’’ Source: <writer,Abdullah>

For this sentence, both annotators agree that there is an objective speech event frame for the writer and a direct subjective frame for Abdullah with the text span “said.” They disagree, however, as to whether an objective speech event or a direct subjective frame should be marked for text span “citing.” Thus, to measure agreement for distinguishing between objective speech event and direct subjective frames, we first match up the *explicit* frame annotations identified by both annotators (i.e., based on overlapping text spans), including the frames for the writer’s speech events. We then measure how well the annotators agree in their classification of that set of annotations as objective speech events or direct subjective frames.

Specifically, let $S1_{all}$ be the set of all objective speech event and direct subjective frames identified by annotator $A1$, and let $S2_{all}$ be the corresponding set of frames for annotator $A2$. Let $S1_{intersection}$ be all the frames in $S1_{all}$ such that there is a frame in $S2_{all}$ with an overlapping text span. $S2_{intersection}$ is defined in the same way. The analysis in this section involves the frames $S1_{intersection}$ and $S2_{intersection}$. For each frame in $S1_{intersection}$, there is a matching frame in $S2_{intersection}$, and the two matching frames reference the same expression in the text. For each matching pair of frames then, we are interested in determining whether the annotators agree on the type of frame — is it an objective

Table VI. Annotators A & M: Contingency table for objective speech event/direct subjective frame type agreement. n_{oo} is the number of frames the annotators agreed were objective speech events. n_{ss} is the number of frames the annotators agreed were direct subjective. n_{so} and n_{os} are their disagreements.

		<i>Tagger M</i>	
		<i>ObjectiveSpeech</i>	<i>DirectSubjective</i>
<i>Tagger A</i>	<i>ObjectiveSpeech</i>	$n_{oo} = 181$	$n_{os} = 25$
	<i>DirectSubjective</i>	$n_{so} = 12$	$n_{ss} = 252$

Table VII. Pairwise Kappa scores and overall percent agreement for objective speech event/direct subjective frame type judgments

	All Expressions		Borderline Removed		
	κ	agree	κ	agree	% removed
A & M	0.84	0.91	0.94	0.96	10
A & S	0.84	0.92	0.90	0.95	8
M & S	0.74	0.87	0.84	0.92	12

speech event or a direct subjective frame? Because the set of expressions being evaluated is the same, we use Kappa (κ) to measure agreement.

Table VI shows the contingency table for these judgments made by annotators A and M. The Kappa scores for all annotator pairs are given in Table VII. The average pairwise κ score is 0.81. Under Krippendorff's scale (Krippendorff, 1980), this allows for definite conclusions.

With many judgments that characterize natural language, one would expect that there are clear cases as well as borderline cases that are more difficult to judge. This seems to be the case with sentence (27) above. Both annotators agree that there is a strong private state being expressed by the speech event "said." But the speech event for "citing" is less clear. One annotator sees only an objective speech event. The other annotator sees a weak expression of a private state (the *strength* and *expression strength* ratings in the frame are *low*). Indeed, the agreement results provide evidence that there are borderline cases for objective versus subjective speech events. Consider the expressions referenced by the frames in $S1_{intersection}$ and $S2_{intersection}$. We consider an expression to be *borderline subjective* if (1) at least one annotator marked the expression with a direct subjective frame and (2) neither

annotator characterized its *strength* as being greater than low. For example, “citing” in sentence (27) is borderline subjective. In sentence (28) below, the expression “observed” is also borderline subjective, whereas the expression “would not like” is not. The objective speech event and direct subjective frames identified by both annotators M and S are also given below.

(28) “The US authorities would not like to have it (Mexico) as a trading partner and, at the same time, close to OPEC,” he Lasserre observed.

Annotator M	Annotator S
Objective speech event frame: Span: the entire sentence Source: <writer> Implicit: true	Objective speech event frame: Span: the entire sentence Source: <writer> Implicit: true
Direct subjective frame: Span: ‘‘observed’’ Source: <writer,Lasserre> Strength: low Expression strength: low	Direct subjective frame: Span: ‘‘observed’’ Source: <writer,Lasserre> Strength: low Expression strength: neutral
Direct subjective frame: Span: ‘‘would not like’’ Source: <writer,authorities> Strength: low Expression strength: low	Direct subjective frame: Span: ‘‘would not like’’ Source: <writer,authorities> Strength: high Expression strength: high

In Table VIII we give the contingency table for the judgments given in Table VII but with the frames for the borderline subjective expressions removed. This removes, on average, only 10% of the expressions. When these are removed, the average pairwise κ climbs to 0.89.

6.3.5. *Agreement for Sentences*

In this section, we use the annotators’ low-level frame annotations to derive sentence-level judgements, and we measure agreement for those judgements.

Measuring agreement using higher-level summary judgements is informative for two reasons. First, objective speech event and direct

Table VIII. Annotators A & M: Contingency table for objective speech event/direct subjective frame type agreement, borderline subjective frames removed.

		<i>Tagger M</i>	
		<i>ObjectiveSpeech</i>	<i>DirectSubjective</i>
<i>Tagger A</i>	<i>ObjectiveSpeech</i>	$n_{oo} = 181$	$n_{os} = 8$
	<i>DirectSubjective</i>	$n_{so} = 11$	$n_{ss} = 224$

subjective frames that were excluded from consideration in Section 6.3.4 because they were identified by only one annotator¹⁰ may now be included. Second, having sentence-level judgements enables us to compare agreement for our annotations with previously published results (Bruce and Wiebe, 1999).

The annotators' sentence-level judgements are defined in terms of their lower-level frame annotations as follows. First, we exclude the objective speech event and direct subjective frames that both annotators marked as *insignificant*. Then, for each sentence, let an annotator's judgement for that sentence be *subjective* if the annotator created one or more direct subjective frames in the sentence; otherwise, let the judgement for the sentence be *objective*.

The pairwise agreement results for these derived sentence-level annotations are given in Table IX. The average pairwise Kappa for sentence-level agreement is 0.77, 8 points higher than the sentence-level agreement reported in (Bruce and Wiebe, 1999). Our new results suggest that adding detail to the annotation task can help annotators perform more reliably.

As with objective speech event versus direct subjective frame judgements, we again test whether removing borderline cases improves agreement. We define a sentence to be *borderline subjective* if (1) at least one annotator marked at least one direct subjective frame in the sentence, and (2) neither annotator marked a direct subjective frame with a *strength* greater than low. When borderline subjective sentences are removed, on average only 11% of sentences, the average Kappa increases to 0.87.

¹⁰ Specifically, the frames that are not in the sets $S1_{intersection}$ and $S2_{intersection}$ were excluded.

Table IX. Pairwise Kappa scores and overall percent agreement for sentence-level objective/subjective judgments.

	All Sentences		Borderline Removed		
	κ	agree	κ	agree	% removed
A & M	0.75	0.89	0.87	0.95	11
A & S	0.84	0.94	0.92	0.97	8
M & S	0.72	0.88	0.83	0.93	13

7. Related Work

Many different fields have contributed to the large body of work on opinionated and emotional language, including linguistics, literary theory, psychology, philosophy, and content analysis. From these fields come a wide variety of relevant terminology: subjectivity, affect, evidentiality, stance, propositional attitudes, opaque contexts, appraisal, point of view, etc. We have adopted the term “private state” from (Quirk et al., 1985) as a general covering term for internal states, as mentioned above in Section 2. At times in general discussion, we use the words “opinions” and “emotions” (these terms cover many types of private states). From literary theory, we adopt the term “subjectivity” (Banfield, 1982) to refer to the linguistic expression of private states.

While a comprehensive survey of research in subjectivity and language is outside the scope of this paper, this section sketches the work most directly relevant to ours.

Our annotation scheme grew out of a model developed to support a project in automatically tracking point of view in narrative (Wiebe, 1994). This model was in turn based on work in literary theory and linguistics, most directly Doležel (1973), Uspensky (1973), Kuroda (1973; 1976), Chatman (1978), Cohn (1978), Fodor (1979), and Banfield (1982). Our work capturing nested levels of attribution (“nested sources”) was inspired by work on propositional attitudes and belief spaces in artificial intelligence (Wilks and Bien, 1983; Asher, 1986; Rapaport, 1986) and linguistics (Fodor, 1979; Fauconnier, 1985).

The importance of strength and type of attitude in characterizing opinions and emotions has been argued by a number of researchers in linguistics (e.g., Labov (1984) and Martin (2000)) and psychology (e.g., Osgood et al. (1957) and Heise (1965)). In psychology, there is a long tradition of using hand compiled emotion lexicons in experiments to help develop or support various models of emotion. One line of research

(e.g., Osgood et al. (1957), Heise (1965), Russell (1980), and Watson and Tellegen (1985)) uses factor analysis to determine dimensions¹¹ for characterizing emotion. Others (e.g., de Rivera (1977), Ortony et al. (1987), and Johnson-Laird and Oatley (1989)) develop taxonomies of emotions. Our goals and the goals of these works in psychology are quite different—we are not interested in building models or taxonomies of emotion. Nevertheless, there is room for cross pollination. The corpus that we have developed, with words and expressions of attitudes and emotions marked in context, may be a good resource to aid in lexicon creation. Likewise, the various typologies of attitude proposed in the literature may be informative for refining our annotation scheme in the future.

The work most similar to ours is the framework proposed by Appraisal Theory (Martin, 2000; White, 2002) for analyzing evaluation and stance in discourse. Appraisal Theory emerged from the field of systemic functional linguistics (see Halliday (1994) and Martin (1992)). The Appraisal framework is composed of the concepts (or *systems* in the terminology of systemic functional linguistics): Affect, Judgement, Appreciation, Engagement, and Amplification. Affect, Judgement, and Appreciation represent different types of positive and negative attitudes. Engagement distinguishes various types of “intersubjective positioning” such as attribution and expectation. Amplification considers the force and focus of the attitudes being expressed.

Appraisal Theory is similar to our annotation scheme in that it, too, is concerned with systematically identifying expressions of opinions and emotions in context, below the level of the sentence. However, the two schemes focus on different things. The Appraisal framework primarily distinguishes different types of private state (e.g., affect versus judgement), and the typology of attitude types it proposes is much richer than ours. Currently, we only consider polarity. On the other hand, 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). The Appraisal framework also does not include a representation for nested levels of attribution.

Besides 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 annotation schemes used in Bruce and Wiebe (1999) and Wiebe et al. (2004) are earlier, less detailed versions of the annotation scheme presented in this paper.

¹¹ Dimensions corresponding to polarity and intensity are among those that have been identified.

The annotations in Bruce and Wiebe (1999) are sentence-level annotations; the annotations in Wiebe et al. (2004) mark only the text spans of expressive subjective elements. In contrast to the detailed, expression-level annotations of our current scheme, the annotations from Yu and Hatzivassiloglou (2003) are sentence level, subjective/objective and polarity judgements.

In some research that involves subjective language, lexicons are created by compiling lists of words and judging the words as to their polarity, strength, or any number of other factors of subjective language. The goals of these works are wide ranging. Osgood et al. (1957) and Heise (1965; 2001), for example, aim to develop dimensional models of affect. In work by Hart (1984), Anderson and McMaster (1982), Biber and Finegan (1989), Subasic and Heuttner (2001), and others, the lexicons are used to automatically characterize political texts, literature, news, and other types of discourse, along various subjective lines. Recent work by Kaufer et al. (2004) is noteworthy. In a multi-year project, judges compiled lists of strings which “prime” readers with respect to various categories, including many subjectivity categories. Although the lexicons that result from these works are valuable, they fail to capture, as our annotation scheme does, how subjective language is used in context in the documents where they appear.

What is innovative about our work is that it pulls together into one linguistic annotation scheme both the concept of private states and the concept of nested sources, and applies the scheme comprehensively to a large corpus, with the goal of annotating expressions in context, below the level of the sentence.

8. Future Work

The main goal behind the annotation scheme presented in this paper is to support the development and evaluation of NLP systems that exploit knowledge of opinions in applications. We have recently begun a project to incorporate knowledge of opinions into automatic question answering systems. The goals are to automatically extract the frames of our annotation scheme from text, using the annotated data for training and testing, and then to incorporate the extracted information into a summary representation of opinions that will summarize the opinions expressed in a document or group of documents (Cardie et al., 2003). Building the summary representations will force us to address questions such as which private states are similar to each other, and which source agents in a text are presented as sharing the same opinion (Bergler,

1992). Initial results in this project are described in Wilson et al. (2004) and Breck and Cardie (2004).

The most immediate refinement we plan for our annotation scheme involves the *attitude-type* attribute. Drawing on work in linguistics, psychology, and content analysis in attitude typologies (e.g., Ortony et al. (1987), Johnson-Laird and Oatley (1989), Martin (2000), White (2002), and Kaufer et al. (2004)), we will refine the *attitude type* attribute to include subtypes such as emotion, warning, stance, uncertainty, condition, cognition, intention, and evaluation. The values will include polarity values and degrees of certainty, enabling us to distinguish among, for example, positive emotions, negative evaluations, strong certainty, and weak uncertainty. A single private state frame will potentially include more than one attitude type and value, since a single expression often implies multiple types of attitudes. There has already been a fair amount of work developing typologies of attitude and emotion types, as mentioned above. Our research goal is not to develop a new typology. Our hope is that including a richer set of attitude types in our corpus annotations will contribute to a new understanding of how attitudes of various types and in various combinations are expressed linguistically in context.

The next aspect of the annotation scheme to be elaborated will be the annotation of targets. To date, only a limited set are annotated, as described above in Section 2. We will include targets that are entities who are not agents, as well as propositional targets (e.g., “Mary is sweet” in “John believes that Mary is sweet”). Frames with more than one type of attitude may contain more than one target, reflecting, for example, that a phrase expresses a negative evaluation toward one thing and a positive evaluation toward another.

Longer term research must incorporate reader and audience expectations and knowledge into the representation. This will involve not only the presumed relationship between the writer and his or her readership, but also relationships among agents mentioned in the text. For example, a quoted speaker may have spoken to a group which is quite different from the intended audience of the article. Work in narratology (Onega and Landes, 1996) will be informative for addressing such issues.

Also in the longer term, the bigger pictures of point of view, bias, and ideology must be considered. Our annotation scheme focuses on explicit linguistic expressions of private states. As such, other manifestations of point of view, bias, and ideology, such as bias reflected in which events in a story are mentioned and which roles are assigned to the participants (e.g., *insurgent* versus *soldier*), are outside the purview of our annotation scheme. Even so, a promising area of investigation would be corpus studies of interactions among these various factors.

9. Conclusions

This paper described a detailed annotation scheme that identifies key components and properties of opinions and emotions in language. The scheme pulls together into one linguistic annotation scheme both the concept of private states and the concept of nested sources, and applies the scheme comprehensively to a large corpus, with the goal of annotating expressions in context, below the level of the sentence. Several examples illustrating the scheme were given, and a corpus annotation project was described, including the results of an inter-annotator agreement study. The annotated corpus is freely available at <http://nrrec.mit.edu/NRRC/publications.htm>. We hope this work will be useful to others working in corpus-based explorations of subjective language and that it will encourage NLP researchers to experiment with subjective language in their applications.

10. Acknowledgements

We are grateful to our annotators: Matthew Bell, Sarah Kura, and Ana Mihalega and to the other participants and supporters of the ARDA Summer Workshop on Multiple Perspective Question Answering: Eric Breck, Chris Buckley, Paul Davis, David Day, Bruce Fraser, Diane Litman, Penny Lehtola, Mark Maybury, David Pierce, John Prange, and Ellen Riloff.

References

- Anderson, C. and G. McMaster: 1982, 'Computer assisted modeling of affective tone in written documents'. *Computers and the Humanities* **16**, 1–9.
- Asher, N.: 1986, 'Belief in Discourse Representation Theory'. *Journal of Philosophical Logic* **15**, 127–189.
- Banfield, A.: 1982, *Unspeakable Sentences*. Boston: Routledge and Kegan Paul.
- Bergler, S.: 1992, 'Evidential Analysis of Reported Speech'. Ph.D. thesis, Brandeis University.
- Biber, D. and E. Finegan: 1989, 'Styles of stance in English: Lexical and grammatical marking of evidentiality and affect'. *Text* **9**(1), 93–124.
- Breck, E. and C. Cardie: 2004, 'Playing the Telephone Game: Determining the Hierarchical Structure of Perspective and Speech Expressions'. In: *Proceedings of the Twentieth International Conference on Computational Linguistics (COLING 2004)*.
- Bruce, R. and J. Wiebe: 1999, 'Recognizing Subjectivity: A Case Study of Manual Tagging'. *Natural Language Engineering* **5**(2), 187–205.

- Cardie, C., J. Wiebe, T. Wilson, and D. Litman: 2003, 'Combining Low-Level and Summary Representations of Opinions for Multi-Perspective Question Answering'. In: *Working Notes — New Directions in Question Answering (AAAI Spring Symposium Series)*.
- Castaneda, H.-N., 'On the Philosophical Foundations of the Theory of Communication: Reference'. *Midwest Studies in Philosophy* **2**.
- Chatman, S.: 1978, *Story and Discourse: Narrative Structure in Fiction and Film*. Ithaca, NY: Cornell University Press.
- Cohn, D.: 1978, *Transparent Minds: Narrative Modes for Representing Consciousness in Fiction*. Princeton, NJ: Princeton University Press.
- Cunningham, H., D. Maynard, K. Bontcheva, and V. Tablan: 2002, 'GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications'. In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*.
- Dave, K., S. Lawrence, and D. M. Pennock: 2003, 'Mining the Peanut Gallery: Opinion Extraction and Semantic Classification of Produce Reviews'. In: *Proceedings of the 12th International World Wide Web Conference (WWW2003)*. Budapest, Hungary. Web Proceedings.
- de Rivera, J.: 1977, *A Structural Theory of Emotions*. New York: International Universities Press.
- Doležel, L.: 1973, *Narrative Modes in Czech Literature*. Toronto, Canada: University of Toronto Press.
- Donnellan, K.: 1966, 'Reference and Definite Descriptions'. *Philosophical Review* **60**, 281–304.
- Fauconnier, G.: 1985, *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge, MA: MIT Press.
- Fodor, J. D.: 1979, *The Linguistic Description of Opaque Contexts*, Outstanding dissertations in linguistics 13. New York & London: Garland.
- Gordon, A., A. Kazemzadeh, A. Nair, and M. Petrova: 2003, 'Recognizing Expressions of Commonsense Psychology in English Text'. In: *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics (ACL-03)*. Sapporo, Japan, pp. 208–215.
- Halliday, M.: 1985/1994, *An Introduction to Functional Grammar*. London: Edward Arnold.
- Hart, R. P.: 1984, 'Systematic Analysis of Political Discourse: The Development of DICTION'. In: K. S. et al. (ed.): *Political Communication Yearbook: 1984*. Carbondale, Illinois: Southern Illinois University Press, pp. 97–134.
- Hatzivassiloglou, V. and K. McKeown: 1997, 'Predicting the Semantic Orientation of Adjectives'. In: *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL-97)*. Madrid, Spain, pp. 174–181.
- Heise, D.: 2001, 'Project Magellan: Collecting Cross-Cultural Affective Meanings Via The Internet'. *Electronic Journal of Sociology* **5**(3).
- Heise, D. R.: 1965, 'Semantic differential profiles for 1000 most frequent English words'. *Psychological Monographs* **79**(601).
- Johnson-Laird, P. and K. Oatley: 1989, 'The language of emotions: An analysis of a semantic field'. *Cognition and Emotion* **3**(2), 81–123.
- Kaufer, D.: 2000, *Flaming: A White Paper*. www.eudora.com.
- Kaufer, D., S. Ishizaki, B. Butler, and J. Collins: 2004, *The Power of Words: Unveiling the Speaker and Writer's Hidden Craft*. Mahwah, New Jersey: Lawrence Erlbaum.

- Krippendorff, K.: 1980, *Content Analysis: An Introduction to its Methodology*. Beverly Hills: Sage Publications.
- Kuroda, S.-Y.: 1973, 'Where Epistemology, Style and Grammar Meet: A Case Study from the Japanese'. In: P. Kiparsky and S. Anderson (eds.): *A Festschrift for Morris Halle*. New York, NY: Holt, Rinehart & Winston, pp. 377–391.
- Kuroda, S.-Y.: 1976, 'Reflections on the Foundations of Narrative Theory—From a Linguistic Point of View'. In: T. van Dijk (ed.): *Pragmatics of Language and Literature*. Amsterdam: North-Holland, pp. 107–140.
- Labov, W.: 1984, 'Intensity'. In: D. Schiffrin (ed.): *Meaning, Form, and Use in Context: Linguistic Applications*. Washington, D.C.: Georgetown University Press, pp. 43–70.
- Martin, J.: 1992, *English Text: System and Structure*. Philadelphia/Amsterdam: John Benjamins.
- Martin, J.: 2000, *Evaluation in Text: Authorial stance and the construction of discourse*, Chapt. Beyond Exchange: APPRAISAL systems in English, pp. 142–175. Oxford: Oxford University Press.
- Onega, S. and J. A. G. Landes: 1996, *Narratology: An Introduction*. Longman.
- Ortony, A., G. L. Clore, and M. A. Foss: 1987, 'The referential structure of the affective lexicon'. *Cognitive Science* **11**, 341–364.
- Osgood, C. E., G. Suci, and P. Tannenbaum: 1957, *The Measurement of Meaning*. Urbana: University of Illinois Press.
- Pang, B., L. Lee, and S. Vaithyanathan: 2002, 'Thumbs up? Sentiment Classification Using Machine Learning Techniques'. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2002)*. Philadelphia, Pennsylvania, pp. 79–86.
- Quine, W. V.: 1956/1976, *The Ways of Paradox and Other Essays, Revised and Enlarged Edition*. Cambridge, MA: Harvard University Press.
- Quirk, R., S. Greenbaum, G. Leech, and J. Svartvik: 1985, *A Comprehensive Grammar of the English Language*. New York: Longman.
- Rapaport, W.: 1986, 'Logical Foundations for Belief Representation'. *Cognitive Science* **10**, 371–422.
- Riloff, E. and J. Wiebe: 2003, 'Learning Extraction Patterns for Subjective Expressions'. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2003)*. Sapporo, Japan, pp. 105–112.
- Riloff, E., J. Wiebe, and T. Wilson: 2003, 'Learning Subjective Nouns Using Extraction Pattern Bootstrapping'. In: *Proceedings of the 7th Conference on Natural Language Learning (CoNLL-2003)*. Edmonton, Canada, pp. 25–32.
- Russell, J.: 1980, 'A circumplex model of affect'. *Journal of Personality and Social Psychology* **39**, 1161–1178.
- Speratus, E.: 1997, 'Smokey: Automatic Recognition of Hostile Messages'. In: *Proceedings of the Eighth Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-97)*. Providence, Rhode Island, pp. 1058–1065.
- Subasic, P. and A. Huettner: 2001, 'Affect Analysis of Text Using Fuzzy Semantic Typing'. *IEEE Transactions on Fuzzy Systems* **9**.
- Turney, P.: 2002, 'Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews'. In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL-2000)*. Philadelphia, Pennsylvania, pp. 417–424.
- Turney, P. and M. L. Littman: 2003, 'Measuring Praise and Criticism: Inference of Semantic Orientation from Association'. *ACM Transactions on Information Systems (TOIS)* **21**(4), 315–346.

- Uspensky, B.: 1973, *A Poetics of Composition*. Berkeley, CA: University of California Press.
- Watson, D. and A. Tellegen: 1985, 'Toward a consensual structure of mood'. *Psychological Bulletin* **98**, 219–235.
- White, P.: 2002, *The Handbook of Pragmatics*, Chapt. Appraisal: The language of attitudinal evaluation and intersubjective stance, pp. 1–27. Amsterdam/Philadelphia: John Benjamins Publishing Company.
- Wiebe, J.: 1991, 'References in Narrative Text'. *Noûs* **25**(4), 457–486.
- Wiebe, J.: 1994, 'Tracking Point of View in Narrative'. *Computational Linguistics* **20**(2), 233–287.
- Wiebe, J.: 2000, 'Learning Subjective Adjectives from Corpora'. In: *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*. Austin, Texas, pp. 735–740.
- Wiebe, J.: 2002, 'Instructions for Annotating Opinions in Newspaper Articles'. Department of computer science technical report tr-02-101, University of Pittsburgh.
- Wiebe, J., R. Bruce, M. Bell, M. Martin, and T. Wilson: 2001a, 'A Corpus Study of Evaluative and Speculative Language'. In: *Proceedings of the 2nd ACL SIGdial Workshop on Discourse and Dialogue (SIGdial-2001)*. Aalborg, Denmark, pp. 186–195.
- Wiebe, J., T. Wilson, and M. Bell: 2001b, 'Identifying Collocations for Recognizing Opinions'. In: *Proceedings of the ACL-01 Workshop on Collocation: Computational Extraction, Analysis, and Exploitation*. Toulouse, France, pp. 24–31.
- Wiebe, J., T. Wilson, R. Bruce, M. Bell, and M. Martin: 2004, 'Learning Subjective Language'. *Computational Linguistics* **30**(3).
- Wilks, Y. and J. Bien: 1983, 'Beliefs, Points of View and Multiple Environments'. *Cognitive Science* **7**, 95–119.
- Wilson, T. and J. Wiebe: 2003, 'Annotating Opinions in the World Press'. In: *Proceedings of the 4th ACL SIGdial Workshop on Discourse and Dialogue (SIGdial-03)*. Sapporo, Japan, pp. 13–22.
- Wilson, T., J. Wiebe, and R. Hwa: 2004, 'Just How Mad Are You? Finding Strong and Weak Opinion Clauses'. In: *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-2004)*.
- Yi, J., T. Nasukawa, R. Bunescu, and W. Niblack: 2003, 'Sentiment Analyzer: Extracting Sentiments about a Given Topic using Natural Language Processing Techniques'. In: *Proceedings of the 3rd IEEE International Conference on Data Mining (ICDM-2003)*. Melbourne, Florida.
- Yu, H. and V. Hatzivassiloglou: 2003, 'Towards Answering Opinion Questions: Separating Facts from Opinions and Identifying the Polarity of Opinion Sentences'. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2003)*. Sapporo, Japan, pp. 129–136.

Appendix: Differences in Terminology

As stated above in Section 5, the terminology in this paper differs somewhat from terminology used in the current release of the corpus, though the two versions are equivalent. Later releases of the corpus will be updated to include the new terminology. In the meantime, to help readers map between the two, this appendix shows a sentence annotated according to both versions of the terminology, which illustrates all of the differences between the two sets of terminology.

The sentence is the following one:

The Foreign Ministry said Thursday that it was “surprised, to put it mildly” by the U.S. State Department’s criticism of Russia’s human rights record and objected in particular to the “odious” section on Chechnya.

Eleven annotations (delimited below by lines) are created for this sentence, representing the subjective expressions and agents in the sentence. For annotations that differ in the old and new versions, both versions are given. Comments are included in bold face.

Same in both versions:

AGENT=The Foreign Ministry id=ministry nested-source=w,ministry

In the terminology of this paper:

OBJECTIVE SPEECH EVENT=the entire sentence implicit=true nested-source=w

In the terminology of the current release of the data:

ON= is-implicit=yes nested-source=w onlyfactive=yes

Descriptions of the differences:

- In the current release of the data, the term ON is used for both OBJECTIVE SPEECH EVENTS and for DIRECT SUBJECTIVE annotations. An OBJECTIVE SPEECH EVENT in the new terminology corresponds to an ON with attribute-value pair “onlyfactive=yes”.
- The attribute “implicit” in the new terminology is the same as attribute “is-implicit” in the current release of the data. In the new representation, the span of an implicit OBJECTIVE SPEECH EVENT (or DIRECT SUBJECTIVE annotation) is the entire sentence; in the current release of the data, the span is empty.

In the terminology of this paper:

DIRECT SUBJECTIVE=said nested-source=w,ministry expression-strength=neutral strength=medium

In the terminology of the current release of the data:

ON=said nested-source=w,ministry onlyfactive=no on-strength=neutral overall-strength=medium

Descriptions of the differences:

- A DIRECT SUBJECTIVE annotation in the new terminology corresponds to an ON with attribute-value pair “onlyfactive=no” in the current release of the data.
- The attribute “expression-strength” in the new terminology is the same as attribute “on-strength” in the current release of the data.
- The attribute “strength” in the new terminology is the same as attribute “overall-strength” in the current release of the data.

Same in both versions:

AGENT=it nested-source=w,ministry,ministry

In the terminology of this paper:

DIRECT SUBJECTIVE=was ”surprised, to put it mildly nested-source=w, ministry, ministry expression-strength=high strength=high attitude-type=negative attitude-toward=ussd

In the terminology of the current release of the data:

ON=was ”surprised, to put it mildly nested-source=w, ministry, ministry onlyfactive=no on-strength=high overall-strength=high attitude-type=negative attitude-toward=ussd

All the differences have been already noted.

In the terminology of this paper:

EXPRESSIVE SUBJECTIVE ELEMENT=to put it mildly nested-source=w,ministry strength=medium attitude-type=negative

In the terminology of the current release of the data:

EXPRESSIVE-SUBJECTIVITY=to put it mildly nested-source=w,ministry strength=medium attitude-type=negative

Difference: The term “EXPRESSIVE SUBJECTIVE ELEMENT” in the terminology of this paper is the same as “EXPRESSIVE-SUBJECTIVITY” in the current release of the data.

Same in both versions:

AGENT=the U.S. State Department id=ussd nested-source=w,ministry,ministry,ussd

In the terminology of this paper:

DIRECT SUBJECTIVE=criticism nested-source=w,ministry,ministry,ussd
 expression-strength=medium strength=medium attitude-type=negative
 attitude-toward=russia

In the terminology of the current release of the data:

ON=criticism nested-source=w,ministry,ministry,ussd onlyfactive=no
 on-strength=medium overall-strength=medium attitude-type=negative
 attitude-toward=russia

All the differences have been already noted.

Same in both versions:

AGENT=Russia id=russia

In the terminology of this paper:

DIRECT SUBJECTIVE=objected nested-source=w,ministry expression-
 strength=medium strength=high attitude-type=negative attitude-toward=ussd

In the terminology of the current release of the data:

ON=objected nested-source=w,ministry onlyfactive=no on-strength=medium
 overall-strength=high attitude-type=negative attitude-toward=ussd

All the differences have been already noted.

In the terminology of this paper:

EXPRESSIVE SUBJECTIVE ELEMENT=odious nested-source=w,ministry
 strength=high attitude-type=negative

In the terminology of the current release of the data:

EXPRESSIVE-SUBJECTIVITY=odious nested-source=w,ministry strength=high
 attitude-type=negative

All the differences have been already noted.