
An Account of Opinion Implicatures

Janyce Wiebe · Lingjia Deng

Abstract While previous sentiment analysis research has concentrated on the interpretation of explicitly stated opinions and attitudes, this work initiates the computational study of a type of opinion implicature (i.e., opinion-oriented inference) in text. This paper describes a rule-based framework for representing and analyzing opinion implicatures which we hope will contribute to deeper automatic interpretation of subjective language. In the course of understanding implicatures, the system recognizes implicit sentiments (and beliefs) toward various events and entities in the sentence, often attributed to different sources (holders) and of mixed polarities; thus, it produces a richer interpretation than is typical in opinion analysis.

Keywords Sentiment analysis · Subjectivity analysis

1 Introduction and Motivation

Opinions are ubiquitous in language: they appear in written and spoken text ranging from editorials, reviews, blogs, and informal conversations to written news reports and broadcast news. We have seen a dramatic surge in *opinion analysis* (*sentiment analysis*) research over the past decade (Liu (2012); Pang and Lee (2008)). And many advances have been made with respect to the automatic identification and characterization of opinions in text. At the document level, methods now exist to determine whether the overall sentiment in a review or otherwise opinion-oriented text is positive or negative (e.g., Das and Chen. (2001), Pang et al (2002), Dave et al (2003)) and whether a sentence is subjective or objective (e.g., Wiebe and Riloff (2005)). There has also been work on fine-grained opinion

Janyce Wiebe
Department of Computer Science
University of Pittsburgh
E-mail: wiebe@cs.pitt.edu

Lingjia Deng
Intelligent Systems Program
University of Pittsburgh
E-mail: lid29@pitt.edu

analysis: to recognize expressions of opinion at the phrase level (e.g., Breck et al (2007), Choi et al (2006)); to determine their polarity (e.g., Wilson et al (2005), Choi and Cardie (2008)); to identify the opinion holder (e.g., Choi et al (2005), Kim and Hovy (2005); Choi et al (2006)) — the person or entity that is the source of the opinion; and to identify the target or topic of the opinion (e.g., Qiu et al (2011), Yi et al (2003), Stoyanov and Cardie (2008)).

Still, research in opinion analysis has plateaued at a somewhat superficial level, providing methods that exhibit a fairly shallow understanding of subjective language as a whole. In particular, past research in NLP has mainly addressed explicit opinion expressions, ignoring implicit opinions expressed via implicatures, i.e., default inferences. Consider, for example, the following sentences:

- (1) John was glad that Mary saved Bill.
- (2) The international community seems to be tolerating the Israeli campaign of suppression against the Palestinians.

While existing opinion analysis techniques might be able to determine the surface opinions (e.g., John has a positive attitude toward Mary’s saving Bill), they would stop short of identifying the sentences’ many *opinion implicatures*, i.e., implied attitudes and opinions, such as the following:

- for (1): *John is positive toward Bill, John is positive towards Mary* (for saving Bill), *John believes that Mary is positive towards Bill* (because she saved him).
 for (2): *The writer is negative toward the Israeli campaign of suppression and toward the International Community* (for tolerating it).

Overall, the goal of this work is to make progress toward a deeper automatic interpretation of opinionated language by developing computational models for the representation and interpretation of opinion implicature in language. This report focuses on a rule-based implementation of a conceptual framework for opinion implicature, specifically implicatures that arise in the presence of explicit sentiments, and events that positively or negatively affect entities. To eliminate interference introduced by the noisy output of automatic NLP components, the system takes as input manually annotated explicit-sentiment and event information, and makes inferences based on that input information. Thus, the purpose of this work is to provide a conceptual understanding of (a type of) opinion implicature, to provide a blueprint for realizing fully automatic systems in the future.

To make the framework more accessible, we begin this report with overview sections, first motivating the framework from an NLP perspective in Section 2, then providing terminology and sketching out its processing in Section 3.

The graphical data structure representing the system’s knowledge is described in Section 4.

The knowledge representation scheme for the nodes and edges of the graph is presented in Section 5; **this section may be safely skipped** for readers not concerned with such details.

The inference mechanisms are described next, in Section 6, namely specifications of (1) the default inference rules used by the system (Subsection 6.1), (2) the mechanism for drawing inferences in belief and sentiment spaces (Subsection 6.2), and (3) the cases when inferences are blocked (Subsection 6.3).

The semantic compositions performed by the system are described in Section 7 and, at last, the actual implicature rules are given in Section 8. The control of execution is described in Section 9.

Section 10 is the heart of the report, in that it steps through examples illustrating the system's lines of reasoning. **Some readers may prefer to jump from the overview sections directly to this section**, referring to the intervening sections according to their interests. Section 11 revisits an earlier example, now that the reader is familiar with the various inference patterns, to illustrate interdependent ambiguities.

Sections 12 and 13 return to issues and possible extensions of the knowledge representation scheme and belief and sentiment space mechanisms, respectively. **These may be safely skipped.**

Section 14 discusses related work. We first consider recent work in NLP on sentiment analysis from the perspective of our framework, and then acknowledge older work in NLP and AI whose ideas we exploited to create the inference architecture. Finally, Section 15 is the conclusion. Appendix A gives the full output of the system for the examples covered in Sections 10.2 through 10.10.

Appendix A is available at
<http://www.cs.pitt.edu/~wiebe/AppendixA.txt>

2 NLP Perspective

As mentioned above, our rule-based system was developed to provide a conceptual understanding of a type of opinion implicature. Between the rule schemas and the mechanisms for inference within sentiment and belief spaces, it produces rich interpretations, as will be seen below in Section 10. However, the system is currently supplied with manual annotations of opinion and event information, and its rule-based architecture is not ideal for practical application to real-world texts. So, before diving into presenting the rule-based system, we pause and consider opinion implicature from the perspective of NLP.

Rather than the explicit application of inference rules, we believe *sentiment propagation* will be a key mechanism in practice.

Consider the following sentence:

The bill would lower health care costs, which would be a tremendous positive change across the entire health-care system.

The writer is clearly positive toward the idea of *lowering health care costs*. But how does s/he feel about the costs? If s/he is **positive** toward the idea of *lowering* them, then, presumably, she is **negative** toward the *costs* themselves (specifically, how high they are). The only explicit sentiment expression, *tremendous positive change*, is positive, yet we can infer a **negative** attitude toward the **object** of the event itself (i.e., *health care costs*).

Going further, since *the bill* is the **agent** of an event toward which the writer is positive, we may (defeasibly) infer that the writer is **positive** toward *the bill*, even though there are no explicit sentiment expressions describing it.

Now, consider *The bill would curb skyrocketing health care costs*. The writer expresses an explicit **negative** sentiment (*skyrocketing*) toward the **object** (*health care costs*) of the event. Note that *curbing* costs, like *lowering* them, is bad for them (the costs are reduced). We can reason that, because the event is bad for something

toward which the writer is negative, the writer is **positive** toward the **event**. We can reason from there, as above, that the writer is **positive** toward *the bill*, since it is the **agent** of the positive event.

These examples illustrate how explicit sentiments toward one entity may be propagated to other entities via opinion implicature rules. In Deng and Wiebe (2014), we incorporate constraints derived from implicature rules into a graph-based model, and use Loopy Belief Propagation (Pearl (1982)) to accomplish sentiment propagation in the graph. We showed that the graph-based model improves over an explicit sentiment classification system.

A fully automatic implicature system will require several computational components, each tackling its own ambiguities, such as explicit sentiment recognition, event extraction, and semantic role labeling. This raises opportunities for interdependent ambiguity resolution. The implicature rules model dependencies among ambiguities, such that the total number of joint interpretations is greatly reduced. Thus, rather than having to take a pipeline approach, an optimization framework may exploit those interdependencies to jointly resolve ambiguities. In a first study,¹ we develop local classifiers to resolve four individual ambiguities, and then use Integer Linear Programming to conduct global inference, resulting in substantial improvements for two of the ambiguities without loss in performance for the others.

The studies so far only address sentiments of the writer, and they only exploit simplified versions of four out of ten rule schemas currently incorporated into the rule-based system. The rule-based system is meant to be a “what-if” system that looks beyond the current capabilities of fully automatic systems toward deeper interpretations that would be possible with improved NLP tools; it provides an understanding that should be helpful in designing future experiments in sentiment analysis.

3 Overview

This section gives an overview, starting with the main concepts and terminology underlying this work (in Subsection 3.1), then introducing the system’s opinion inferences in Subsection 3.2.

3.1 Some Terminology

The fundamental building blocks of our opinion implicature framework are *subjectivity*, *inferred private states*, and *benefactive/malefactive* events and states.

3.1.1 Subjectivity

In our work (Wiebe et al (2005); Wiebe (1994)), *subjectivity* is defined as the expression of private states in language, where private states are mental and emotional states such as speculations, evaluations, sentiments, and beliefs (Quirk et al (1985)). We focus on three main types of subjective expressions:

¹ Currently in submission.

References to *private states*:

(1) Japan has been **eager** for a sign that Mr. Bush is **concerned** about the issue.

References to *speech or writing* events expressing private states:

(2) UCC/Disciples leaders **roundly condemned** the Iranian President's **verbal assault** on Israel.

Expressive subjective elements, which do not refer to private states, but rather evoke them through wording and description:

(3) The **ill-conceived** plan is based on **little more** than continuation of a **business-as-usual** path.

Subjective expressions have *sources* (or *holders*): the entity or entities whose private states are being expressed. For example, in (1) the source of the private state *eager* is Japan and the source of *concerned* is Mr. Bush. In (2), the source of the subjective expressions is *UCC/Disciples leaders* and in (3) it is the writer. Sources are, in a sense, nested: for example, in (1), the source of *eager* is (writer, Japan), i.e., it is according to the writer that Japan is eager; and the source of *concerned* is (writer, Japan, Mr. Bush).

Thus, in our approach, a **private state** is an attitude held by a source toward (optionally) a target. **Sentiment and belief are types of attitudes**. Subjectivity is the linguistic expression of private states. **Subjectivity is a pragmatic notion**: as the sentence is interpreted in context, a private state is attributed to a source in that context (Banfield (1982)). By **sentiment expression** or **explicit sentiment**, we mean a subjective expression where the attitude type of the expressed private state is sentiment. We use **opinion** when a general/vague term is appropriate.

There are many types of linguistic clues that contribute to recognizing subjective expressions (Wiebe (1994)). In the clearest case, some word **senses** give rise to subjectivity whenever they are used in discourse (Wiebe and Mihalcea (2006)), for example the *hindrance* meaning of *catch* (*what's the catch?*). Other clues are not as definitive. For example, researchers in NLP have recently begun to develop lexicons of connotations (Feng et al (2011, 2013)), i.e., words associated with positive and negative polarities, out of context. For example, *war* has negative connotation and *sunshine* has positive connotation. Conceptually, though, *war* (used with an objective sense referring to physical warfare) is not itself subjective. While it's likely that *war* distributes more frequently with negative subjective expressions, positive subjectivity is certainly possible, as in *Ghenghis Kan likes war*. When we refer to *subjectivity* or *subjective expressions*, we mean that, pragmatically, attitudes are attributed to sources in that context as part of the message being conveyed.

3.1.2 Inferred Private States and Opinion Implicatures

Consider the following example from the MPQA corpus (Wiebe et al (2005)):

I think people are **happy** because Chavez has fallen.

Happy clearly indicates (according to the writer) a positive sentiment of the people toward Chavez's falling. At the same time, we might also *infer* a negative sentiment of the people toward Chavez himself, since they are happy about an event harmful to him.

We address private states inferred from other private states, where the attitude type of both is sentiment. Inference is initiated by explicit sentiment subjectivity,

either toward a gbf event (as in this example), or toward the agent or object of a gbf event (examples are given in Sections 10.6, 10.8.2, and 10.8.3).

We borrow the term *implicature* from linguistics, specifically *generalized conversational implicature*. Grice (1967, 1989) introduced the notion to account for how more can be pragmatically communicated than what is strictly said – what is implicated vs. what is said (Doran et al (2012)). Generalized conversational implicatures are cancellable, or defeasible.

Analogously, we can treat **subjectivity** as **part of what is said**,² and the **private-state inferences** we address to be **part of what is implicated**. Opinion implicatures are default inferences that may not go through in context.

Though much previous work on sentiment analysis is relevant to our work, there is almost no previous work in NLP that focuses on opinion implicature. The closest is research on plot units and affect interpretation (Lehnert (1981); Goyal et al (2010)) and related work on inferring goals in the interpretation of narratives (Schank and Abelson (1977); Wilensky (1978)). Recent research in linguistics, however, investigates one of the opinion implicature rules that we propose (see RS1 in 3.2). In particular, Reschke and Anand (2011) carry out a corpus study of the application of the inference rule via sentences that match a set of fixed linguistic templates and find that, in general, the predicted inferences hold in context. Their results bode well for our general approach.

3.1.3 Benefactive/Malefactive Events and States

This work addresses sentiments toward, in general, states and events which positively or negatively affect entities. Various lexical items and semantic roles evoke such situations. We focus on one clear case that occurs frequently in opinion sentences: $\langle agent, event, object \rangle$ triples, where *event* negatively or positively affects the *object*. Our terms for such events are *benefactive* and *malefactive*, or, for ease of writing, *goodFor* and *badFor* (hereafter *gbf*). Focusing on this clear case will make developing a fully automatic, end-to-end system more feasible. As other cases become clear, they may be incorporated into the framework in the future.

A GOODFOR event is an event that positively affects an entity (similarly, for BADFOR events). Note that gbf objects are not equivalent to benefactive/malefactive semantic roles. For example, *She baked a cake for me*: *a cake* is the object of GOODFOR event *baked* (creating something is good for it (Anand and Reschke (2010))), while *me* is the filler of its benefactive semantic role (Zúñiga and Kittilä (2010)).

A *reverser* is an expression that that reverses the polarity of a gbf event (from BADFOR to GOODFOR, or vice versa).

We have annotated a corpus with gbf information and the speaker’s sentiments toward the agents and objects of gbf events (Deng et al (2013)).³

² While the focus in the literature on what is said is semantics, Grice and people later working on the topic acknowledge that what is said must include pragmatics such as co-reference and indexical resolution (Doran et al (2012)), and subjectivity arises from deixis (Bruder and Wiebe (1995); Stein and Wright (1995)). However, as long as what is said is conceived of as only truth evaluable propositions, then it is not exactly the notion for our setting.

³ Available at <http://mpqa.cs.pitt.edu>

3.2 Opinion Inference

In this section, we introduce opinion inferences by stepping through examples of inference. The system includes default inference rules which apply if there is no evidence to the contrary. The system requires as input explicit sentiment and gfbf information (plus any evidence that is contrary to the inferences).

Consider this example:

EX(1) Why would [President Obama] **support** [health care reform]? Because [reform] could **lower** [skyrocketing health care costs], and **prohibit** [private insurance companies] from **overcharging** [patients].

Suppose an explicit-sentiment analysis system recognizes only one explicit sentiment expression, (*skyrocketing*). There are several gfbf events - *lower*, *overcharge*, and *support*.

The input to the system for EX(1) is the following. The first line, for example, represents a gfbf event, E_1 , whose agent is *reform* and whose object is *costs*. The gfbf term is *lower*, and the polarity of the gfbf event is BADFOR. The last line represents the fact that the writer's sentiment toward the costs is negative.

E_1 : $\langle \text{reform, lower:BADFOR, costs} \rangle$
 E_2 : $\langle \text{reform, prohibit:REVERSE, } E_3 \rangle$
 E_3 : $\langle \text{companies, overcharge:BADFOR, patients} \rangle$
 E_4 : $\langle \text{Obama, support:GOODFOR, reform} \rangle$
 S_1 : $\text{sent}(\text{writer, costs}) = \text{neg}$

There are 10 rule schemas implemented in the rule-based system. Since this is an overview, we give four somewhat simplified schemas here.

In the following, $\text{sent}(X, \alpha) = \beta$ means that X's sentiment toward α is β , where α is a GOODFOR event, a BADFOR event, or the agent or object of a gfbf event, and β is either *positive* or *negative* (*pos* or *neg*, for short). $P \rightarrow Q$ means to infer Q from P.

RS1: $\text{sent}(S, \text{gfbf event}) \rightarrow \text{sent}(S, \text{object})$

- 1.1 $\text{sent}(S, \text{GOODFOR}) = \text{pos} \rightarrow \text{sent}(S, \text{object}) = \text{pos}$
- 1.2 $\text{sent}(S, \text{GOODFOR}) = \text{neg} \rightarrow \text{sent}(S, \text{object}) = \text{neg}$
- 1.3 $\text{sent}(S, \text{BADFOR}) = \text{pos} \rightarrow \text{sent}(S, \text{object}) = \text{neg}$
- 1.4 $\text{sent}(S, \text{BADFOR}) = \text{neg} \rightarrow \text{sent}(S, \text{object}) = \text{pos}$

RS2: $\text{sent}(S, \text{object}) \rightarrow \text{sent}(S, \text{gfbf event})$

- 2.1 $\text{sent}(S, \text{object}) = \text{pos} \rightarrow \text{sent}(S, \text{GOODFOR}) = \text{pos}$
- 2.2 $\text{sent}(S, \text{object}) = \text{neg} \rightarrow \text{sent}(S, \text{GOODFOR}) = \text{neg}$
- 2.3 $\text{sent}(S, \text{object}) = \text{pos} \rightarrow \text{sent}(S, \text{BADFOR}) = \text{neg}$
- 2.4 $\text{sent}(S, \text{object}) = \text{neg} \rightarrow \text{sent}(S, \text{BADFOR}) = \text{pos}$

RS3: $\text{sent}(S, \text{gfbf event}) \rightarrow \text{sent}(S, \text{agent})$

- 3.1 $\text{sent}(S, \text{GOODFOR}) = \text{pos} \rightarrow \text{sent}(S, \text{agent}) = \text{pos}$
- 3.2 $\text{sent}(S, \text{GOODFOR}) = \text{neg} \rightarrow \text{sent}(S, \text{agent}) = \text{neg}$

3.3 $\text{sent}(S, \text{BADFOR}) = \text{pos} \rightarrow \text{sent}(S, \text{agent}) = \text{pos}$

3.4 $\text{sent}(S, \text{BADFOR}) = \text{neg} \rightarrow \text{sent}(S, \text{agent}) = \text{neg}$

RS4: $\text{sent}(S, \text{agent}) \rightarrow \text{sent}(S, \text{gfbf event})$

4.1 $\text{sent}(S, \text{agent}) = \text{pos} \rightarrow \text{sent}(S, \text{GOODFOR}) = \text{pos}$

4.2 $\text{sent}(S, \text{agent}) = \text{neg} \rightarrow \text{sent}(S, \text{GOODFOR}) = \text{neg}$

4.3 $\text{sent}(S, \text{agent}) = \text{pos} \rightarrow \text{sent}(S, \text{BADFOR}) = \text{pos}$

4.4 $\text{sent}(S, \text{agent}) = \text{neg} \rightarrow \text{sent}(S, \text{BADFOR}) = \text{neg}$

Applying the rule schemas to EX(1), we have the following inferences.

In E_1 , from the negative sentiment expressed by *skyrocketing* (the writer is negative toward the costs because they are too high), and the fact that *costs* is the object of a BADFOR event (*lower*), Rule2.4 infers a **positive attitude toward** E_1 .

Now, Rule3.3 applies. We infer the writer is **positive toward the reform**, since it is the agent of E_1 , toward which the writer is positive.

E_2 is an event that reverses the polarity of its object, E_3 . E_3 is the event of companies overcharging patients. And, while companies overcharging patients is BADFOR them, reform **preventing** companies from doing so is GOODFOR patients. Thus, compositionally, we have a new event:

E_{3R} : $\langle \text{reform}, \text{GOODFOR}, \text{patients} \rangle$

Above, we inferred that the writer is positive toward *reform*, the agent of E_{3R} . By Rule 4.1, the writer is **positive toward** E_{3R} . Since the writer is positive toward E_{3R} , and E_{3R} is GOODFOR patients, Rule 1.1 infers that the writer is **positive toward patients**.

Turning to E_4 , *support health care reform* is GOODFOR reform. We already inferred the writer is positive toward *reform*. Rule 2.1 infers that the writer is **positive toward** E_4 . Rule 3.1 then infers that the writer is **positive toward the agent of** E_4 , *Obama*.

In summary, RS1-RS4 infer that the writer is positive toward E_1 , health care reform, E_{3R} , patients, E_4 , and Obama. Thus, from a single explicit **negative** sentiment, the system infers several **positive** sentiments.

We now apply RS1-RS4 to an example from a political blog on the site red-state.com.

EX(2) It is no surprise then that [MoveOn] would **attack** [Senator McCain].

In the context of the blog containing it, this sentences expresses negative subjectivity toward the event. (The linguistic clues, which are ambiguous, are fronting, *surprise*, and *then* (Wiebe (1994)). We return to this example in Section 11.)

Following are the inputs for EX (2).⁴

E_5 : (MoveOn,attack:BADFOR,McCain)
 S_2 : sent(writer, E_5) = neg

Among RS1-RS4, two rules apply: Rule 1.4 infers that the writer is positive toward McCain, and Rule 3.4 infers that the writer is negative toward MoveOn.

So far in this section, the source of all the private states (both explicit and inferred) has been the writer. We now consider private states whose sources are entities mentioned in the sentence.

As will be seen below, in addition to the inferences laid out above, the system also **ascribes attitudes to the agents** of the gfbf events. For Ex(2), through a sequence of inferences, the system infers that the writer believes that (1) MoveOn is negative toward McCain, (2) MoveOn intentionally performed the action that is BADFOR McCain, and (3) MoveOn wanted to (i.e., is positive toward) perform the action. This is accomplished by the fuller set of rules used by the system (as mentioned above, RS1-RS4 are simplifications). Importantly, as will be seen in Section 6.3, the rules used to infer that the writer believes (1)-(3) are the **same rules** involved when inferences toward agents are **blocked**. The inferences are blocked by evidence that breaks the inference chains leading to (1)-(3). Thus, we do not need separate rules for **richer inference**, on the one hand, and for **defeating implicatures**, on the other; defeated inference occurs when expected inference is blocked (Levinson (1983); Schank and Abelson (1977)).

Further, if a rule matches a sentiment or event that is the target of a private state, the nesting structure is preserved when generating the conclusions. We say that inference is carried out in *private state spaces*. In addition to drawing **conclusions** within private state spaces, the system **defeasibly infers that the rule premises and assumptions are in those spaces as well**. private-state spaces are described in Section 6.2.

Note that rules are applied repeatedly until no new conclusions can be drawn. Thus, even though there are a total of only ten rule schemas, the number of inferences may be quite high.

Consider Ex(3), in which the writer expresses a sentiment toward a sentiment.

Ex(3) However, it appears as if the international community (IC) is tolerating the [Israeli] **campaign of suppression against** [the Palestinians].

The input annotations are the following:

E_6 : (Israeli,suppression:BADFOR, palestinians)
 S_3 : sent(IC, E_6) = positive

⁴ Note that the word sense of *attack* in EX(2) is subjective (Wiebe and Mihalcea (2006)), so there is also negative subjectivity of MoveOn toward McCain in the sentence, nested in the subjectivity of the writer. To keep this example from being too complex, we did not include it in the input to the system. When we add it to the input and re-run the system, the new inferences are consistent with the previous ones, and the additional sentiment reinforces the overall interpretation.

S_4 : $\text{sent}(\text{writer}, S_3) = \text{negative}$

The IC is positive toward the event in the sense that they tolerate it. *However* and *appears as if* are clues that the writer is negative toward IC's positive sentiment.

The following are the sentiments inferred **just toward the entities** in the sentence; note that many of the sentiments are nested in private-state spaces:

writer is positive toward the Palestinians
writer is negative toward Israel
writer is negative toward international community
writer believes that Israel is negative toward the Palestinians
writer believes that international community is negative toward the Palestinians
writer believes that international community is positive toward Israel
writer believes that international community believes that Israel is negative toward the Palestinians

This section presented an overview. Four rule schemas were presented, which show that inference may proceed from sentiment toward events to sentiment toward entities (RS1 and RS3), or from sentiment toward entities to sentiment toward events (RS2 and RS4). We introduced the idea that richer inference provides the basis for defeasibility, and the notion that inference is carried out in private-state spaces. In total, the rule-based system has 10 rule schemas; all of the inferences are made via (repeated) applications of rules in these schemas. If a rule matches a sentiment or event that is the target of a private state, the nesting structure is preserved. All of the inferences carried out by the system are default rules which may be blocked by evidence to the contrary.

4 Data Structure

The system builds a graphical representation of what it knows and infers about the meaning of a sentence.

Sentences are processed independently; so, the graphs for two different sentences do not share any nodes.

The system begins by building a graphical representation of the inputs for a sentence. Below is example input for a sentence. Note that the the id's in the input are only used by the system to build the initial graph. The nodes of the graph are numbered independently from the input id's.

```
"Mayor Blooming idiot urges Congress to vote for gun control."
E1 gfbf (Congress, goodFor (voting for), gun control)
S1 subjectivity (Mayor-Blooming-idiot, positive sentiment (urges), E1)
B1 privateState (writer, positive believesTrue (""), S1)
S2 subjectivity (writer, negative sentiment (blooming idiot), Mayor-Blooming-idiot)
```

The system builds these nodes (as printed by the system):

43 writer positive believesTrue
 41 Mayor-Blooming-idiot positive sentiment
 38 Congress voting for gun control
 45 writer negative sentiment
 42 Mayor-Blooming-idiot

The system's printout does not show all the structure of a node. Consider node 41. It has a *source* edge to the node representing Mayor-Blooming-idiot, and a *target* edge to node 38, which in turn has an *agent* edge to the node representing Congress and a *goodFor* edge to the node representing gun control (for convenience, there is also an edge labeled *object* from node 38 to the node representing gun control). The nodes also have attributes which record, e.g., what type of node it is (node 41 is a *privateState* and node 38 is a *gfbf*), polarity (if relevant), etc.

The graph is directed. For example, node 38 is a child of 41, which is a child of 43.

A specification for the input is that each root node must be a sentiment or believesTrue node whose source is the writer.

Inference proceeds by matching rules to the graph built so far and, when a rule successfully fires, adding nodes to the graph.

The following sections give the knowledge representation scheme for nodes (in 5), specify the form and meaning of explicit inferences rules (in 6.1), and describe a mechanism for attitude ascription (in 6.2). The various specifications give us a graphical data structure consisting of chains of believesTrue and sentiment nodes, where nodes are linked in a chain via the *target* relationship. The source of the root of each chain is the writer, and the target of the rightmost node of a chain is not a believesTrue or sentiment node. As we will see in 6.2, a node that is the target of the rightmost node on a chain is considered to be in the *private-state space* defined by that chain.

5 Knowledge Representation Scheme

Node X having attribute A with value V means $A(X,V)$.

Node P having child C, where the label of the edge from P to C is L, means $L(P,C)$.

A node represents a concept of something. Each node has a *type* attribute, specifying what type of thing the node represents.

Below, we present each type of node, its attributes, and children.

- *Type=anim*: An animate agent. A node with no attributes or children.
 - *Type=thing*: A thing. Like anim, a node no attributes or children.
 - *Type=state*: a state other than a privateState.
- Children:** experiencer, object (type: anim or thing).

Attributes: None.

- *Type=event*: an event other than a gfbf.

Children: agent, object (type: anim or thing).

Attributes: None.

- *Type=gfbf*: An event with benefactive or malefactive effect on something or someone.

Children: agent, object (type: anim or thing). If the event is goodFor (badFor) the object, then the node also has a goodFor (badFor) edge to the object.

Attributes: None.

- *Type=ideaOf*: The idea of a gfbf.

Child: ideaObject (type: gfbf).

Attributes: None.

X ideaObject Y means that X is the idea of Y.

- *Type=p(x)*:

Attribute: property (isBad, isGood, isTrue, isFalse, should, or shouldNot).

Child: x (type: anim, thing, ideaOf, agreement, privateState)

- *Type=agreement*:

Attribute: polarity (positive or negative).

Children: source (type: anim), withWhom (type: anim), target (type: p(x)).

Suppose X is an agreement node with polarity = negative; source S; withWhom W; and target p(x). This means that S disagrees with W that p(x). The disagreement is from S's perspective (i.e., according to S, W disagrees with him about p(x)).

- *Type=privateState*:

Attributes: attType (believesTrue, sentiment, intends, or believesShould); polarity (positive or negative).

Children: source (anim), target (which types are allowed depends on the attType of the node).

The source is the immediate source of the private state. Nested sources are not represented explicitly; they are defined dynamically by the nesting of privateState and agreement nodes created so far for the current sentence.

The attitude types are the following:

- *attType=believesTrue*:

Target types: privateState, agreement, p(x), or gfbf.

What S positive believesTrue T means depends on the type of T.

Note that privateStates, agreements, and p(x)'s are all propositions. It makes sense to say, for example, that the writer believes that X is positive toward something (a private state), the writer believes that X disagrees

with Y about something (an agreement), or that the writer believes that something is good (a $p(x)$). Thus, if T is one of those (i.e., not a gfbf), then S positive believesTrue T simply means that S believes that T.

Events, on the other hand, are not themselves propositions. For an event to be the object of *believes that*, it needs to be coerced into a proposition. We handle this by saying that the source has to believe *something* about the event. Thus, if T is a gfbf, then S positive believesTrue T means that there exists some p such that S believes that $p(T)$. Typically, we don't commit to what "p" is. The exception is $p = \text{"substantial"}$. S positive believesTrue substantial(T) means that S believes that T is "real," i.e., that the event happened or will happen. That's what we *mean* by the sentence "John believes that Mary killed Bill," i.e., John believes that the event happened.⁵

Finally, let's consider what S **negative** believesTrue T means. It means (A) or (B):

- (A) S does not believe that T, in the sense that T is not in S's belief space. Among the things that S believes are true, T is not one of them. S may not believe anything about T.
- (B) S does not believe that T, in the sense that S believes that not T, i.e., S believes that T is false.

Section 12 addresses this issue in some detail (that section may be safely skipped for those not interested). Suffice it to say here that we have one representation that means (A) or (B); our knowledge representation does not distinguish between the two cases. We could extend our KR to distinguish between them if needed in the future.

– *attType=sentiment:*

Target types: gfbf, $p(x)$, privateState, anim, thing, agreement, or ideaOf.

S positive sentiment toward T means that S is positive toward T (similarly for negative sentiment).

– *attType=intends*

Target type: gfbf.

S positive intends T means that S intends or intended to do T.

S negative intends T means that S does not/did not intend to do T. It does not mean that S intends to do the opposite of T. That is, "negative intends" is the absence of intention.

– *attType=believesShould:*

⁵ In general, in developing the framework, I tried to avoid requiring judgments from the future NLP system as to whether events are substantial or not. Thus, $p=\text{substantial}$ is only provided on the input when there is a reason to. As you will see below, sentiment inferences toward targets and animate agents do not hinge on whether the gfbfs are substantial. Further, by allowing p to be unspecified, a belief toward a gfbf in the input does not commit to the belief that the event is substantial/real.

Target type: gfbf.

S positive believesShould T means that S believes that the agent of T should do T.

S negative believesShould T means that S believes that the agent of T should not do T.

6 Inference

This section describes the inferences carried out by the system. The basis of the inference is, as expected, a set of rules (described next in 6.1). But, the rules alone are not sufficient, because the inferences are carried out in the context of private states. Once a rule fires, what attitudes toward the conclusions should be ascribed, and to whom? These questions are addressed in 6.2. Finally, all of the reasoning is done by default. The final subsection, 6.3, describes the types of negative evidence which may block an inference.

It should be emphasized that sentences are processed independently from one another; private states are not carried over from sentence to sentence. However, one could process an entire discourse by treating it as one long sentence.

Further, all of the reasoning mechanisms and rules assume that, within a private-state space, all the attitudes of a particular source are consistent with each other. For example, within a private-state space, there cannot be two nodes with the same sources, targets, and attitude types, but different polarities (we return to this point in Section 6.3).

6.1 Rules

The rules are default rules: a conclusion cannot be drawn if there is evidence against it.

A rule may include assumptions. For example, suppose a rule would successfully fire if S believes P for some S mentioned in the sentence. If the writer believes P, and there is no evidence to the contrary, then we'll assume that S believes it as well, if a rule "asks us to".

Thus, our rules are conceptually of the form:

$$P_1, \dots, P_j: A_1, \dots, A_k / Q_1, \dots, Q_m$$

where the Ps are preconditions, the As are assumptions, and the Qs are conclusions. For the Qs to be concluded, the Ps must already hold; there must be a basis for assuming each A; and there must be no evidence against any of the As or Qs.

Here is an example rule (as displayed by the system), which has one precondition and one conclusion.

$$S \text{ sentiment toward } A \text{ goodFor/badFor } T \implies$$

S sentiment toward the idea of A goodFor/badFor T

Assumptions are indicated using the term “Assume”, as in this rule:

S sentiment toward A goodFor/badFor T, where A is a thing &
 (Assume S positive believesTrue substantial) A goodFor/badFor T \implies
 S sentiment toward A

The first line contains a precondition, the second an assumption, and the third a conclusion.

6.2 Private-State Spaces

Inference is carried out in *private-state spaces*. The system’s knowledge is represented as chains of believesTrue and sentiment nodes, where nodes are linked in a chain via the *target* relationship. The source of the root of each chain is the writer, and the target of the rightmost node of a chain is not a believesTrue or sentiment node. A node that is the target of the rightmost node on a chain is considered to be in the *private-state space* defined by that chain.

Other than private states of the writer, all propositions and events must be the target of *some* private state. In the simplest case, the proposition or event is simply positive believedTrue by the writer.

We want to carry out inferences within private-state spaces so that, for example, from S positive believesTrue P, and $P \rightarrow Q$, the system may infer S positive believesTrue Q. However, we are working with sentiment, not only belief, and we want to allow, as appropriate, these types of inferences: from S sentiment toward P, and $P \rightarrow Q$, infer S sentiment toward Q. For example, If I’m upset my computer is infected with a virus, then I’m also upset with the consequences (e.g., that my files are corrupted).

Not all inferences are carried out via explicit rules; those such as the one just described involve ascription in private-state spaces.

The Ps, As, and Qs of the rules match nodes in the graph where their ancestors are all sentiment and believesTrue nodes. A private-state space is defined by a path where the root is a believesTrue or sentiment node whose source is the writer, and each node on the path is a believesTrue or sentiment node. Two paths define the same private-state space if, at each corresponding position, they have the same attType, polarity, and source. P is *in* a private-state space if P is the target of the rightmost node on a path defining that space.

So, suppose that (1) the preconditions, Ps, of a rule hold, (2) there is a basis to make each assumption A, and (3) there is no evidence against any of the As or the Qs. For each private-state space containing all of the Ps, nodes are built to place the As and Qs in that space as well. We call this process *space extension*.

However, if there is a negative believesTrue on the path defining a space, space extension is not carried out. I.e., this is one way an inference may be blocked. The other two ways inferences may be blocked are described in the following subsection.

6.3 Blocking of Inferences

There are two additional ways inferences may be blocked: within the context of a private-state space, and based on evidence from outside the rule-based system.

6.3.1 *Explicit Private States in Private-State Spaces*

An inference is blocked if space expansion fails, i.e., if there is no private-state space (1) whose path does not include any negative believesTrue nodes, (2) which includes all of the Ps, and (3) to which it would be valid to add all of the As and Qs. The third item still needs to be explained. It would not be valid to add a proposition P to a private-state space S if:

1. The rightmost node of S is a negative believesTrue with target P. (This is redundant with the restriction that the private-state space path cannot contain negative believesTrue nodes; we leave this as a separate condition for now.)
2. There is a node P1 already in S such that the sources and attitude types of P and P1 are the same; their targets have the same structure; but their polarities are opposite. For example, in the same space, we cannot have both a positive and a negative sentiment toward the same target.⁶

6.3.2 *Evidence Against*

In a practical setting, where the input will come from an NLP system and not manual annotations, it will generally not be reasonable to expect a system to make fine-grained private-state space distinctions.

Thus, we also have “evidence” nodes on the input, which indicate attitude type, polarity, and target, but which are not placed in private-state spaces. Rule inferences are blocked if there is evidence against an assumption or conclusion. That is, if there is evidence from the NLP system against, e.g., Q1, then the assumption or inference of Q1 is blocked from going through in any of the private-state spaces.

Importantly, the mechanisms for blocking inferences based on explicit private states in private-state spaces **cannot** be replaced by inference blocking based on evidence recognized by an NLP system. The private-state-space-blocking mechanism is needed for the internal reasoning of the system, i.e., for the proper coordination of inference rules and space extensions.

6.4 Attitude Ascription

This section provides more detail about attitude ascription.

There are two places where attitude ascription is performed: when making assumptions of the As of a rule, and during space extension.

⁶ If in the future it appears there is a need to do so, that would mean that the targets need to be more fine-grained; i.e., if we want a space to include both a positive and negative sentiment toward something, we need to split into two sentiments, one toward the positive aspect and one toward the negative aspect.

First, consider assumptions. There are two bases for making an assumption. First, we have a basis for assuming S positive believesTrue P if the writer positive believesTrue P. Second, we have a basis for assuming S has an attitude of type AT toward target T if S already has an attitude toward T which is *not* of type AT (of course, as long as the existing attitude is not a negative believesTrue). For example, we have a basis for assuming S positive believesTrue T if S has a sentiment toward T.

Remember, assumptions may later be blocked; they are actually concluded (in appropriate private-state spaces) only if the entire rule successfully fires.

Turning to space extension, there is one more piece of information to give.

As described above, if a rule successfully fires, the As and Qs are placed in all of the private-state spaces that already contain all the Ps. There is one more step: for each of those spaces which includes a sentiment to the left of the path's end, the Ps, As, and Qs are added to the spaces created by replacing sentiments with beliefs on the path. The idea is, if S has a sentiment toward T, then S must have a positive believesTrue attitude toward T. Of course, people may have sentiments toward events that never happen. But that's fine. Remember, if T is a gfbf, S positive believesTrue T means that S positive believesTrue $p(T)$, for some p . I.e., S believes *something* about the gfbf; this does not imply that S believes the event really happened (i.e., we are not committing to $p == \text{substantial}$).

The reader may notice when going through the examples in Section 10 that there are some additional beliefs the system may be warranted in inferring. For the interested reader, these are discussed in Section 13. However, that section may be safely skipped for those not interested.

7 Compositionality

The main type of compositionality performed by the system is composition of *influencers* and gfbfs. Influencers are either *reversers* or *retainers*. Here are some examples. In *John didn't kill Bill*, the negation is a reverser influencer on the gfbf event of killing Bill. The system creates a new gfbf event, *John goodFor Bill*. The new gfbf is the one that participates in the inferencing; the original one from the input is ignored by the system from this point forward.

In *John helped Mary to save Bill*, *helped* is a retainer influencer. The system builds a new gfbf: *John goodFor Bill*.

Now, because influencers are reversers or retainers, they are also, conceptually, gfbf's! Further, there may be chains of them: for example, the structure of *He stopped trying to kill Bill* is a reverser with object retainer (*trying*) which has object *badFor Bill*. In our gfbf and speaker-attitude annotation scheme (Deng et al (2013)), we opted to treat all such chains of length N as $N-1$ influencers followed by a single gfbf.

The system then starts from the gfbf, and follows the backward chain of influencers, creating new gfbf nodes as it goes. Only the final one is added to the set of nodes being retained to represent the sentence.

The question arises, how should evidence nodes be handled in such cases? It turns out that creating appropriate new evidence nodes as the new gfbf nodes are created during compositionality is straightforward.

For example:

"" Oh no! The tech staff tried to stop the virus, but they failed.

- 115 writer negative sentiment
 - 114 the tech staff (reverse)
 - 111 the tech staff stop the virus
- 118 There is evidence that the following is not substantial
 - 111 the tech staff stop the virus
- 117 There is evidence that the following is intentional:
 - 111 the tech staff stop the virus

In the input, we have a gfbf that is intentional (they tried to stop the virus), but is not substantial (they failed to do so).

The new gfbf:

influencer node:

- 115 writer negative sentiment
 - 114 the tech staff (reverse)
 - 111 the tech staff stop the virus

new gfbf node:

- 1262 writer negative sentiment
 - 1259 the tech staff goodFor the virus

new evidence node:

- 1260 There is evidence that the following is not intentional:
 - 1259 the tech staff goodFor the virus
- 1261 There is evidence that the following is substantial
 - 1259 the tech staff goodFor the virus

The new gfbf is 1259, the tech staff goodFor the virus. The new evidence nodes are 1260 and 1261; they have opposite polarities from their counterparts in the input. Gfbf event 1259 is substantial, because a good thing did happen to the virus: it wasn't stopped. But it is not intentional: the tech staff didn't try to do something goodFor the virus; in fact, they tried to do something badFor it (stop it).

The next point is not concerned with compositionality, per se, but this is a good place to point out another case of goodFor/badFor relationships the system can handle if appropriate entries are made in the lexicon:

"The people are angry that the leader deprived the children of food"

- 194 writer positive believesTrue
 - 192 the people negative sentiment
 - 193 the leader deprived the children of
 - 191 food; which is goodFor the children

The lexical entry for *deprive* notes that depriving is badFor the object (here, the children). *Deprive* also takes another semantic role, here filled by *food*; *deprive*'s entry states that the filler of that second role is goodFor the filler of the object role, whatever it is. If someone is deprived of something, then that thing is good for them (Moilanen and Pulman (2007), among others, have rules for such cases).

8 The Rules

This section gives the rules. They are described as we step through examples below; here, we address two questions.

First, why **ideaOf** A goodFor/badFor T? Because the purview of this work is making inferences about attitudes, not about events themselves. Conceptually, *ideaOf* is important for making the knowledge representation scheme consistent. An intuition is that it coerces an event into an idea, raising it into the realm of private-state spaces.

Second, why are assumptions sometimes (but not always) included in rules? The explanation is that, for a rule to fire, all of the preconditions must be at the “same level” with respect to private-state spaces, in order for the operation of adding the conclusions to private-state spaces to be well defined. Consider a rule: $P1 \ \& \ P2 \rightarrow Q$. As discussed in Section 6.4, Q is added to each private-state space that already contains both $P1$ and $P2$. If $P1$ is an attitude and $P2$ is not, then *adding Q to the spaces that contain both $P1$ and $P2$* does not make sense. In this case, the rule includes an assumption toward $P2$ that nests $P2$ in a private-state space, leveling things out.

The ten rule schemas are the following:

rule8

S positive believesTrue A goodFor/badFor T & S sentiment toward T
 \Rightarrow S sentiment toward A goodFor/badFor T

rule1

S sentiment toward A goodFor/badFor T
 \Rightarrow S sentiment toward the idea of A goodFor/badFor T

rule2

S sentiment toward the idea of A goodFor/badFor T
 \Rightarrow S sentiment toward T

rule3.1

S1 sentiment toward S2 sentiment toward Z
 \Rightarrow S1 agrees/disagrees with S2 that isGood/isBad Z & S1 sentiment toward Z

rule3.2

S1 sentiment toward S2 pos/neg believesTrue substantial Z
 \Rightarrow S1 agrees/disagrees with S2 that isTrue/isFalse Z & S1 pos/neg believesTrue substantial Z

rule3.3

S1 sentiment toward S2 pos/neg believesShould Z
 \Rightarrow S1 agrees/disagrees with S2 that should/shouldNot Z & S1 pos/neg believesShould Z

rule4

S1 agrees/disagrees with S2 that *
 \Rightarrow S1 sentiment toward S2

rule6

A goodFor/badFor T, where A is animate
 \implies A intended A goodFor/badFor T

rule7

S intended S goodFor/badFor T
 \implies S positive-sentiment toward ideaOf S goodFor/badFor T

rule9

S sentiment toward A goodFor/badFor T, where A is a thing &
 (Assume S positive believesTrue substantial) A goodFor/badFor T
 \implies S sentiment toward A

rule10

(Assume Writer positive believesTrue) A goodFor/badFor T & T in connotation lexicon
 \implies Writer sentiment toward T

The following two rules are not general inference rules, but rather capture the fact that, in English, explicit subjectivity is often expressed toward an individual rather than a proposition or event in which the individual participates. Note that they apply only to nodes that are directly from the input.

rule5source

S1 sentiment toward S2 in the input &
 (Assume S1 positive believesTrue) S2 privateState toward Z in input
 \implies S1 sentiment toward S2 privateState toward Z

rule5agent

S1 sentiment toward A in input &
 (Assume S1 sentiment toward) A goodFor/badFor T in input
 \implies S1 sentiment toward A goodFor/badFor T

9 Control

The rules are ordered. The process is iterative. On each iteration, all of the rules are applied, in the same order. Inference stops when none of the rules returns any nodes during an iteration.

10 Examples

We will present examples in a uniform way. First, the input to the system will be shown:

"Is it no surprise then that MoveOn would attack Senator McCain!?"

E1 gfbf (MoveOn, badFor (attack,attack:lexEntry), Senator McCain)

S1 subjectivity (writer, negative sentiment (surprise & then & the question), E1)

Then, the system's display of the graph built from the input:

```
4 writer negative sentiment
  1 MoveOn attack Senator McCain
```

Finally, excerpts of the system's output will be shown, with comments interspersed. The full system output may be found in Appendix A.

Below, we step through examples to motivate the rules and highlight significant reasoning chains. To make this feasible, we created (or at least simplified) examples. However, in the next subsection, we show an example from the MPQA corpus, which illustrates the case where annotators perceive subjectivity which they cannot anchor to a clear text signal; the implicature rules provide an explanation.

10.1 An example from MPQA corpus

Note that the types of inferences in this section are covered in turn in the following subsections.

Consider the following sentence from MPQA:

Ex(5) [He] is therefore planning to **trigger** [wars] here and there to **revive** [the flagging arms industry].

There are two gfbf events in this sentence: $\langle \text{He, trigger, wars} \rangle$ and $\langle \text{He, revive, arms industry} \rangle$.

The manual inputs for this sentence are:

```
E1 gfbf (He, goodFor (trigger), wars (war:lexEntry))
E2 gfbf (He, goodFor (revive), flagging arms industry)
B1 privateState (writer, positive believesTrue (""), E1)
B2 privateState (writer, positive believesTrue (""), E2)
```

The input nodes built for the inference are:

```
8 writer positive believesTrue
  4 He revive flagging arms industry
6 writer positive believesTrue
  1 He trigger wars
```

The attribute *lexEntry* in the input is a signal to retrieve information from a lexicon. In this case, the lexical entry for *war* indicates it has a negative connotation. The first inference is from connotation to sentiment:

rule10

```
(Assume Writer positive believesTrue) A goodFor/badFor T &
T's anchor is in connotation lexicon
⇒ Writer sentiment toward T
```

Assumptions:

6 writer positive believesTrue
1 He trigger wars

 \Rightarrow **Infer Node:**

17 writer negative sentiment
2 wars

From the writer's negative sentiment toward *wars*, the system infers a negative sentiment toward *trigger wars*, since triggering wars is goodFor them:

rule8

S positive believesTrue A goodFor/badFor T &
S sentiment toward T
 \Rightarrow S sentiment toward A goodFor/badFor T

Preconditions:

6 writer positive believesTrue
1 He trigger wars
17 writer negative sentiment
2 wars

 \Rightarrow **Infer Node:**

28 writer negative sentiment
1 He trigger wars

On the other hand, since the agent, *He*, is animate and there is no evidence to the contrary, the system infers that the triggering event is intentional, and that *He* is positive toward the idea of his performing the event:

rule6

A goodFor/badFor T, where A is animate
 \Rightarrow A intended A goodFor/badFor T

Preconditions:

6 writer positive believesTrue
1 He trigger wars
28 writer negative sentiment
1 He trigger wars

 \Rightarrow **Infer Node:**

38 writer negative sentiment
20 He positive intends
1 He trigger wars

rule7

S intended S goodFor/badFor T
 \Rightarrow S positive-sentiment toward ideaOf S goodFor/badFor T

Preconditions:

21 writer positive believesTrue
20 He positive intends
1 He trigger wars
38 writer negative sentiment
20 He positive intends
1 He trigger wars

 \Rightarrow **Infer Node:**

41 writer negative sentiment
25 He positive sentiment
26 ideaOf
1 He trigger wars

Let us explain the output just shown for the application of rule 6. The precondition match is node 1. Node 1 appears in two private-state spaces, *writer positive believesTrue* (node 6) and *writer negative sentiment* (node 28). Nodes 6 and 28 are both printed, so we can see which private-state spaces the precondition is in. The node matching the conclusion of the rule is node 20. The node actually inferred by the system is 38 – the conclusion has been placed in the private-state space *writer negative sentiment*. Node 20 is also in the other private-state space, *writer positive believesTrue*, but the system had already inferred *writer positive believesTrue Node 20* by this point in the inference process. In the body of this paper, we are not showing the existing nodes that are inferred, but only the newly created nodes (Appendix A shows inferences of both existing and inferred nodes).

Similarly, for the application of rule 7, node 20 matches the precondition, and 20 is in the same private-state spaces as node 1. Node 25 is the node matching the conclusion, and node 41 is the only new node created at this point; the system had already inferred *writer positive believes True Node 25*.

Continuing with inference, since the writer has a negative sentiment toward the agent's positive sentiment, the system infers that the writer disagrees with him and thus that the writer is negative toward him:

rule3.1

S1 sentiment toward (S2 sentiment toward Z)
 \implies S1 agrees/disagrees with S2 that isGood/isBad Z & S1 attitude toward Z

Preconditions:

41 writer negative sentiment
 25 He positive sentiment
 26 ideaOf
 1 He trigger wars

 \implies **Infer Node:**

50 writer disagrees with He that
 49 isGood
 26 ideaOf
 1 He trigger wars

rule4

S1 agrees/disagrees with S2 that *
 \implies S1 sentiment toward S2

Preconditions:

50 writer disagrees with He that
 49 isGood
 26 ideaOf
 1 He trigger wars

 \implies **Infer Node:**

55 writer negative sentiment
 3 He

The MPQA annotators marked the writer's negative sentiment, choosing the long spans *is therefore ... industry* and *therefore planning ... here and there* as attitude and expressive subjective element spans, respectively. They were not able to pinpoint any clear sentiment phrases. A machine learning system trained on such examples would have difficulty recognizing the sentiments. The system, relying on the negative connotation of *war* and the gfbf information in the sentence, is ultimately able to infer several sentiments, including the writer's negative sentiment toward the *trigger* event.

10.2 Inference toward the object of a gfbf

We now step through significant chains of reasoning, beginning with perhaps the most basic inference: inferring an attitude toward the object of a gfbf which is the target of a sentiment. We will use an example given above:

"Is it no surprise then that MoveOn would attack Senator McCain!?"

4 writer negative sentiment

1 MoveOn attack Senator McCain

Two rules infer the writer's sentiment toward Senator McCain, rule1 then rule2:

rule1

S sentiment toward A goodFor/badFor T
 \implies S sentiment toward the idea of A goodFor/badFor T

Preconditions:	\implies Infer Node:
4 writer negative sentiment	139 writer negative sentiment
1 MoveOn attack Senator McCain	138 ideaOf
	1 MoveOn attack Senator McCain

rule2

S sentiment toward the idea of A goodFor/badFor T
 \implies S sentiment toward T

Preconditions:	\implies Infer Node:
139 writer negative sentiment	140 writer positive sentiment
138 ideaOf	2 Senator McCain
1 MoveOn attack Senator McCain	

10.3 Inferences toward an animate agent of a gfbf

Inference of an attitude toward an animate agent of a gfbf proceeds by first inferring that the action is intentional. Continuing with our current example:

"Is it no surprise then that MoveOn would attack Senator McCain!?"

rule6

A goodFor/badFor T, where A is animate
 \implies A intended A goodFor/badFor T

Preconditions:	\implies Infer Node:
4 writer negative sentiment	142 writer negative sentiment
1 MoveOn attack Senator McCain	141 MoveOn positive intends
	1 MoveOn attack Senator McCain
	143 writer positive believesTrue
	141 MoveOn positive intends
	1 MoveOn attack Senator McCain

And then that MoveOn is positive toward their intentional event:

rule7

S intended S goodFor/badFor T
 \implies S positive-sentiment toward ideaOf S goodFor/badFor T

Preconditions:	\implies Infer Node:
142 writer negative sentiment	145 writer negative sentiment
141 MoveOn positive intends	144 MoveOn positive sentiment
1 MoveOn attack Senator McCain	138 ideaOf
143 writer positive believesTrue	1 MoveOn attack Senator McCain
141 MoveOn positive intends	146 writer positive believesTrue
1 MoveOn attack Senator McCain	144 MoveOn positive sentiment
	138 ideaOf
	1 MoveOn attack Senator McCain

A fundamental rule now applies: if the writer has a negative sentiment toward MoveOn having a positive sentiment toward X, *the writer disagrees with MoveOn* about X. Further, by definition (of disagreement), the writer has a negative sentiment toward X.

rule3.1

S1 sentiment toward S2 sentiment toward Z
 \implies S1 agrees/disagrees with S2 that isGood/isBad Z & S1 sentiment toward Z

Preconditions:	\implies Infer Node:
145 writer negative sentiment	153 writer disagrees with MoveOn that
144 MoveOn positive sentiment	152 isGood
138 ideaOf	138 ideaOf
1 MoveOn attack Senator McCain	1 MoveOn attack Senator McCain

Finally, the system infers that the writer is negative toward MoveOn, *since the writer disagrees with them*:

rule4	
S1 agrees/disagrees with S2 that *	
\implies S1 sentiment toward S2	
Preconditions:	\implies Infer Node:
153 writer disagrees with MoveOn that	158 writer negative sentiment
152 isGood	3 MoveOn
138 ideaOf	
1 MoveOn attack Senator McCain	

10.4 Attitude Ascription

The system ascribes its own reasoning to the sources in the sentence. Above, the system inferred that the writer believes that MoveOn is positive toward the idea of MoveOn attacking Senator McCain (node 146). Now, it applies rule2 *again* (note that the precondition, 144, is a child of node 146):

"Is it no surprise then that MoveOn would attack Senator McCain.!"

rule2	
S sentiment toward the idea of A goodFor/badFor T	
\implies S sentiment toward T	
Preconditions:	\implies Infer Node:
145 writer negative sentiment	150 writer negative sentiment
144 MoveOn positive sentiment	149 MoveOn negative sentiment
138 ideaOf	2 Senator McCain
1 MoveOn attack Senator McCain	151 writer positive believesTrue
146 writer positive believesTrue	149 MoveOn negative sentiment
144 MoveOn positive sentiment	2 Senator McCain
138 ideaOf	
1 MoveOn attack Senator McCain	

Recall that the system inferred above that **the writer** is **positive** toward Senator McCain; here it just inferred that the writer believes that **MoveOn** is **negative** toward him.

Moreover, the system inferred that the writer is negative that MoveOn is negative toward Senator McCain (node 150). Thus, the system infers another disagreement:

rule3.1
S1 sentiment toward S2 sentiment toward Z
\implies S1 agrees/disagrees with S2 that isGood/isBad Z & S1 sentiment toward Z

Preconditions:

150 writer negative sentiment
 149 MoveOn negative sentiment
 2 Senator McCain

⇒ Infer Node:

156 writer disagrees with MoveOn that
 155 isBad
 2 Senator McCain

We won't show it here, but rule4 fires a second time, inferring node 158 again (that the writer is negative toward MoveOn, since the writer disagrees with him).

10.5 Inferences toward an inanimate (*thing*) agent

Consider the sentence *Mother is upset that the tree fell on the boy*. Without evidence to the contrary, Mother is negative toward the tree *because* it fell on the boy; before the incident, she may have loved the tree. She doesn't "disagree" or blame the tree, since it isn't animate. Thus, in these cases, the system requires that the event actually happened, i.e., that the event is *substantial*. And, only one simple inference is made concerning the agent (the tree), namely that the writer believes that Mother is negative toward it.

Here is the input provided to the system:

```
"Mother is upset that the tree fell on the boy"
E1 gfbf (the tree:thing, badFor (fell on,fall on:lexEntry), the boy)
S1 subjectivity (mother, negative sentiment (upset), E1)
B1 privateState (writer, positive believesTrue (""), S1)
B2 privateState (writer, positive believesTrue (""), E1)
Prop1 p(B2,substantial)
```

Notice the "p=substantial" on the last line (see *believesTrue* in Section 5) and the indication that the agent is a thing (*the tree:thing*). Here are the nodes built for the input:

```
13 writer positive believesTrue substantial
  6 the tree:thing fell on the boy
11 writer positive believesTrue
  9 mother negative sentiment
    6 the tree:thing fell on the boy
```

Note that we have only told the system that *the writer* believesTrue substantial that the tree fell on the boy. The rule involved, rule9, has an assumption. For this example, the assumption is that *Mother* believesTrue substantial that the tree fell on the boy. Since the writer's belief provides a basis for the assumption, and since there is no evidence to the contrary, the rule successfully fires.

rule9

S sentiment toward A goodFor/badFor T, where A is a thing &
 (Assume S positive believesTrue substantial) A goodFor/badFor T
 ⇒ S sentiment toward A

Preconditions:

11 *writer positive believesTrue*
 9 *mother negative sentiment*
 6 *the tree fell on the boy*

Assumptions:

178 *mother positive believesTrue substantial*
 6 *the tree fell on the boy*

⇒ Infer Node:

180 *writer positive believesTrue*
 178 *mother positive believesTrue substantial*
 6 *the tree fell on the boy*
 181 *writer positive believesTrue*
 179 *mother negative sentiment*
 8 *the tree*

Other rules then fire for this example and the system infers, for example, that mother is positive toward the boy (since she is upset that something bad happened to him).

This is a good place to introduce the *by spaces* summary display of the nodes built for a sentence. The *by spaces* display consists of repetitions of the following. One or more lines in square brackets is printed; each of these is a private-state space. Then, a node is printed. This means that that node is in all of those private-state spaces. The numbers are the node numbers used elsewhere in the output.

Below is the *by spaces* display for the current sentence. The tree, for example, is only in one space: *writer positive believesTrue* (*writer +B*) *mother* has a negative sentiment (*mother -S*) (see node 181, which we inferred just above).

The end of the display shows that node 6 is in three spaces, two from the input, and the third from a rule application. In particular, node 180 was built due to the assumption in rule9.

[177 *writer +B mother +S*]

7 *the boy*

[181 *writer +B mother -S*]

8 *the tree*

[175 *writer +B mother -S*]

173 *ideaOf*

6 *the tree fell on the boy*

From Input: [11 writer +B mother -S]

From Input: [13 writer +B]

[180 *writer +B mother +B*]

6 *the tree fell on the boy*

10.6 From sentiment toward the object of a *gfbf* to sentiment toward the *gfbf* itself

It is good if something bad happens to something bad, or if something good happens to something good; it's bad if something bad happens to something good, or if something good happens to something bad. In our framework, “something bad” or “something good” is manifested as sentiment toward that thing. If that thing is the object of a *gfbf*, we can infer a sentiment toward the *gfbf* as well.

We illustrate this inference with the following example:

“Obama will finally bring skyrocketing health care costs under control”

E1 *gfbf* ⟨Obama, badFor (bring-under-control,bringUnderControl:lexEntry), skyrocketing-health-care-costs)⟩

B1 privateState ⟨writer, positive believesTrue (""), E1⟩
 S1 subjectivity ⟨writer, negative sentiment (skyrocketing), skyrocketing-health-care-costs⟩

17 writer positive believesTrue
 14 Obama bring-under-control skyrocketing-health-care-costs
 19 writer negative sentiment
 15 skyrocketing-health-care-costs

Skyrocketing is a clear negative evaluation of the object of the gbf. Bringing costs under control is badFor the costs. Thus, the system infers that the writer is **positive** toward the gbf.

rule8

S positive believesTrue A goodFor/badFor T & S sentiment toward T
 ⇒ S sentiment toward A goodFor/badFor T

Preconditions:

17 writer positive believesTrue
 14 Obama bring-under-control costs⁷
 19 writer negative sentiment
 15 skyrocketing-health-care-costs

⇒ Infer Node:

194 writer positive sentiment
 14 Obama bring-under-control costs

Now that the system has inferred a sentiment toward a gbf, inference proceeds from here as it does from the **start** for the first example above, *Is it no surprise then that MoveOn would attack Senator McCain.!?*

10.7 Complex Attitude Ascription

This section shows nodes built for a sentence with nested sentiment. First, here are the inputs and the system's display of the graph built from the input.

"MoveOn is livid that the Republicans keep opposing Obama's efforts to raise taxes on the rich"

E1 gbf ⟨Obama, goodFor (raise,raise:lexEntry), taxes on the rich⟩
 S1 subjectivity ⟨the Republicans, negative sentiment (opposing), E1⟩
 S2 subjectivity ⟨MoveOn, negative sentiment (livid), S1⟩
 B1 privateState ⟨writer, positive believesTrue (""), S2⟩

27 writer positive believesTrue
 25 MoveOn negative sentiment
 23 the Republicans negative sentiment
 20 Obama raising taxes on the rich

Now, here is the "by spaces" representation of several of the nodes built for this sentence (see Appendix A for the full listing):

[315 writer +B]

313 MoveOn disagrees with the Republicans that
 312 isBad

⁷ We use *costs* to represent *skyrocketing-health-care-costs* for short in node 14.

281 Obama positive sentiment
 21 taxes on the rich
 [370 writer +B]
 368 MoveOn disagrees with the Republicans that
 367 isBad
 22 Obama
 [321 writer +B MoveOn -S]
 [323 writer +B MoveOn +B]
 318 the Republicans disagrees with Obama that
 317 isGood
 21 taxes on the rich
 [366 writer +B]
 364 MoveOn agrees with Obama that
 317 isGood
 21 taxes on the rich
 [235 writer +B MoveOn -S the Republicans -S]
 [236 writer +B MoveOn +B the Republicans -S]
 [249 writer +B MoveOn +S]
 [271 writer +B MoveOn -S the Republicans -S Obama +S]
 [274 writer +B MoveOn +B the Republicans -S Obama +S]
 [277 writer +B MoveOn +B the Republicans +B Obama +S]
 [304 writer +B MoveOn +S Obama +S]
 [337 writer +B MoveOn +B Obama +S]
 232 ideaOf
 20 Obama raise taxes on the rich
 [256 writer +B MoveOn -S]
 24 the Republicans
 [330 writer +B MoveOn -S the Republicans -S]
 [332 writer +B MoveOn +B the Republicans -S]
 [371 writer +B MoveOn +S]
 22 Obama
 [242 writer +B MoveOn -S the Republicans -S]
 [244 writer +B MoveOn +B the Republicans -S]
 [254 writer +B MoveOn +S]
 [284 writer +B MoveOn -S the Republicans -S Obama +S]
 [287 writer +B MoveOn +B the Republicans -S Obama +S]
 [290 writer +B MoveOn +B the Republicans +B Obama +S]
 [316 writer +B MoveOn +S Obama +S]
 [346 writer +B MoveOn +B Obama +S]
 21 taxes on the rich

10.8 Rules that Fire Only on Input

10.8.1 From connotation to sentiment

Rule8 (which infers sentiment toward a gfbf from sentiment toward its object) fires only if there is **sentiment** toward the object of a gfbf which, as with all

attitudes, is attributed to a source. What if there is merely connotation in place of the sentiment, for example, *The attack is a fight against justice* or *He started a war* (*justice* has positive connotation while *war* has negative connotation). The way to handle this is to infer sentiment from connotation. As with all the rules, an explicit sentiment to the contrary would block the inference of a sentiment from a connotation. For example, the writer might be *glad* that someone started a war, if he is an arms dealer (even though *war* has negative connotation, and starting a war is goodFor it). This section illustrates the inference from connotation to sentiment, using the following sentence.

“GOP Attack on Health Care Reform Is a Fight Against Racial Justice.”

```
E1 gbf (GOP, badFor (attack), Health-Care-Reform)
B1 privateState (writer, positive believesTrue (""), E1)
E2 gbf (GOP-Attack, badFor (fight), racial-justice)
B2 privateState (writer, positive believesTrue (""), E2)
```

```
32 writer positive believesTrue
   29 GOP attack Health-Care-Reform
37 writer positive believesTrue
   34 GOP-Attack fight racial-justice
```

This sentence has two gbfbs: the GOP attacking health care reform and the GOP’s attack fighting racial justice. The lexicon has an entry for *justice* indicating that it has positive connotation. However, it does not have an entry for *health care reform* because, in our healthcare dataset (Deng et al (2013)), both negative and positive attitudes toward reform are common.

Note that a rule fires on *all* possible matches. So, in the full output (Appendix A) you will see that the system makes inferences for both gbfbs. The connotation rule, rule10, only matches the object of the second gbf, since only its text anchor has a connotation in the lexicon.

Note that rule10 may only infer sentiments of the writer – not other sources in the sentence. The reason is that it is *the writer’s* words that appear in the sentence. (At least, the current system assumes that all words in the sentence are the writer’s words. In the future, a component should be added that recognizes quotations and rule10 should be modified appropriately.)

rule10

(Assume Writer positive believesTrue) A goodFor/badFor T &
T’s anchor is in connotation lexicon
⇒ Writer sentiment toward T

Assumptions:

```
32 writer positive believesTrue
   29 GOP-Attack fight racial-justice
```

⇒ Infer Node:

```
418 writer positive sentiment
   34 racial-justice
```

Now that this sentiment (node 418) has been inferred, inference proceeds as in 10.6 for the sentence *Obama will finally bring skyrocketing health care costs under control*. Following is a subset of the *by spaces* display of the nodes built for the sentence.

```
454 writer disagrees with GOP-Attack that
```

453 isBad
 34 racial-justice
 “GOP Attack on Health Care Reform Is a Fight Against Racial Justice.”
 [418 writer +S]
 [435 writer +B GOP-Attack -S]
 [448 writer -S GOP-Attack -S]
 34 racial-justice

The inferences often bring into sharp relief the fact that the words are the writer’s words and not those of the other sources in the sentence. It is doubtful that the GOP would agree that it has a negative sentiment toward *racial justice*.

10.8.2 From sentiment toward the source to sentiment toward the private state

In this section, we return to the example we started in Section 4.

“Mayor Blooming idiot urges Congress to vote for gun control.”
 E1 gfbf (Congress, goodFor (voting for,voteFor:lexEntry), gun control)
 S1 subjectivity (Mayor-Blooming-idiot, positive sentiment (urges), E1)
 B1 privateState (writer, positive believesTrue (“”), S1)
 S2 subjectivity (writer, negative sentiment (blooming idiot), Mayor-Blooming-idiot)

 43 writer positive believesTrue
 41 Mayor-Blooming-idiot positive sentiment
 38 Congress voting for gun control
 45 writer negative sentiment
 42 Mayor-Blooming-idiot

In this sentence, the writer opts to explicitly express his negative sentiment toward Mayor Bloomberg (*Blooming idiot*), not toward Mayor Bloomberg’s attitude toward the gfbf (i.e., his positive attitude toward Congress doing something good-For gun control). Rule rule5source infers a sentiment toward an attitude from a sentiment toward the source of that attitude:

rule5source	
S1 sentiment toward S2 in the input & (Assume S1 positive believesTrue) S2 privateState toward Z in input ⇒ S1 sentiment toward S2 privateState toward Z	
Preconditions:	⇒ Infer Node:
45 writer negative sentiment	444 writer negative sentiment
42 Mayor-Blooming-idiot	41 Mayor-Blooming-idiot positive sentiment
Assumptions:	38 Congress voting for gun control
43 writer positive believesTrue	
41 Mayor-Blooming-idiot positive sentiment	
38 Congress voting for gun control	

From this point, inference proceeds as it did for the example in 10.7, *MoveOn is livid that the Republicans keep opposing Obama’s efforts to raise taxes on the rich*.

This rule (rule5source) is not a general conceptual rule such as the ones we’ve seen so far. It is a rule concerning which explicit sentiment expressions writers

choose to include in a sentence. Thus, the rule applies only to nodes built to represent the input.⁸

Note that rule5source includes an assumption. Here is an example for which the assumption comes into play.

“Republicans oppose Obama because he supports the states legalizing gay marriage.”

E1 gfbf (the states, goodFor (legalizing,legalize:lexEntry), gay marriage)

S1 subjectivity (Obama, positive sentiment (supports), E1)

B1 privateState (writer, positive believesTrue (“”), S1)

S2 subjectivity (Republicans, negative sentiment (supports), Obama)

B2 privateState (writer, positive believesTrue (“”), S2)

55 writer positive believesTrue

53 Republicans negative sentiment

50 Obama

51 writer positive believesTrue

49 Obama positive sentiment

46 the states legalizing gay marriage

The preconditions of rule5source have to be “at the same level” in terms of private-state spaces. Via the assumption, the system ascribes to the Republicans the writer’s belief (node 51) that Obama has a positive sentiment toward the gfbf (legalizing gay marriage).

rule5source

S1 sentiment toward S2 in the input &
 (Assume S1 positive believesTrue) S2 privateState toward Z in input
 \implies S1 sentiment toward S2 privateState toward Z

Preconditions:

55 *writer positive believesTrue*

53 Republicans negative sentiment

50 Obama

Assumptions:

592 Republicans positive believesTrue

49 Obama positive sentiment

46 the states legalizing gay marriage

\implies Infer Node:

594 *writer positive believesTrue*

592 Republicans positive believesTrue

49 Obama positive sentiment

46 the states legalizing gay marriage

595 *writer positive believesTrue*

593 Republicans negative sentiment

49 Obama positive sentiment

46 the states legalizing gay marriage

Following is an example with a source other than the writer for which the belief is explicitly in the sentence, so a new assumption does not need to be made.

“Muslims also hate Obama because they think he supported the Koran burning event by Pastor Terry Jones in Florida on March 2011.”

S1 subjectivity (Muslims, negative sentiment (hate), Obama)

B1 privateState (writer, positive believesTrue (“”), S1)

E1 gfbf (Pastor-Terry-Jones, badFor (burning,burn:lexEntry), Koran)

S2 subjectivity (Obama, positive sentiment (supports), E1)

⁸ As indicated in a footnote above, because the rule5 rules are not general inference rules, currently the system will only allow the rule to fire once on a given precondition, even if two assumptions are possible with the same precondition. The idea is that one fire gives us an explanation or reason for the precondition. This option can be changed by flipping a variable value from 1 to 0.

B2 privateState ⟨Muslims, positive believesTrue (""), S2⟩
 B3 privateState ⟨writer, positive believesTrue (""), B2⟩

66 writer positive believesTrue
 65 Muslims positive believesTrue
 64 Obama positive sentiment
 61 Pastor-Terry-Jones burning Koran
 59 writer positive believesTrue
 56 Muslims negative sentiment
 57 Obama

rule5source

S1 sentiment toward S2 in the input &
 (Assume S1 positive believesTrue) S2 privateState toward Z in input
 ⇒ S1 sentiment toward S2 privateState toward Z

Preconditions:

59 *writer positive believesTrue*
 56 Muslims negative sentiment
 57 Obama

⇒ Infer Node:

791 *writer positive believesTrue*
 790 Muslims negative sentiment
 64 Obama positive sentiment
 61 Pastor-Terry-Jones burning Koran

Assumptions:

66 *writer positive believesTrue*
 65 Muslims positive believesTrue
 64 Obama positive sentiment
 61 Pastor-Terry-Jones burning Koran

10.8.3 From sentiment toward the agent of a gfbf to sentiment toward the gfbf

This rule is very similar to the one just above. It infers a sentiment toward a gfbf from a sentiment toward the agent of the event. This rule also applies only to nodes built from the input, and includes the same type of assumption as rule5source does. Thus, we only show one example here.

“Muslims started hating Obama when he ordered the US Army to kill Osama bin Laden”

S1 subjectivity ⟨Muslims, negative sentiment (hate), Obama⟩
 B1 privateState ⟨writer, positive believesTrue (""), S1⟩
 E1 gfbf ⟨US Army, badFor (kill,kill:lexEntry), Osama bin Laden⟩
 I1 influencer ⟨Obama, retain (ordered,order:lexEntry), E1⟩
 B2 privateState ⟨writer, positive believesTrue (""), I1⟩

76 writer positive believesTrue
 75 Obama ⟨retain⟩
 72 US Army kill Osama bin Laden
 70 writer positive believesTrue
 67 Muslims negative sentiment
 68 Obama

influencer node:

76 *writer positive believesTrue*
 75 Obama ⟨retain⟩
 72 US Army kill Osama bin Laden

new gfbf:

976 *writer positive believesTrue*
 975 Obama badFor Osama bin Laden

rule5agent

S1 sentiment toward A in input &
 (Assume S1 sentiment toward) A goodFor/badFor T in input
 ⇒ S1 sentiment toward A goodFor/badFor T

Preconditions:

70 *writer positive believesTrue*
 67 Muslims negative sentiment
 68 Obama

Assumptions:

977 Muslims negative sentiment
 975 Obama badFor Osama bin Laden

⇒ Infer Node:

978 *writer positive believesTrue*
 977 Muslims negative sentiment
 975 Obama badFor Osama bin Laden

This completes the section on rules that fire only on input nodes. Looking to the future, we believe that the rule5* rules might be a good place to start when extending the approach to the level of the discourse. Often, there are patterns such as *The Muslims hate Obama. He did X, he supports Y, etc.*

10.9 Arguing Subjectivity

So far in this document, we've considered *sentiment subjectivity*: positive and negative emotions, evaluations, and judgments. Another important type of subjectivity is *arguing subjectivity*, where someone argues (1) that something is or isn't true, or (2) that something should or should not be done.

The first type of arguing subjectivity is given in the input as subjectivity where the attitude type of the private state is *believesTrue* and, if the target of the private state is a *gfbf*, the attribute *p* is given the value *substantial*. That is, the interpretation supported under the framework is that, if someone arguesTrue an event, then what they are arguing is that the event actually occurred (or will occur). (Down the road, I foresee the NLP system's task as recognizing arguing subjectivity based on lexical items such as "accuse"; the *p=substantial* attribute would simply be added automatically whenever arguing subjectivity is placed on the input.)

The second is subjectivity where the attitude type of the private state is *believesShould*.

To date, we haven't addressed the second type. But, the system is able to make an interesting inference with respect to the first type.

Consider the following example (we will consider the crossed-out lines below).

"Republicans roared onto the post-State-of-the-Union morning shows accusing President Obama of waging class warfare against the rich"

E1 *gfbf* (obama, badFor (waging class warfare against,wagingClassWarfare:lexEntry), the rich)

B1 *subjectivity* (republicans, positive believesTrue (accusing), E1)

B2 *privateState* (writer, positive believesTrue (""), B1)

Prop1 *p*(B1,substantial)

S1 *subjectivity* (writer, negative sentiment (roared), republicans)

~~S2 *subjectivity* (republicans, negative sentiment (accusing), obama)~~

~~B3 *privateState* (writer, positive believesTrue (""), S2)~~

82 *writer positive believesTrue*

80 *republicans positive believesTrue substantial*

77 *obama waging class warfare against the rich*

84 *writer negative sentiment*

81 republicans

Inference begins with an application of rule5source:

rule5source

S1 sentiment toward S2 in the input &
 (Assume S1 positive believesTrue) S2 privateState toward Z in input
 ⇒ S1 sentiment toward S2 privateState toward Z

Resulting in this attitude of the writer's:

1027 writer negative sentiment
 80 republicans positive believesTrue substantial
 77 obama waging class warfare against the rich

In the examples so far, the system has inferred positive and negative sentiments (toward gfbfs and their agents and objects). However, *the writer isn't negative due to sentiment; he's negative because he doesn't think it's true!*

rule3.2

S1 sentiment toward S2 positive/negative believesTrue substantial Z
 ⇒ S1 agrees/disagrees with S2 that isTrue/isFalse Z &
 S1 pos/neg believesTrue substantial Z

Preconditions:

1027 writer negative sentiment
 80 republicans positive believesTrue substantial
 77 obama waging class warfare against rich

⇒ Infer Node:

1029 writer disagrees with republicans that
 1028 isTrue
 77 obama waging class warfare against rich
 1030 writer negative believesTrue substantial
 77 obama waging class warfare against rich

Let's step through the rest of the inferences for this sentence, to see how things fit together (some comments have been added).

rule4

S1 agrees/disagrees with S2 that *
 ⇒ S1 sentiment toward S2

Preconditions:

1029 writer disagrees with republicans that
 1028 isTrue
 77 obama waging class warfare against rich

⇒ Infer Node:

Existing:
 84 writer negative sentiment
 81 republicans

The system infers that the writer is negative toward the republicans because he disagrees with them about something. This matches the negative subjectivity in the input.⁹

Continuing on, the system infers an intention and then the agent's positive attitude toward his intended action. Note that node 1030 is a negative believesTrue private state. Negative beliefs block space-extension inferences. Thus, the system does not infer that the writer believes that Obama intends the action.

⁹ The system often infers nodes that already exist. On the full printout (see Appendix A), inferences of existing nodes are included, marked by "Existing:" as above for Node 84.

rule6

A goodFor/badFor T, where A is animate
 \implies A intended A goodFor/badFor T

Preconditions:

1030 *writer negative believesTrue substantial*
 77 obama waging class warfare against rich
 82 *writer positive believesTrue*
 80 *republicans positive believesTrue substantial*
 77 obama waging class warfare against rich
 1027 *writer negative sentiment*
 80 *republicans positive believesTrue substantial*
 77 obama waging class warfare against rich

 \implies **Infer Node:**

1034 *writer positive believesTrue*
 1033 *republicans positive believesTrue*
 1032 obama positive intends
 77 obama waging class warfare ...
 1036 *writer negative sentiment*
 1033 *republicans positive believesTrue*
 1032 obama positive intends
 77 obama waging class warfare ...
 Inference blocked in space [writer -B]:
 because it contains a negative believesTrue.

rule7

S intended S goodFor/badFor T
 \implies S positive-sentiment toward ideaOf S goodFor/badFor T

Preconditions:

1034 *writer positive believesTrue*
 1033 *republicans positive believesTrue*
 1032 obama positive intends
 77 obama waging class warfare ...
 1036 *writer negative sentiment*
 1033 *republicans positive believesTrue*
 1032 obama positive intends
 77 obama waging class warfare ...

 \implies **Infer Node:**

1040 *writer positive believesTrue*
 1039 *republicans positive believesTrue*
 1037 obama positive sentiment
 1038 ideaOf
 77 obama waging class warfare ...
 1042 *writer negative sentiment*
 1039 *republicans positive believesTrue*
 1037 obama positive sentiment
 1038 ideaOf
 77 obama waging class warfare ...

The system has just inferred (within the spaces [writer +B republicans +B] and [writer -S republications +B]) that Obama has a positive sentiment toward the idea of the gfbf. Thus, the system infers he has a negative sentiment toward the object of the gfbf (the rich):

rule2

S sentiment toward the idea of A goodFor/badFor T
 \implies S sentiment toward T

Preconditions:

1040 *writer positive believesTrue*
 1039 *republicans positive believesTrue*
 1037 obama positive sentiment
 1038 ideaOf
 77 obama waging class warfare...
 1042 *writer negative sentiment*
 1039 *republicans positive believesTrue*
 1037 obama positive sentiment
 1038 ideaOf
 77 obama waging class warfare...

 \implies **Infer Node:**

1045 *writer positive believesTrue*
 1044 *republicans positive believesTrue*
 1043 obama negative sentiment
 78 the rich
 1047 *writer negative sentiment*
 1044 *republicans positive believesTrue*
 1043 obama negative sentiment
 78 the rich

That is the end of the inferences for this sentence.

For the sake of comparison, consider the sentence above, *Mayor Blooming idiot urges Congress to vote for gun control*. The *Blooming idiot* sentence has the same

structure as the one we just looked at, except it has sentiment subjectivity (*urges*) where the current sentence has arguing subjectivity (*accusing*). A segment of the *by spaces* representation will remind you of the inferences made for that sentence (note that *gun control* is the object of the gfbf):

“Mayor Blooming idiot urges Congress to vote for gun control.”

[450 writer +B Mayor-Blooming-idiot +S]

[451 writer -S Mayor-Blooming-idiot +S]

[457 writer -S]

[479 writer +B Mayor-Blooming-idiot +S Congress +S]

[481 writer -S Mayor-Blooming-idiot +S Congress +S]

[483 writer +B Mayor-Blooming-idiot +B Congress +S]

[504 writer -S Congress +S]

[522 writer +B Congress +S]

39 gun control

If you look at the full *by spaces* representation for the *Mayor Blooming idiot* sentence, you will find 10 agrees/disagrees nodes.

For our current sentence, however, we do not have sentiments nested in sentiments. Here is the segment of the *by spaces* display for the object of the gfbf (*the rich*). *The rich* is only in two spaces.

“Republicans roared onto the post-State-of-the-Union morning shows accusing President Obama of waging class warfare against the rich”

[1045 writer +B republicans +B obama -S]

[1047 writer -S republicans +B obama -S]

78 the rich

Further, there is only a single agrees/disagrees node, the one concerning whether Obama is waging class warfare:

1029 writer disagrees with republicans that

1028 isTrue

77 obama waging class warfare against the rich

Recall that two lines were crossed out in the input above. Those lines correspond to the Republicans having negative sentiment toward Obama. Here is the full set of input lines and the nodes built to represent them:

“(alt) Republicans roared onto the post-State-of-the-Union morning shows accusing President Obama of waging class warfare against the rich”

E1 gfbf (obama, badFor (waging class warfare against,wagingClassWarfare:lexEntry), the rich)

B1 subjectivity (republicans, positive believesTrue (accusing), E1)

B2 privateState (writer, positive believesTrue (""), B1)

Prop1 p(B1,substantial)

S1 subjectivity (writer, negative sentiment (roared), republicans)

S2 subjectivity (republicans, negative sentiment (accusing), obama)

B3 privateState (writer, positive believesTrue (""), S2)

94 writer positive believesTrue
 93 republicans negative sentiment
 87 obama

90 writer positive believesTrue
 88 republicans positive believesTrue substantial
 85 obama waging class warfare against the rich

92 writer negative sentiment
 89 republicans

For this version of the input, the system's inferences are superset of the inferences made for the original version. The additional inferences come from the added sentiment (node 93). Basically, rule rule5agent infers, from node 93, a republican sentiment toward the gfbf, and inference continues from there (as for several examples above).

Finally, here an example where the input attitude corresponding to arguing subjectivity is believesTrue negative.

"Republicans roared onto the post-State-of-the-Union morning shows denying that President Obama helped the middle class"

E1 gfbf (obama, goodFor (helping,help:lexEntry), the middle class)
 B1 subjectivity (republicans, negative believesTrue (accusing), E1)
 B2 privateState (writer, positive believesTrue (""), B1)
 Prop1 p(B1,substantial)
 S1 subjectivity (writer, negative sentiment (roared), republicans)

100 writer positive believesTrue
 98 republicans negative believesTrue substantial
 95 obama helping the middle class

102 writer negative sentiment
 99 republicans

rule5source

S1 sentiment toward S2 in the input &
 (Assume S1 positive believesTrue) S2 privateState toward Z in input
 ⇒ S1 sentiment toward S2 privateState toward Z

Preconditions:

102 writer negative sentiment
 99 republicans

Assumptions:

100 writer positive believesTrue
 98 republicans negative believesTrue substantial
 95 obama helping the middle class

⇒ Infer Node:

1229 writer negative sentiment
 98 republicans negative believesTrue substantial
 95 obama helping the middle class

rule3.2

S1 sentiment toward S2 positive/negative believesTrue substantial Z
 ⇒ S1 agrees/disagrees with S2 that isTrue/isFalse Z &
 S1 pos/neg believesTrue substantial Z

Preconditions:

1229 writer negative sentiment
 98 republicans negative believesTrue substantial
 95 obama helping the middle class

⇒ Infer Node:

1231 writer disagrees with republicans that
 1230 isFalse
 95 obama helping the middle class
 1232 writer positive believesTrue substantial
 95 obama helping the middle class

The system has just inferred that the writer disagrees with the republicans that it is false that Obama is helping the middle class (node 1231) and, thus, that the writer believes that Obama *is* helping the middle class (node 1232).

10.10 Blocked Inferences

Currently, we do not have general criteria for blocking inferences toward the *object of a gfbf* which is in turn the target of a sentiment. Only ad hoc evidence to the contrary does so.

The framework commits to one case of blocking inference toward an *inanimate agent of a gfbf* which is in turn the target of a sentiment: when there is evidence against the gfbf being *substantial* (i.e., that it happened). Consider this example:

“Mother was worried that the tree might fall on the boy, but it didn't”

E1 gfbf (the tree:thing, badFor (fell on,fall on:lexEntry), the boy)

S1 subjectivity (mother, negative sentiment (worried), E1)

B1 privateState (writer, positive believesTrue (“”), S1)

B2 privateState (writer, negative believesTrue (“”), E1)

Prop1 p(B2,substantial)

110 writer negative believesTrue substantial

103 the tree fell on the boy

108 writer positive believesTrue

106 mother negative sentiment

103 the tree fell on the boy

Actually, the inference that mother is negative toward the tree is not blocked by negative evidence; rather, the relevant rule, rule9 does not fire. Here is rule9:

rule9

S sentiment toward A goodFor/badFor T, where A is a thing &

(Assume S positive believesTrue substantial) A goodFor/badFor T

⇒ S sentiment toward A

The assumption cannot be satisfied, because the only belief toward the gfbf is writer **negative** believesTrue; for the assumption to be made, the polarity would need to be positive. Thus, the system infers that mother has a positive sentiment toward the boy, but not that she has a negative sentiment toward the tree (see Appendix A).

There are two types of evidence we hypothesize systematically block inferences toward an *animate agent of a gfbf* which is in turn the target of a sentiment: evidence against the gfbf being intentional (e.g., not an accident), and evidence against the agent being positive toward the idea of the gfbf. Either type of evidence blocks the reasoning chain illustrated in Section 10.3, namely the chain from gfbf, to intention toward the gfbf, toward positive sentiment toward the idea of the gfbf, to agreement/disagreement with the agent, to a positive/negative attitude toward the agent.

Consider the following example, which we looked at above in Section 7:

“Oh no! The tech staff tried to stop the virus, but they failed.”
 E1 gfbf ⟨the tech staff, badFor (stop,stop:lexEntry), the virus (virus:lexEntry)⟩
 I1 influencer ⟨the tech staff, reverse (failed,fail:lexEntry), E1⟩
 S1 subjectivity ⟨writer, negative sentiment (Oh no!),I1⟩
 V1 evidence ⟨none,positive intends (tried),E1⟩
 V2 evidence ⟨none,negative believesTrue (failed),E1⟩

115 writer negative sentiment
 114 the tech staff ⟨reverse⟩
 111 the tech staff stop the virus
 118 There is evidence that the following is not substantial
 111 the tech staff stop the virus
 117 There is evidence that the following is intentional:
 111 the tech staff stop the virus

influencer node:

115 *writer negative sentiment*
 114 the tech staff ⟨reverse⟩
 111 the tech staff stop the virus

new gfbf node:

1262 *writer negative sentiment*
 1259 the tech staff goodFor the virus

new evidence nodes:

1260 There is evidence that the following is not intentional:
 1259 the tech staff goodFor the virus
 1261 There is evidence that the following is substantial
 1259 the tech staff goodFor the virus

Recall from Section 7 that *failure* is a *reverser* influencer, resulting in new gfbf and evidence nodes, as just shown.

The system infers that the writer is negative toward the virus, but not that he is negative toward the tech staff. Specifically, rule6 is blocked from firing:

rule6

A goodFor/badFor T, where A is animate
 \Rightarrow A intended A goodFor/badFor T

Preconditions:

1262 *writer negative sentiment*
 1259 the tech staff goodFor the virus

\Rightarrow **Infer Node:**

Blocked by evidence node
 The evidence node:
 1260 ... the following is not intentional:
 1259 the tech staff goodFor the virus

The following sentence illustrates an inference being blocked because the agent is not positive toward their intentional event.

“Thanks to the Affordable Care Act, consumers will receive more value for their premium dollar because insurance companies will be required to spend 80 to 85 percent of premium dollars on medical care and health care quality improvement, rather than on administrative costs, starting in 2011.”
 E1 gfbf ⟨insurance-companies, goodFor (spend on,spendOn:lexEntry), health-care-quality-improvement)⟩

B1 privateState ⟨writer, positive believesTrue (""), E1⟩
 S1 subjectivity ⟨writer, positive sentiment (Thanks & value), E1⟩
 V1 evidence ⟨insurance-companies, negative sentiment (required), E1⟩

 122 writer positive believesTrue
 119 insurance-companies spend on health-care-quality-improvement
 124 writer positive sentiment
 119 insurance-companies spend on health-care-quality-improvement
 125 (evidence,insurance-companies,negative,sentiment)
 126 ideaOf
 119 insurance-companies spend on health-care-quality-improvement

The input includes one gfbf, node 119; sentiment on the part of the writer, node 124 (*Thanks to, value*); and evidence against the agent being positive toward the event. The gfbf is intentional – it will not be an accident that the insurance companies will spend their premium dollars – but, they will be *forced to* (and won't like it). The evidence node blocks rule7 from firing, and, thus, the system does *not* infer that the writer is positive toward the insurance companies, even though it does infer the writer is positive toward health-care-quality-improvement, and the insurance companies are doing something goodFor it (see Appendix A).

11 Exploring Interdependent Ambiguities

Recall this example from Section 10.2:

"Is it no surprise then that MoveOn would attack Senator McCain!?"
 4 writer negative sentiment
 1 MoveOn attack Senator McCain

Note that the subjectivity clues in this sentence – fronting, *surprise*, and *then* – are ambiguous with respect to whether the writer is expressing a positive or negative sentiment. If the writer were positive rather than negative toward MoveOn attacking McCain, then the polarities of the inferred sentiments would be reversed. Considering all the rules together, we really only have two choices rather than eight: (1) $\text{polarity}(\textit{surprise}) = \text{negative}$; $\text{sent}(\text{writer}, \text{McCain}) = \text{positive}$; $\text{sent}(\text{writer}, \text{MoveOn}) = \text{negative}$, or (2) $\text{polarity}(\textit{surprise}) = \text{positive}$, $\text{sent}(\text{writer}, \text{McCain}) = \text{negative}$; $\text{sent}(\text{writer}, \text{MoveOn}) = \text{positive}$. From an NLP perspective, the interdependencies captured by the implicature rules may be encoded as constraints to support sentiment propagation among entities (Deng and Wiebe (2014)) and as constraints in an optimization framework for joint disambiguation.¹⁰

Further, this example illustrates that evidence from the larger discourse or pragmatic situation could be exploited to improve sentence-level processing. Any outside evidence (for example, that the writer is politically liberal or conserva-

¹⁰ In submission.

tive) concerning the writer’s attitude toward McCain or MoveOn could be used to determine which **set** of attitudes is more probable.

12 A Return to the KR Scheme: S **negative believesTrue** T

This section may safely be skipped.

We stated above in Section 5 that S **negative believesTrue** T means (A) or (B):

- (A) S does not believe that T, in the sense that T is not in S’s belief space. Among the things that S believes are true, T is not one of them. S may not believe anything about T.
- (B) S does not believe that T, in the sense that S believes that not T, i.e., S believes that T is false.

Consider this sentence given in Section 10.10:

"Mother was worried that the tree might fall on the boy, but it didn't"

E1 gfbf ⟨the tree:thing, badFor (fell on,fall on:lexEntry), the boy⟩

S1 subjectivity ⟨mother, negative sentiment (worried), E1⟩

B1 privateState ⟨writer, positive believesTrue (""), S1⟩

B2 privateState ⟨writer, negative believesTrue (""), E1⟩

Prop1 p(B2,substantial)

In this sentence, the tree did not in fact fall on the boy. Thus, we have case (B). The input line corresponding to the tree not falling is B2. This line means either (A) the writer has an absence of belief that the tree fell on the boy, or (B) the writer believes that it is false that the tree fell on the boy. Our KR scheme does not “nail down” which of (A) or (B) applies to this sentence, even though a human can perceive that it is (B).

Contrast this with the following (which is included in the appendix output file):

"Mother dislikes the judge; by the way, he freed the murderer, but Mother doesn't know he did."

S1 subjectivity ⟨Mother, negative sentiment (dislikes), the judge⟩

B1 privateState ⟨writer, positive believesTrue (""), S1⟩

E1 gfbf ⟨the judge, goodFor (freeing,free:lexEntry), the murderer⟩

B2 privateState ⟨writer, positive believesTrue (""), E1⟩

B3 privateState ⟨mother, negative believesTrue (doesn't know), E1⟩

B4 privateState ⟨writer, positive believesTrue (""), B3⟩

The relevant private state is B3, that mother negative believesTrue that the judge freed the murderer. This sentence illustrates case (A): we only know she has the absence of belief that the judge freed the murderer. The sentence does not suggest that mother believes that it is false that the judge freed the murderer. Again, our KR scheme does not commit to which of (A) or (B) applies, even though a human can perceive that it is (A).

Consider the arguing subjectivity examples from Section 10.9. The positive arguing true example we gave is this:

"Republicans roared onto the post-State-of-the-Union morning shows accusing President Obama of waging class warfare against the rich"

The relevant private state is that the Republicans positive believeTrue that Obama badFor the rich. The system makes several inferences, leading to these:

```
1029 writer disagrees with republicans that
    1028 isTrue
        77 obama waging class warfare against the rich
1030 writer negative believesTrue substantial
    77 obama waging class warfare against the rich
```

Together, 1029 and 1030 demonstrate a (conceptual) weakness of our KR scheme combined with this inference rule. Node 1029 implies that the writer believes it is false that obama badFor the rich (since the writer disagrees with the republicans that it is true). That is, 1029 suggests that case (B) applies. But, node 1030 loses that information: all we know from 1030 is that case (A) or case (B) applies.

We do not anticipate this being a problem for our project as it proceeds, as we are not aiming to make such fine-grained belief distinctions. If it does become a problem, we can refine the KR scheme. This would require adding logical negation (and rules to reason with it) to the rule-based framework, which would increase its overall complexity.

Note that we *are* incorporating negation into the framework via the compositionality component (see Section 7). Compositional processing is performed *prior* to the inference process. An advantage of this is that the two processes - semantic compositionality and pragmatic inference - do not need to share the same representations or mechanisms.

13 A Return to Space Extension: Extensions that could potentially be added

This section may safely be skipped.

As described in Section 6.4, when a rule fires, the As and Qs are placed into all the spaces that already contain all the Ps. Then, the Ps, As, and Q's are added to the spaces created by replacing sentiments with beliefs on those paths. The idea is, if S has a sentiment toward T, then S must have a positive believesTrue attitude toward T. In this section, we will call the new spaces the *expected spaces*.

There are cases where we opt not to attribute beliefs, when arguably it would be valid to attribute them. These cases are beliefs toward gfbfs, ideaOfs, animate entities, and things.

Consider these two segments from a *by spaces* display:

```
Mother is upset that the judge freed the murderer.
[215 writer +B mother -S]
```

```
[216 writer +B mother +B]
213 judge positive intends
    5 judge freeing the murderer
[212 writer +B mother -S]
[225 writer +B mother -S judge +S]
[227 writer +B mother +B judge +S]
6 the murderer
```

Node 216 is the result of adding a conclusion to one of the *expected* spaces, i.e., the result of replacing -S in the path of 215 with +B. Note that such nodes are excluded for node 6: the rightmost node of all the spaces in which 6 appears is a sentiment.

The difference is that 213 is a private state, and 6 is an entity. private states are propositions; they naturally may be the objects of *believes that*. While we have a means for representing *believes that* having an object that is not a proposition (we specify that the source believes that p(object) for some p) it seems more natural not to introduce them without a reason to.

The case of *believes that* objects that are not propositions which currently arises is when they are added to the input, for example, node 87 in the following:

```
"" Mother dislikes the judge for freeing the murderer.
88 writer positive believesTrue
    83 Mother negative sentiment
        84 the judge
87 writer positive believesTrue
    86 the judge freeing the murderer
```

We wanted to allow inputs such as 87; we tried some other more complex input schemes, but the reactions to them suggested we should stick to the simpler scheme.

Now, perhaps all possible beliefs should be generated. After all, if you have a sentiment toward something you have to believe *something* about it. In that case, the entry for 6 would be the following (the starred nodes would be the additional one).

```
Mother is upset that the judge freed the murderer.
[212 writer +B mother -S]
[225 writer +B mother -S judge +S]
[227 writer +B mother +B judge +S]
[* writer +B mother +B]
[* writer +B mother +B judge +B]
6 the murderer
```

This would be a straightforward change to make.

14 Related Work

In this section, we first discuss recent related work on sentiment analysis in NLP, then acknowledge older work in NLP and AI whose ideas we exploited to create the inference architecture.

14.1 Sentiment Analysis

While most work in NLP addresses explicit sentiment, there is work that addresses implicit sentiment. Though none presents, as we do, general implicature rules relating gfbf events and explicit and implicit sentiments ascribed to sources/holders, several previous works address relevant aspects of our overall framework. By design, our framework abstracts away from specific linguistic realizations, to capture general underlying inference patterns (for example, gfbf spans may be verbs as well as nouns, and gfbf polarity reversal may involve multiple compositions involving words of different parts of speech). Many previous papers address relevant topics in more linguistic depth than we do. Our hope is that their results may be exploited in the future to realize fully automatic framework components, and that our work will help integrate their various findings.

Researchers have investigated identifying objective words that have positive or negative connotations (e.g., Feng et al (2013)) and identifying noun product features that imply opinions (e.g., Popescu and Etzioni (2005); Zhang and Liu (2011)). Rule 10 in our schema shows where connotation is brought into the framework: in the absence of evidence to the contrary, Rule 10 infers sentiment from connotation. The simple lexicon test for connotation in the precondition of Rule 10 could be replaced by a more sophisticated recognition component that handles related notions such as polar features.

Several papers apply compositional semantics to determine polarity (e.g., Moilanen and Pulman (2007); Choi and Cardie (2008); Moilanen and Pulman (2009); Moilanen et al (2010); Neviarouskaya et al (2011); Ruppenhofer and Rehbein (2012)). The goal of such work is to determine one overall polarity of an expression or sentence (though polarities may be assigned to intermediate entities along the way). Conceptually, in terms of the framework, such composition is performed before inference begins. (Section 7 describes the compositionality currently performed by the system, namely processing influencer chains ending in gfbf events. Compositional processing should also be incorporated into its recognition of connotation for Rule 10, building on the previous work just cited.) In contrast to compositional processing, the implicature rules infer sentiments of *different* sources/holders toward various events and entities in the sentence. In addition, their inferences are defeasible.

States and events (such as gfbf events) which positively or negatively affect entities have figured in several works in sentiment analysis. For example, two papers mentioned above Zhang and Liu (2011); Choi and Cardie (2008) include linguistic patterns for the tasks that they address that include gfbf events. Goyal et al (2012), in their work toward automating plot units, generate a lexicon of *patient polarity verbs*, which correspond to gfbf events whose spans are verbs. Riloff et al (2013), in their work on recognizing sarcasm on Twitter, learn phrases describing situations which are negative for the Tweeter.

Turning to the *inference* of implicit from explicit sentiment, Zhang and Liu (2011) introduce two specific rules: *DecreasedNeg* \rightarrow *Positive* and *DecreasedPos* \rightarrow *Negative* (p. 577), which apply to phrases such as *reduce the fun of driving*. Goyal et al (2012) infer the affective states of characters in fables, for example, that Mary has positive affect from *Mary laughed* and that John has positive affect from *John was rewarded*. These are different inferences from the ones we address: we do not infer the affective states of entities mentioned in the text, but rather sentiments held by the writer and other entities toward the entities and events in the sentence. Though Riloff et al (2013) do not perform inference, per se, they do address contrasts between polarities of explicit and implicit sentiments as we do – in their model, sarcasm arises from positive sentiment toward negative situations.

In Section 6.3, we saw that inferences toward agents may be blocked by evidence breaking an inference chain such as *Agent gfbf Object* \rightarrow *Agent intended Agent gfbf Object* \rightarrow *Agent positive sentiment toward ideaOf Agent gfbf Object*. These inference chains were partially inspired by the study of implicit sentiment by Greene and Resnik (2009). They investigated a “connection between implicit sentiment and grammatically relevant semantic properties . . . by varying the syntactic form of event descriptions” (p. 505) and show that the semantic properties of descriptions predict perceived sentiment. They constructed stimuli of the form *X verb of killing Y* in the context of news reports about crimes, where *Y* is a victim. Since killing a victim is rarely viewed positively, their stimuli correspond, in our terms, to:

E_1 : $\langle X, \text{kill:BADFOR}, Y \rangle$
 S_1 : $\text{sent}(\text{writer}, E_1) = \text{negative}$

They varied the syntactic form in ways corresponding to semantic properties. Subjects were asked to rate how sympathetic they perceive a stimulus to be toward the agent, *X*. They found that volition is negatively correlated with sympathy – the more volitional the act, the more negative the judgment against *X*. Consistent with their findings, in our system, evidence against *X*'s action being volitional blocks a negative inference toward *X*. Their work will be relevant in future work on detecting linguistic evidence against agent inferences, i.e., detecting evidence for defeated implicatures.

14.2 Inference Architecture

Our inference architecture involves explicit rules and mechanisms for default inference and inference within private-state spaces.

The assumptions in the rules are inspired by Hobbs et al (1993), where interpreting a text is cast in the form of abduction, and assumptions are made as necessary to derive an interpretation. All assumptions in our rules are assumptions of private states, either beliefs or sentiments. Our ascription of belief was inspired by work of Wilks and colleagues (e.g., Wilks and Bien (1983); Wilks and Ballim (1987)) who devise a *default ascriptional rule* which assumes “one’s view of another person’s view is the same as one’s own *except where there is explicit evidence to the contrary*” (Wilks and Ballim (1987)(p. 119)).

Our main inspiration for default inference is work done in the 1970's by Schank and his research group. The inference performed by our system is forward inference, triggered by the input, resulting in reasoning chains that proceed to the end unless evidence to the contrary breaks the chain (halts inference). This is similar to the type of reasoning performed by SAM (Script Applier Mechanism) (Schank and Abelson (1977)), for example (though our rules do not instantiate roles along the way).

Our main inspiration for inference within private-state spaces is work by Stuart Shapiro, William Rapaport, and their research group (Rapaport (1986); Shapiro and Rapaport (1987)) on their knowledge representation system, SNePS, considered as a fully intensional propositional semantic network, and the extension of SNePS, created by João Martins, to handle reasoning in multiple belief spaces

15 Conclusions

While previous sentiment analysis research has concentrated on the interpretation of explicitly stated opinions and attitudes, this work initiates the computational study of a type of opinion implicature (i.e., opinion-oriented inference) in text. This paper described a rule-based framework for representing and analyzing opinion implicatures which we hope will contribute to deeper automatic interpretation of subjective language. In the course understanding implicatures, the system recognizes implicit sentiments (and beliefs) toward various events and entities in the sentence, often attributed to different sources and of mixed polarities; thus, it produces a much richer interpretation than is typical in opinion analysis.

References

- Anand P, Reschke K (2010) Verb classes as evaluativity functor classes. In: Interdisciplinary Workshop on Verbs. The Identification and Representation of Verb Features
- Banfield A (1982) *Unspeakable Sentences*. Routledge and Kegan Paul, Boston
- Breck E, Choi Y, Cardie C (2007) Identifying expressions of opinion in context. In: IJCAI'07: Proceedings of the 20th international joint conference on Artificial intelligence, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp 2683–2688
- Bruder G, Wiebe J (1995) Recognizing subjectivity and identifying subjective characters in third-person fictional narrative. In: Duchan J, Bruder G, Hewitt L (eds) *Deixis in Narrative: A Cognitive Science Perspective*, Lawrence Erlbaum Associates
- Choi Y, Cardie C (2008) Learning with compositional semantics as structural inference for subsentential sentiment analysis. In: Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Honolulu, Hawaii, pp 793–801, URL <http://www.aclweb.org/anthology/D08-1083>
- Choi Y, Cardie C, Riloff E, Patwardhan S (2005) Identifying sources of opinions with conditional random fields and extraction patterns. In: HLT '05: Proceedings of the conference on Human Language Technology and Empirical Methods in

- Natural Language Processing, Association for Computational Linguistics, Morristown, NJ, USA, pp 355–362
- Choi Y, Breck E, Cardie C (2006) Joint extraction of entities and relations for opinion recognition. In: Proceedings of EMNLP
- Das S, Chen M (2001) Yahoo for amazon: Extracting market sentiment from stock message boards. In: Proceedings of the 8th Asia Pacific Finance Association Annual Conference.
- Dave K, Lawrence S, Pennock DM (2003) Mining the peanut gallery: opinion extraction and semantic classification of product reviews. In: WWW '03: Proceedings of the 12th international conference on World Wide Web, ACM, New York, NY, USA, pp 519–528
- Deng L, Wiebe J (2014) Sentiment propagation via implicature constraints. In: Meeting of the European Chapter of the Association for Computational Linguistics (EACL-2014)
- Deng L, Choi Y, Wiebe J (2013) Benefactive/malefactive event and writer attitude annotation. In: Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Association for Computational Linguistics, Sofia, Bulgaria, pp 120–125, URL <http://www.aclweb.org/anthology/P13-2022>
- Doran R, Ward G, Larson M, McNabb Y, Baker RE (2012) A novel experimental paradigm for distinguishing between ‘what is said’ and ‘what is implicated’. *Language* 88(1):124–154
- Feng S, Bose R, Choi Y (2011) Learning general connotation of words using graph-based algorithms. In: Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Edinburgh, Scotland, UK., pp 1092–1103, URL <http://www.aclweb.org/anthology/D11-1101>
- Feng S, Kang JS, Kuznetsova P, Choi Y (2013) Connotation lexicon: A dash of sentiment beneath the surface meaning. In: Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Sofia, Bulgaria, pp 1774–1784, URL <http://www.aclweb.org/anthology/P13-1174>
- Goyal A, Riloff E, Daumé III H (2010) Automatically producing plot unit representations for narrative text. In: Empirical Methods in Natural Language Processing (EMNLP), Boston, MA, URL <http://pub.ha13.name/#daume10daal>
- Goyal A, Riloff E, Daume III H (2012) A computational model for plot units. *Computational Intelligence*
- Greene S, Resnik P (2009) More than words: Syntactic packaging and implicit sentiment. In: Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Association for Computational Linguistics, Boulder, Colorado, pp 503–511, URL <http://www.aclweb.org/anthology/N/N09/N09-1057>
- Grice HP (1967) *Logic and conversation*. The William James lectures
- Grice HP (1989) *Studies in the Way of Words*. Harvard University Press
- Hobbs JR, Stickel ME, Appelt DE, Martin P (1993) Interpretation as abduction. *Artificial Intelligence* 63(1-2):69–142
- Kim SM, Hovy E (2005) Automatic detection of opinion bearing words and sentences. In: Companion Volume to the Proceedings of IJCNLP-05, the Second International Joint Conference on Natural Language Processing, Jeju Island,

- KR, pp 61–66
- Lehnert WG (1981) Plot Units and Narrative Summarization. *Cognitive Science* 5(4):293–331
- Levinson S (1983) *Pragmatics*. Cambridge Textbooks in Linguistics, Cambridge University Press, URL <http://books.google.com/books?id=SJXr9w\1VLUC>
- Liu B (2012) *Sentiment Analysis and Opinion Mining*. Morgan & Claypool
- Moilanen K, Pulman S (2007) Sentiment composition. In: *Proceedings of RANLP 2007*, Borovets, Bulgaria
- Moilanen K, Pulman S (2009) Multi-entity sentiment scoring. In: *Proceedings of the International Conference RANLP-2009*, Association for Computational Linguistics, Borovets, Bulgaria, pp 258–263, URL <http://www.aclweb.org/anthology/R09-1048>
- Moilanen K, Pulman S, Zhang Y (2010) Packed feelings and ordered sentiments: Sentiment parsing with quasi-compositional polarity sequencing and compression. In: *Proceedings of the 1st Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA 2010)*, pp 36–43
- Neviarouskaya A, Prendinger H, Ishizuka M (2011) Affect analysis model: novel rule-based approach to affect sensing from text. *Natural Language Engineering* 17(01):95–135
- Pang B, Lee L (2008) Opinion mining and sentiment analysis. *Found Trends Inf Retr* 2(1-2):1–135
- Pang B, Lee L, Vaithyanathan S (2002) Thumbs up?: sentiment classification using machine learning techniques. In: *EMNLP '02: Proceedings of the ACL-02 conference on Empirical methods in natural language processing*, Association for Computational Linguistics, Morristown, NJ, USA, pp 79–86
- Pearl J (1982) Reverend bayes on inference engines: A distributed hierarchical approach. In: *Proceedings of the American Association of Artificial Intelligence National Conference on AI*, Pittsburgh, PA, pp 133–136
- Popescu AM, Etzioni O (2005) Extracting product features and opinions from reviews. In: *Proceedings of the Human Language Technologies Conference/Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP-2005)*, Vancouver, Canada, pp 339–346, URL <http://www.aclweb.org/anthology/H/H05/H05-1043.pdf>
- Qiu G, Liu B, Bu J, Chen C (2011) Opinion word expansion and target extraction through double propagation. *Comput Linguist* 37:9–27
- Quirk R, Greenbaum S, Leech G, Svartvik J (1985) *A Comprehensive Grammar of the English Language*. Longman, New York
- Rapaport WJ (1986) Logical foundations for belief representation. *Cognitive Science* 10(4):371–422
- Reschke K, Anand P (2011) Extracting contextual evaluativity. In: *Ninth International Conference on Computational Semantics (IWCS 2011)*
- Riloff E, Qadir A, Surve P, De Silva L, Gilbert N, Huang R (2013) Sarcasm as contrast between a positive sentiment and negative situation. In: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Seattle, Washington, USA, pp 704–714, URL <http://www.aclweb.org/anthology/D13-1066>
- Ruppenhofer J, Rehbein I (2012) Semantic frames as an anchor representation for sentiment analysis. In: *Proceedings of the 3rd Workshop in Computational Approaches to Subjectivity and Sentiment Analysis*, Association for Computational

- Linguistics, Jeju, Korea, pp 104–109, URL <http://www.aclweb.org/anthology/W12-3716>
- Schank R, Abelson R (1977) *Scripts, Plans, Goals and Understanding*. Lawrence Erlbaum Associates, Hillsdale, New Jersey
- Shapiro SC, Rapaport WJ (1987) Sneys considered as a fully intensional propositional semantic network. In: Cercone N, McCalla G (eds) *The Knowledge Frontier*, Springer, New York, pp 262–315
- Stein D, Wright S (eds) (1995) *Subjectivity and Subjectivisation*. Cambridge University Press, Cambridge
- Stoyanov V, Cardie C (2008) Topic identification for fine-grained opinion analysis. In: *Proceedings of COLING*
- Wiebe J (1994) Tracking point of view in narrative. *Computational Linguistics* 20(2):233–287
- Wiebe J, Mihalcea R (2006) Word sense and subjectivity. In: *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Sydney, Australia, pp 1065–1072, DOI 10.3115/1220175.1220309, URL <http://www.aclweb.org/anthology/P06-1134.pdf>
- Wiebe J, Riloff E (2005) Creating subjective and objective sentence classifiers from unannotated texts. In: *Sixth International Conference on Intelligent Text Processing and Computational Linguistics (CICLing-2005)*, Mexico City, Mexico
- Wiebe J, Wilson T, Cardie C (2005) Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation (formerly Computers and the Humanities)* 39(2/3):164–210
- Wilensky R (1978) *Understanding goal-based stories*. PhD thesis, Yale University, Also available as Tech Report YALEU/CSD/RR #140
- Wilks Y, Ballim A (1987) Multiple agents and the heuristic ascription of belief. In: *Proceedings of the 10th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI'87*, pp 118–124
- Wilks Y, Bien J (1983) Beliefs, points of view and multiple environments. *Cognitive Science* 7:95–119
- Wilson T, Wiebe J, Hoffmann P (2005) Recognizing contextual polarity in phrase-level sentiment analysis. In: *Proceedings of the Human Language Technologies Conference/Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP-2005)*, Vancouver, Canada, pp 347–354
- Yi J, Nasukawa T, Bunescu R, Niblack W (2003) Sentiment analyzer: Extracting sentiments about a given topic using natural language processing techniques. *Data Mining, IEEE International Conference on* 0:427
- Zhang L, Liu B (2011) Identifying noun product features that imply opinions. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Portland, Oregon, USA, pp 575–580, URL <http://www.aclweb.org/anthology/P11-2101>
- Zúñiga F, Kittilä S (2010) Introduction. In: Zúñiga F, Kittilä S (eds) *Benefactives and malefactives, Typological studies in language*, J. Benjamins Publishing Company, URL <http://books.google.com/books?id=B8tzvbjurfwC>