

Finding the Sources and Targets of Subjective Expressions

Josef Ruppenhofer, Swapna Somasundaran, Janyce Wiebe

Intelligent Systems Program, Dept. of Computer Science, Dept. of Computer Science
University of Pittsburgh, Pittsburgh, PA 15260
josefr@cs.pitt.edu, swapna@cs.pitt.edu, wiebe@cs.pitt.edu

Abstract

As many popular text genres such as blogs or news contain opinions by multiple sources and about multiple targets, finding the sources and targets of subjective expressions becomes an important sub-task for automatic opinion analysis systems. We argue that while automatic semantic role labeling systems (ASRL) have an important contribution to make, they cannot solve the problem for all cases. Based on the experience of manually annotating opinions, sources, and targets in various genres, we present linguistic phenomena that require knowledge beyond that of ASRL systems. In particular, we address issues relating to the attribution of opinions to sources; sources and targets that are realized as zero-forms; and inferred opinions. We also discuss in some depth that for arguing attitudes we need to be able to recover propositions and not only argued-about entities. A recurrent theme of the discussion is that close attention to specific discourse contexts is needed to identify sources and targets correctly.

1. Introduction

Sentiment analysis is an active research area encompassing work on acquiring lexica of opinion expressions, recognizing such expressions in text and classifying phrases, sentences, or documents as objective, positive, or negative (Hatzivassiloglou and McKeown, 1997; Riloff and Wiebe, 2003; Pang and Lee, 2004; Esuli and Sebastiani, 2006; Andreevskaia and Bergler, 2006; Takamura et al., 2007). Other efforts focus on applications like movie and product review mining (e.g. Turney, 2002; Hu & Liu, 2004).

Here we want to contribute towards progress on the sub-tasks of recognizing opinion sources (holders) and targets (topics). As research turns to new genres, these tasks become more important. For instance, in meetings and blogs several sources may be present and several issues under discussion. Consider example (1) below, a post taken from the site of the New York Times commentator Stanley Fish.

- (1) This is right on the money. The only question is: will anyone from the media - from beat reporters and bloggers to editors and the commentariat and punditocracy (including those of the NY-Times) read this essay and understand that Mr Fish is talking about them.
They never have gotten the message when Paul Krugman takes them to task, so here's betting the answer is no.

In this post, the writer never explicitly refers to himself even though he expresses several opinions, e.g. "This is right on the money". The poster also attributes opinions and speech to other people, namely "Mr Fish", who wrote the article that the post is in response to, and "Paul Krugman" who is another New York Times commentator. Clearly, a system that processes such posts for Question Answering or summarization purposes ought to distinguish the various sources and how the individual opinions are attributed to

them. For instance, it should treat a report of Paul Krugman's opinion on the Fish blog differently from an expression of opinion by Krugman in an article authored by him. Such systems also need a way of identifying unstated sources of opinion, such as the writer in the case of (1).

While automatic semantic role labeling is known to help source and target recognition, it cannot be the whole solution. This is also suggested by the example above: the writer who does express his opinions never figures as the filler of any semantic role in the passage. In this work, we want to deepen the understanding of what finding sources and targets requires, and discuss issues that need to be addressed for a more complete solution. The contribution we offer is based on the practical experience of annotating subjective expressions, sources, and targets in various written and spoken genres.

The organization of this paper is as follows. In section 2 we describe our annotation scheme. Section 3 discusses how role labeling systems can help in finding sources and targets of subjective expressions. Sections 4 and 5 present cases in which role labeling systems cannot identify (all) the opinion roles. In section 6, we situate our work in the research context. We present ideas for further research in section 7 and conclude in section 8.

2. Annotation scheme

The basic MPQA scheme (Wiebe et al., 2005) concerns words and phrases that are used in particular contexts to express **private states** such as emotions, evaluations, speculations, etc. Following Quirk et al. (1985), private states are defined as states that are not open to objective observation or verification. They are the states of **sources** holding **attitudes**, optionally toward **targets**.

The annotation scheme distinguishes three types of private state expressions:

- (2) explicit mentions of private states:
He was **boiling with anger**
- (3) speech events expressing private states:
The paper's editors **attacked** the new House

This research was supported in part by the Department of Homeland Security under grant N000140710152.

Speaker.

- (4) expressive subjective elements:
That doctor is a **quack**.

We group explicit mentions of private states such as *boiling with anger* in (2) and speech events expressing private states such as *attacked* in (3) together as direct subjective expressions (DSEs) but call *quack* in (4) an expressive subjective element (ESE). The key difference between direct subjective expressions and expressive subjective elements is that only the former explicitly introduce a private state and its source while the latter presuppose a private state and source but do not introduce them.

The scheme records the sources of private states in terms of nested levels.

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

In (5), on the top level, the writer presents the whole sentence. One level down, we have the speaking of Xirao-Nima, for which we record <writer,Xirao-Nima> as source since we know of this event only *by way of* the author. On the lowest level, we have *fear* with source <writer, Xirao-Nima, US>. Unlike *fear*, the two speaking events in (5) do not express private states of their sources. They are therefore assigned to a separate category of objective speaking events.

In an extension to the basic scheme, attitudes and targets were added to the annotation scheme (Wilson and Wiebe, 2005; Wilson, 2008). The attitude types that Wilson (2008) defines include Arguing, Sentiment, Agreement, Speculation, Intention, and an Other category. In this paper, we will mainly refer to the first two types, which are illustrated in (6) and (7). The notation used here and throughout the paper is the following: sources appear in small caps, private state expressions appear in boldface, attitude spans are rendered in italics, and target spans appear as plain text, enclosed by square brackets.

- (6) AFRICAN OBSERVERS *generally approved* of [his victory], while WESTERN GOVERNMENTS **denounced** [it].
- (7) She highlighted the situation of the ethnic Uighur Muslim population of Xinjiang, Chinas furthest west region, where rights groups **have claimed** [*a crackdown against separatists*] *has been stepped up since the September 11 attacks on the United States*.

Example (6) contains two sentiment attitudes attached to the private state expressions “approved” and “denounced”; the former expresses a positive sentiment, and the latter a negative one. Example (7) displays an instance of positive arguing that a crackdown has intensified. Arguing, like Sentiment, can be positive or negative, and it can either be concerned with what should or should not be done, or with what is or is not the case.

Conceptually, one or several attitudes together make up a private state. In the annotation scheme, this relationship is captured by linking each attitude annotation to the appropriate direct subjective annotation. Targets are associated

with the individual attitudes rather than directly with the direct subjective annotation. An attitude annotation is intended to capture the span of text that expresses a particular attitude that is part of a given private state.

3. Automatic semantic role labeling

Sources and targets are pragmatic roles that occur when words and phrases are *used*. They are different from the semantic or argument roles that characterize frames or word senses *lexically*. Nevertheless, for many lexical items that are typically used as opinion expressions, we can map sources and targets to semantic roles and then use automatic semantic role labeling (ASRL) to identify them in context. Consider the verbs *fear* and *please* in (8) and (9).

- (8) WE **fear** [an early death] much more.
- (9) [This] **pleased** THE MAINLY FEMALE AUDIENCE.

These verbs realize source and target differently. Relative to PropBank (Palmer et al., 2005), the source for *fear*, in that particular syntactic frame, maps to the verb’s Arg0 and the target to its Arg1. For *please* in (9), source and target correspond to Arg1 and Arg0, respectively. Comparable mappings are available for FrameNet (Baker et al., 1998). The source of *fear* maps to the Experiencer frame element, and the target to the Content frame element. For *please*, the source maps to the Experiencer frame element and the target to the Stimulus frame element.

As with semantic role labeling, it is important to identify the correct sense of an opinion expression when identifying sources and targets: the realization of opinion roles may differ between senses along with the inventory and realization of semantic roles. For example, in (10), *thankfully* is a sentence adverb with the quoted speaker *he* as its source and the whole clause “the policemen had ... at bay” as its target, whereas, in (11), *thankfully* functions as a manner adverb with *he* as source and with no expression of the target in the clause “he nodded thankfully”.

- (10) HE **said**: “**Thankfully**, the policeman had the presence of mind to pull out one of the organ pipes and use it to hold the man at bay.”
- (11) I handed them to him and HE nodded **thankfully**.

In (11), the first of the two conjoined clauses contains the event that “he” is thankful for: being handed aspirin and a glass of water (“them”).

4. Beyond role labeling

Challenges that exceed role labeling roughly fall into four classes: 1) attribution, 2) referent identification, 3) inferences concerning attitudes and their sources and targets, and 4) targets of a less studied subjectivity type, arguing. The first three topics are discussed in this section, the last will be addressed separately in section 5.

4.1. Attribution

Attribution is the relation between beliefs and assertions expressed in text and their sources, that is the agents holding or making them. Expressions such as *condemn* or *argue* grammatically identify the sources of their associated private states. But when ESEs like *quack* in (4) are used in statements of opinion, sources cannot be found via role labeling since they are not expressed as syntactic dependents of ESEs. Instead, they must be identified from among the psychological subjects in the discourse. Explicitly quoted speakers and thinkers are good candidates to be the sources of subjectivity in speech attributed to them.

(12) SENIOR MIKE SHEEHY **said**, “[It] was a **blast**”.

(13) [Car sales] were **said** to be **disappointing**.

In (12), the subjectivity of *blast* is attributed to <writer, Sheehy>. Sources may also be generic, arbitrary, or non-specific, as in (13), where *said* and the ESE *disappointing* share the same non-specific source, <writer, implicit>.

When a quoted sentence contains no specific attribution, knowledge of orthographic conventions can help resolve the right source as in (14), where the sentences following the first express opinions also belonging to the quoted source Joshu Shapiro.

(14) “It’s a necessary thing,” said Joshua Shapiro, the chief United States economist at MFR, a New York economic research firm. “It’s like the mess going down in financial markets. You gotta get through it. The sooner you get through it you can look for better times.”

For unquoted content, a text’s author is a prime candidate for being the source of the opinions expressed by ESEs (cf. 4). However, identifying the right source in the presence of multiple possible sources in the text can be difficult. A valuable resource for studying attribution is the Penn Discourse Treebank (Prasad et al., 2006). Its annotations show that attribution is realized in many different ways and that parts of the same sentence may be attributed to different sources.

We now consider some specific cases in which the attribution of an attitude to a source is not straightforward.

4.1.1. Subjective expressions with multiple sets of opinion roles

Some subjective expressions pair one set of semantic roles with multiple sources and targets. As direct subjective expressions, *slander* and *slime* denote speaking events in which a source (HE; LIBERAL GROUP) criticizes a target ([city];[Petraus]), as shown in (15) and (16).

(15) It is a shame that there is no jury that can mete out justice for [a city] HE has **slandered** for far too long.

(16) THE LIBERAL GROUP’S AD IN YESTERDAY’S NEW YORK TIMES **sliming** [General Petraus] has caused a monumental backlash against the left.

Additionally, *slander* and *slime* serve as ESEs: by describing the two acts of speech as *slander* and *slime*, the authors express negative sentiment towards the speakers ([he];[liberal group]). Examples (17) and (18) show the relevant attitude annotations.

(17) It is a shame that there is no jury that can mete out justice for *a city* [he] has **slandered** for far too long.

(18) [The liberal group’s ad in yesterday’s New York Times] **sliming** General Petraus has caused a monumental backlash against the left.

Note that (17) and (18) contain no linguistic expression of the sources and private states to which the negative sentiment attitudes belong. These are captured as zero-span annotations within the scheme (not shown here).

It is important to capture the two sets of opinion roles and to keep them distinct.

4.1.2. Sources crossing levels

When we have multiple levels of nesting and multiple private states are expressed in the same sentence, we typically find that ESEs are associated with the source of the private state expression into whose immediate syntactic scope they fall. This can be illustrated by the earlier examples (12) and (13). In (12), the positive evaluation expressed by *blast* has as its source *Mike Sheehy*, who is also the source of the speech event expressed by *said*. Similarly, in (13), the negative evaluation expressed by *disappointing* has as its source an unexpressed group of people, who are also the source of the speaking event referred to by *said*. Now compare the above cases with examples (19) and (20):¹

(19) Bill told me that Anna is going to marry [his **idiot** of a son].

(20) To think that some people enjoy having [**these wretched fur balls**] as pets is such a repulsive thought to me.

While it’s possible that in (19) the negative judgment of Bill’s son comes from Bill, it is more plausible that the source of the negative sentiment is the speaker. In accord with that latter interpretation, the negative sentiment of “idiot” is tied to a zero-width direct-subjective annotation placed at the beginning of the sentence that has the speaker as source. The verb “told” of course still refers to an objective speech event by Bill and is annotated accordingly, though that annotation is not displayed in (19). Similarly, the negative sentiment towards cats in (20) is that of the speaker and not that of the people enjoying the company of cats.

Examples (19) and (20) are cases in which an opinion expression of the ESE type does not share the same source as the direct subjective expression that governs it syntactically.

¹To keep the examples simple, only the private state expressions and attitudes that are relevant for the discussion are displayed.

Instead the source of the ESE is “higher” up, typically the writer or speaker presenting the whole sentence.²

4.1.3. Attribution by anaphora

Normally, we find evidence for a speech event or a private state and for its contents in the same sentence. However, sometimes a writer first presents the content in a way that suggests she is the source before attributing it to another source, as in (21).

- (21) I taught best-selling author Brad Meltzer what he knows about reporting gossip. At least that’s what HE **says** in the foreword of his new thriller, “The Book of Fate”.

In (21), “that’s” refers back to the content of the first sentence. Such ‘delayed’ attributions by anaphora may imply that the writer who makes the attribution disagrees with the source about the attributed content. A clear case of this occurs in (22). Here the writer not only uses the anaphor “so” in the second sentence to identify Moyles as the source of the glowing description in the first sentence, but through the *or-so-X-says* construction she also implies that she does not share Moyles’ view of himself.

- (22) [Chris Moyles] is a **brilliant broadcaster, the saviour of Radio 1, a comedian, a best-selling author, and, in fact, a genius**. Or so HE *says*.

4.2. Problems of reference

We now turn from problems associated with the sources of expressive subjective elements to issues concerning the sources and targets of direct subjective expressions. For this latter type of subjective expression, we can typically map opinion roles to semantic roles, as discussed in section 3. However, a major difficulty arises when sources and targets, though in principle expressible and conceptually necessary, are not overtly expressed. Consider examples (23) to (25).

- (23) [A white Ford F-150 pickup] is **believed** to have been involved.
- (24) **Think** about [it] and then let me know.
- (25) Some people say the Steelers are contenders, but I’M **not convinced**.

In (23), the source is an indefinite entity omitted in the passive construction. By contrast, for the imperative in (24), the omitted source of the desired thinking event is the addressee. In (25), a specific target is unexpressed due to the lexical properties of *convinced*. In (24) and (25), unlike (23), the omitted roles can in principle be retrieved from the discourse context. However, as far as we know, neither role labeling nor co-reference systems account for zero-referents.

Another challenge consists of cases where the targets of subjective expressions are entities and occurrences in the physical discourse setting. These so-called exophoric entities serve as targets without having been introduced into the

discourse. For instance, in face-to-face conversation data like multiparty meetings, we often find expressions such as *thanks, sorry, wow* and various interjections, which take such extra-linguistic targets as graphics being displayed on a screen, actions taken while operating equipment, etc. What’s more, extra-linguistic targets are often realized as zero forms. In some cases, speech transcripts are still sufficient to let us infer the entity or event that is the target of such an expression. But in many others, the target cannot be determined with confidence and the video stream needs to be consulted. For example in (26), which is taken from the AMI meeting corpus, it is not clear what speaker D is apologizing for.

- (26) C :: I was a little short on time , but B :: Uh .
 A :: Yeah , me too , so that’s not No no no , I just fi first my
 B :: Yeah , same here .
 D :: Oh .**Sorry** .
 B :: Uh let’s see . Which one was mine ?

A multi-modal annotation scheme keeping track of actions, gestures, and objects in the speech setting would be needed to relate a subjective expression such as “sorry” in (26) to its target.

4.3. Inferences concerning attitudes, sources, and targets

As part of our annotations, we have noticed that under some circumstances we can infer an additional attitude based on one that is explicitly signaled. We record such inferred attitudes in our annotations. Another issue that we discuss here but which is not addressed by our annotation scheme is how attitudes about related entities combine to have an overall effect on the discourse.

4.3.1. Event participants as targets of inferred attitudes

For some events about which opinions are expressed, we can infer additional attitudes towards the participants of these events. Consider examples (27) and (28).

- (27) a. I am not a Colts fan - I am a Bears fan - but I am **glad** [the Colts beat the Patriots].
 b. I am not a Colts fan - I am a Bears fan - but I am **glad** the Colts beat [the Patriots].
- (28) a. I am **sorry** that [she lost that much money].
 b. I am **sorry** that [she] lost that much money.

In (27), in addition to the basic positive attitude towards the game’s outcome (a), we may infer that the speaker feels negatively towards the Patriots, who are negatively affected by the loss (b). Both of these attitudes are part of the same private state, namely the one expressed by “glad”, which therefore appears in boldface in both (a) and (b). Notice that the speaker preempts the possible inference of positive sentiment for the Colts. By contrast, in (28) we have a negative attitude towards the loss of money (a), but can infer a positive attitude towards “she”, who is negatively affected by that event (b).

²This phenomenon is similar to the so-called *de re* versus *de dicto* distinction in intensional contexts (e.g. Quine, 1965).

Inferences like the above do not always go through. In some instances, the opposite inference is in fact supported: in (29), we can infer, based on the report of a “hostile reception”, that the losing team is being blamed for the defeat it suffered rather than pitied.³

- (29) Members of the Pakistan cricket team have arrived home to a hostile reception from FANS *disappointed* about [their] defeat in the final of the World Cup.

In fact, basic attitudes need not be accompanied by any inferred attitudes at all. For instance, if we modified (27) above to say “I am glad the Colts beat the Patriots. I bet \$100 on that result”, possible inferences about the speaker’s sentiments towards the Colts and Patriots become much less likely.

Other inferences target not the participants that are positively or negatively effected by an event but volitional agents that bring about, or fail to bring about, such effects. For instance, in (30), we can infer a negative attitude towards the government for its action, while in (31), the inferred negative attitude towards the Chancellor is based on a failure to act.

- (30) “WE are *disappointed* that [the government] gave the factory permission to deal with more animal waste,” comments Green Party spokesperson Sue Paylor.
- (31) Tony Vine-Lott gave a brief overview of the progress on certain policy issues , such as the demise of the tax credit , stating that PIMA was *disappointed* that [the Chancellor] failed to respond to an industry call to extend the 10% tax credit beyond April 2004 .

Handling such inferred attitudes correctly requires us to use additional knowledge. First, we need to know for as many predicates as possible which of their arguments are negatively or positively affected by an event and which participants are causally responsible. Second, we need to assess what contextual support there is for the possible inferences. Recall for instance, that in (27) above, the speaker blocks the possible inference of positive sentiment towards the Colts.

4.3.2. Target Relations

We can often observe a distinction between global targets and local targets. For instance, in product reviews such as (32), we see opinions expressed both towards the product as a whole, the global target, and towards some of its features, the local targets.

- (32) [The computer] is *very good and very easy to use*. [It] has a built in camera, bluetooth; *the all singing and dancing machine*. Love [it]. *The only glitch* is [the scrolling pad] is *not as smooth* as my last Toshiba notebook. One other thing is that [Vista] is a *nightmare*...

In (32), the computer is the global target and its features, the scrolling pad, the Vista OS, the camera, and the bluetooth interface, are the local targets. Our annotation scheme does not capture which targets are related and we also do not attempt to capture how individual opinions combine. Some work on particular applications does, however, try to solve these issues.

For instance, in the product review genre, it seems to be relatively easy to identify local and global targets given that reviews focus on one product and its features, and given that sets of reviews can be used to identify via statistical analysis what the important features are (e.g. Hu & Liu, 2004) .

In other genres such as movie reviews and task-oriented meetings, the distinction is harder to make. In movie reviews, evaluations and arguments inside the world of the movie have to be distinguished from those that are about the movie and they do not combine: an evil character may make for a positive movie watching experience (Turney, 2002). In meetings, agenda items are not always neatly separated and even if they are, speakers may still engage in side conversations, jokes, or coordination-activities that do not relate to the the task at hand.

One possible approach to the problem of distinguishing but relating local and global targets correctly is to use taxonomies. Bloom et al. (2007) use hand-built domain-dependent taxonomies for movie and product reviews. They point out, however, that the precision of the target classification is the area of most concern to them.

A second challenge for exploiting target relations consists of relating sentiments or arguments about local targets to the global targets of a discourse. In reviews or meetings, a number of positive sentiments or arguments towards local targets may not add up in a simple linear fashion: an overriding sentiment or argument may sway the attitude towards the global target from what one might expect based simply on counting opinions. Detecting such situations requires knowledge of the linguistic mechanisms by which people weight the importance of sentiments and arguments and by which they mark agreement with or acceptance of others’ statements.

5. Targets of arguing

So far we have largely focused on sentiment expressions. Another important kind of subjectivity concerns arguing, that is the expression of beliefs about what is (not) true or what should (not) be done. For instance, the writer of (33) argues that certain readers should take a particular action.

- (33) [Anyone with large credit card debts] *should definitely take up this offer*.

Arguing attitudes are of interest here for two reasons. First, in conversation they are often conveyed by conditional and causal constructions (a) that can have other, non-argumentative interpretations and (b) in which the target expression may occur in different places. Second, we need to interpret the targets of arguing attitudes against the background of complete propositions in order to fully understand the import of the arguments.

³For examples (29) to (31), we only display inferred attitudes.

5.1. Finding Targets in conditional and causal constructions

Some classes of connectives are highly relevant to arguing. Let us consider first how conditionals and sentiment expressions may be combined to argue.

- (34) Your presentation will be better **if** you [put this on the first slide].
- (35) [Vote YES] **if** you want to keep the cost of government in Lewiston low.

Example (34) argues for putting certain material on the first slide. The argument works by relating the argued-for cause to a positive consequence. In (35), the speaker argues for voting affirmatively on a ballot proposition. The argument works by presenting the yes-vote as the logical consequence of a desire that the addressee is assumed to have.

Note that the two examples differ in where the target of the arguing attitudes is located. In (34), the target is contained in the if-clause, whereas in (35), the target is contained in the main clause. This shows that one must consider both clauses when trying to find the target of an argumentative conditional, which underscores the pragmatic nature of subjectivity.

Like conditionals, causal connectives can also be used to make arguments. Very similar meta-patterns are found: causes are blamed for negative effects and credited for positive ones. Again, depending on how the (un)desirable effect is presented, the cause that enables it may be expressed in the *because*-clause or in the main clause.

- (36) [Easy finder with the a whistle function or something , or rechargeable station] **because** it's a pain when you run out of batteries.
- (37) If you're not a good cook, then taking your girlfriend out to an expensive restaurant might be the next best romantic date idea. ... You'll feel good **because** [you've made her happy with a romantic date].

In (36), the target of an argumentative use of *because* is in the (verbless) main clause. The argument presents the argued-for action (including an easy finder function in the design of a remote control) as the result of a desire to avoid a negative future situation. By contrast, in (37), the target is in the subordinate *because*-clause. The argument presents the argued-for action as the cause of a positive future situation.

5.2. Propositions versus Entities as Targets

In the discussion of targets so far, we have tacitly assumed that we know what kinds of things are appropriate targets for particular kinds of opinion expressions. In practice, however, it is often the case that speakers analyzing a discourse might disagree on what span best represents the target of an opinion expression.

5.2.1. The need for propositions

One key difference that often appears is whether speakers consider the target of an opinion expression to be a proposition or an entity. Consider these examples:

- (38) A. What's your favorite color?
B. I really *love* [green]. That's my favorite color.
- (39) A. Who should become chairman?
B. I'M *for* [Sue].
- (40) A. Who do you think took the money.
B. I'D *guess* [Bill]. He is the only one besides you and me who knows the combination.
- (41) A. What's the right punishment for Sue?
B. I **think** [*doing the dishes for the next three weeks*] would be **appropriate**.

In (38), we have a general evaluation of colors, whereas in (39) we find positive sentiment towards Sue's becoming chairman; in (40) we have an argument that Bill was the one who stole the money; and in (41) we have a proposal that the punishment of Sue consist in her doing the dishes for three weeks. Examples (39) to (41) are thus fundamentally different from (38): they are about propositions regarding what is true or what should be done. Syntactically, sentences like (39)-(41), however, give the impression that we are arguing directly about people or types of actions. This impression is superficial; it arises only because the language allows us to omit the contextually presupposed parts of the relevant propositions. For instance, B's answer in (39) specifies the X in a presupposed proposition of the form "X should become chairman".

In some application contexts, the difference between entities and propositions may not matter much. If we know beforehand what an argument in a text is about, or relative to what behavior or domain an entity is evaluated, we may be able to construct the relevant complete propositions relatively easily. For instance, if we track arguments about the candidates of the 2008 Democratic primary, we can assume that most utterances such as "I am for Obama" or "Everybody should get behind Clinton" are, respectively, arguments for nominating Obama or Clinton as the Democratic presidential candidate. However, if at some point we analyze a larger body of texts that cover both the primary phase and the general election campaign, then statements like "I think Clinton has to be our first choice" could also be about the choice of a running mate rather than presidential candidate. These quite different arguments would be conflated if there was no way of identifying the actual argued-for propositions.

According to our annotation scheme, we annotate prominent entities as the targets of arguments or sentiments. The combination of different annotations does, however, allow us to recover the full propositions that are argued for or against.

The proposition associated with arguing attitudes can be reconstructed by combining the target and attitude spans and, if needed, removing the span of the private state expression. Consider examples (42) and (43).

- (42) In his view, Kao **said** [the cross-strait balance of military power] *is not critical to the ROC's national security*.
- (43) She highlighted the situation of the ethnic Uighur Muslim population of Xinjiang, China's furthest west region, where rights groups have **claimed** [a crackdown against separatists] *has been stepped up since the September 11 attacks on the United States*.

In (42), the argued-for proposition is “the cross-strait balance of military power is not critical to the ROC’s national security”. It corresponds to the merged target and arguing-attitude spans. Since the private state span (“said”) does not overlap the other two spans, it does not have to be removed from the merged span. In (43), the argued-for span is “a crackdown against separatists has been stepped up since the September 11 attacks on the United States”. In this case, the arguing-attitude span overlaps the span of the private state expression and that of the target. With the private state span (“claimed”) excluded, the remainder of the attitude span is the argued-for proposition.

5.2.2. The prominence of entities

Assuming now that in many cases the proper target of an subjective expression is a proposition and not an entity, what explains the strong intuitive appeal of entities as targets? While this is a question beyond the scope of this paper, there are two factors that seem to be key parts of the final answer. One is the idea of causation, which we already encountered earlier in the discussion of inferred attitudes (see section 4.3.1.). People or entities that have causal abilities naturally attract attention and are liable to receive blame or credit for events.

A second factor that appears relevant is information structure. Referents that are prominent as pragmatic foci or topics (in the pragmatic sense of e.g. Lambrecht, 1994) are easy to perceive as targets. For instance, if in a discourse an open proposition is set up where only the filler of one particular semantic role is contested, then the candidate entities strongly appear to be the targets that the argument is about. An example of this is (44), where the open proposition is “X should be voted off”.

- (44) Q: Who should be voted off Idol?
 A1: [Anthony] **should** go next. He is just so original, his voice is great, but he doesn't stand out. I love Carrie Underwood.
 A2: [Anwar] **should** go. Nice guy but never, ever in tune.

In example (44), the referents of the noun phrases *Anthony* and *Anwar* are the pragmatic foci of their respective propositions. Note also that the relevant semantic role of “Anthony” and “Anwar” in (44) is not agent but patient, which shows that information structure exerts a distinct influence from that of causation.

A different kind of prominence also makes entities good candidates for target status, namely when they are topics. Consider example (45), which contains a speculation attitude:

- (45) Let's take Apple's iPod as an example. What is going to happen to iPod sales (i.e., iPod Nano, video and the micromini version)? [These Apple products] **will probably** become irrelevant to a big chunk of Apple's target iPhone market upon its release.

In (45), the question posed in the first sentence introduces iPod sales as the topic of concern and “these Apple products” are the topic of the second sentence. The focus of the second sentence is what will happen to the products, that is, the verb phrase “will probably become irrelevant to a big chunk of Apple’s target iPhone market upon its release”.⁴ Our annotation scheme does not distinguish among the various reasons why an expression referring to an entity might be annotated as the target. However, as shown above, it does allow the recovery of the full details of an arguing or sentiment-attitude where that is needed.

6. Related work

The most closely related research focuses on automatic source and target recognition. Bethard et al. (2004) use semantic parsing to identify the opinion holders and targets for verbs that express targets in propositional form. Their coverage is narrower: they work on a subset of direct subjective elements, and they mainly use written data (FrameNet and PropBank sentences). Choi et al. (2006) investigate the joint extraction of opinions and sources from the MPQA corpus. Their results show that using PropBank argument role labeling is beneficial for source extraction. Kim & Hovy (2006) work on source and target recognition using FrameNet semantic role labeling and manual mappings from semantic to opinion roles. Their evaluation on FrameNet data shows results similar to Choi et al. for sources and newly demonstrates the value of role labeling for target recognition.

Other research tackles source and target recognition as part of building complete review mining systems (Popescu and Etzioni, 2005; Yi and Niblack, 2005; Bloom et al., 2007). Yi & Niblack’s sentiment pattern database and Bloom et al.’s linkage specifications provide mappings of argument to opinion roles similar to the mappings in Kim & Hovy’s work. However, when linking up targets and opinion expressions, Bloom et al. and Popescu & Etzioni also use domain knowledge about product features that are likely targets of evaluation. As our work is more general, we do not use feature taxonomies.

7. Ideas for further research

We have argued that source and target recognition poses challenges that go beyond the capabilities of automatic semantic role labeling. The issues that we have discussed suggest some avenues for future research. One is to study the co-reference relations between unexpressed semantic arguments with definite reference and the overt phrases they co-refer with (cf. 25).

⁴Notice that in (45) topic and focus line up the opposite way from (44): in the former the subject is the topic and the predicate the focus, whereas in the latter the subject is the focus and the predicate is the topic.

More research on attribution is needed to capture the nesting of sources. Such work might build on the existing annotations of the Penn Discourse Treebank, which already handles immediate sources.

Other research ideas concern lexical resources such as PropBank and FrameNet. For one, it would be desirable to have generally available mappings between semantic roles and (possibly multiple sets of) opinion roles for a large number of predicates. Further, it would be interesting to explore how these resources might support the inference of private states and their sources and targets. We would, for instance, like to be able to look up which arguments of a predicate are negatively or positively affected, and which participants are causally responsible.

8. Conclusion

We have argued here that the identification of sources and targets of opinions cannot be solved completely by semantic role labeling systems alone, although that technology may go a long way in certain genres. In particular, role labeling systems face problems with attitudes that are connected to their sources indirectly via attribution; in cases where sources or targets of direct-subjective expressions are realized as zero forms; and when the targets of inferred attitudes need to be found. Further, we have shown that when dealing with arguing-attitudes, labeling prominent entities as targets by themselves is not sufficient unless complete argued-for propositions can be reconstructed either from the text or from extra-linguistic knowledge about the task or domain at hand. Throughout the discussion, we have presented evidence that the recognition of targets often cannot be done deterministically. Rather, attention needs to be paid to the specific context.

9. References

- A. Andreevskaia and S. Bergler. 2006. Semantic Tag Extraction from WordNet Glosses. In *Proceedings of LREC-2006*.
- C. Baker, C. Fillmore, and J. Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of the 17th international conference on Computational linguistics*.
- S. Bethard, H. Yu, A. Thornton, V. Hatzivassiloglou, and D. Jurafsky. 2004. Automatic Extraction of Opinion Propositions and their Holders. In *Proceedings of AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*.
- K. Bloom, N. Garg, and S. Argamon. 2007. Extracting Appraisal Expressions. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*.
- Y. Choi, E. Breck, and C. Cardie. 2006. Joint Extraction of Entities and Relations for Opinion Recognition. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*.
- A. Esuli and F. Sebastiani. 2006. Determining Term Subjectivity and Term Orientation for Opinion Mining. In *Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- V. Hatzivassiloglou and K. McKeown. 1997. Predicting the Semantic Orientation of Adjectives. In *Proceedings of the Thirty-Fifth Annual Meeting of the Association for Computational Linguistics and Eighth Conference of the European Chapter of the Association for Computational Linguistics*.
- M. Hu and B. Liu. 2004. Mining Opinion Features in Customer Reviews. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-2004)*.
- S. Kim and E. Hovy. 2006. Extracting Opinions, Opinion Holders, and Topics Expressed in Online News Media Text. In *ACL Workshop on Sentiment and Subjectivity in Text*.
- K. Lambrecht. 1994. *Information structure and sentence form*. Cambridge University Press, New York, N.Y.
- M. Palmer, P. Kingsbury, and D. Gildea. 2005. The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics*, 31.
- B. Pang and L. Lee. 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. In *Proceedings of the ACL 2004, Barcelona, Spain, Main Volume*.
- A. Popescu and O. Etzioni. 2005. Extracting Product Features and Opinions from Reviews. In *HLT/EMNLP 2005*.
- R. Prasad, N. Dinesh, A. Lee, A. Joshi, and B. Weber. 2006. Annotating Attribution in the Penn Discourse TreeBank. In *Proceedings of the COLING/ACL Workshop on Sentiment and Subjectivity in Text*.
- W.V. Quine. 1956. Quantifiers and propositional attitudes. *Journal of Philosophy*, 53.
- R. Quirk, S. Greenbaum, G. Leech, and J. Svartvik. 1985. *A Comprehensive Grammar of the English Language*. Longman, New York.
- E. Riloff and J. Wiebe. 2003. Learning Extraction Patterns for Subjective Expressions. In *EMNLP 2003*.
- H. Takamura, T. Inui, and M. Okumura. 2007. Extracting Semantic Orientations of Phrases from Dictionary. In *Proceedings of NAACL/HLT 2007*.
- P. Turney. 2002. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 417–424, Philadelphia, Pennsylvania.
- J. Wiebe, T. Wilson, and C. Cardie. 2005. Annotating Expressions of Opinions and Emotions in Language. *Language Resources and Evaluation*, 39.
- T. Wilson and J. Wiebe. 2005. Annotating attributions and private states. In *Proceedings of the ACL Workshop on Frontiers in Corpus Annotation II: Pie in the Sky*, Ann Arbor, Michigan.
- T. Wilson. 2008. *Fine-Grained Subjectivity And Sentiment Analysis: Recognizing The Intensity, Polarity, And Attitudes Of Private States*. Ph.D. thesis, University of Pittsburgh.
- J. Yi and W. Niblack. 2005. Sentiment Mining in Web-Fountain. In *Proceedings of the 21st International Conference on Data Engineering*, pages 1073–1083.