

Instructions for Annotating Topic of Opinions

Introduction

Our ultimate goal is to create algorithms and methods that can automatically extract opinions from text. For the purpose of the discussion, we use the term *opinion* to refer to opinions, beliefs, emotions, sentiment, and other private states expressed in text. Private state is a general term used to refer to mental and emotional states that cannot be directly observed or verified (Quirk et al. 1985).

In order to be able to automatically extract opinions from text we will rely on Machine Learning (ML) techniques. Both for the development and evaluation of these techniques, however, we need a corpus of documents manually annotated with information about opinions. Fortunately, such a corpus already exists – the MPQA corpus contains documents manually annotated with information about opinions.

Documents in the MPQA corpus are manually annotated by designating all expressions of opinions at the fine-grained level of individual expressions of opinions. The manual annotations of opinions include a number of attributes of the opinions such as the *source* of the opinion, the *opinion trigger* or the words that signal the expression of opinion, the *polarity* or favorability of the opinion, the *strength* of the opinion, as well as the *target* of the opinion. However, in the current version of the MPQA corpus, the target attribute of opinion is included for very few of the actual annotations. The absence of more comprehensive marking of the target attributes is due to the challenging nature of the target annotation task in its original definition.

The target attribute, however, is an integral part of the each expression of opinion. As such, it is desired that automatic extractors of opinions are able to extract the target of each opinion. To facilitate the creation and evaluation of automatic targets, the purpose of this annotation task is to add the target attribute to the opinion annotations in the MPQA corpus.

In order to avoid shortcomings of previous approaches to target annotations we chose a different definition (and name) for the target annotation task. For the remainder of the document, the target of an opinion will be referred to as the *topic*. As a term topic carries more or less the same meaning, but we prefer it to target because it is a more general and vague term. Target generally carries the connotation that it refers to a specific well-defined concrete entity.

We define *topic* of an opinion as the (physical or abstract) entity, action, event, artifact, ideology, matter, etc. that is targeted by an opinion. Here are several examples of topics of opinions.

(1) President Chen Shui-bian has on many occasions [expressed goodwill](#) to **mainland China**. (Topic: mainland China)

(2) The IHRC said in a statement that the international community has formulated numerous documents to honor human rights after 50 years of bitter experience and the heavy losses that the humanity has suffered. (Topic: human rights)

(3) The IHRC said in a statement that the international community has formulated numerous documents to honor human rights after 50 years of bitter experience and the heavy losses that the humanity has suffered. (Topic: period of force and violence in international relations. **Note that here the topic is clear only from the context.**)

In all examples in this document, opinion attributes are marked as follows:

- sources of opinions are underlined and in bold.
- opinion triggers (words that signal the presence of opinion) are also underlined and shown in either blue for positive opinions, red for negative opinions or gray for neutral or non-sentiment carrying opinions (more detail about the non-sentiment carrying opinions to come).
- spans that signal the topic are highlighted in yellow.

The last example – example (3) -- hints at why topic annotation is difficult. To circumvent some of the difficulty of topic annotation, we introduce and use the notion of *topic coreference* of opinions. We say that two opinions are *topic coreferent* if they share the same general topic. For example, the opinion from (3) is coreferent with the following opinion in the same document:

(4) Tehran-based Islamic Human Rights Commission (IHRC) on Sunday expressed concern about return of the period in which force and weapon had the last say in international relations.

Armed with the notion of topic coreference, the goal of our annotation task is to group (cluster) together those opinions that concern the same topic (are topic coreferent) and label every group (cluster) with the topic of the cluster. Additionally, we would like to mark the text spans that signal the expression of the topic (corresponding to the yellow highlights in the examples).

In the next section we give a brief background of the existing opinion information, followed by sections containing general instructions of what opinions should be considered topic coreferent and how to form labels for clusters of topic coreferent opinions.

Background

As mentioned in the introduction, documents in the MPQA corpus are annotated with expressions of opinions. We will augment the existing opinion annotations with information about the topic of opinions. For this purpose we will use a special annotation tool to display and augment existing opinion annotations. In this section, we give a really brief overview of the important parts of the existing opinion annotations.

In general opinions in language can be expressed either directly – e.g. “John hates Mary” – or through the choice of style and words in the language used – e.g. “John whined about school all the way to the cafeteria” (the choice of the verb “whine” rather than the more neutral “complain” signals the author’s negative opinion of John). Documents in the MPQA corpus are annotated with both types of opinions, to which we will refer as *direct opinion* and *expressive subjectivity* respectively.

As mentioned previously opinions have a number of attributes such as opinion trigger, source, polarity, and strength. In the examples in this document we show all of these attributes using underlining, highlighting, and color following a template that is similar to the one used by the topic annotation tool. It is worth noting, however, that sometimes not all of the opinion’s attributes are present in the context. For example, for some opinions the source attribute is not explicitly mentioned, but rather inferred from the text (e.g. the source of the opinion “John whined about school all the way to the cafeteria” is the writer, who is not explicitly mentioned in the sentence). Additionally, for most of the opinions the polarity attribute is missing. For instance, in the sentence:

(5) John predicted that it will take a defensive mistake for one of the two teams to score in the game.

Clearly, John’s opinion (or private state) is expressed. However, this sentence does not express any form of positive or negative sentiment, but rather his belief toward the subject matter. In that respect, our use of the term opinion can be somewhat misleading as it typically implies the expression of some form of sentiment. Remember, however, that we use the term *opinion* (arguably quite loosely) to refer to *private state* or a mental state that cannot be directly observed or verified. In this definition of opinion, the private state does not have to express sentiment to be considered opinion.

The significance of the missing polarity should become clear in the next section, as we have found out that topics of opinions that express some form of sentiment are easier to judge as compared to opinions that express just any belief. In the next section, we look into more detail of how opinion annotations should be performed.

Annotation Instructions

Topic Coreference

Recall from the definition that we consider two opinions to be topic coreferent if the general topic of the opinions is the same. The topic might be a concrete person or object such as:

(6) I hate John.

(7) Sue is very found of John.

Or it could be an abstract concept such as event, idea, etc.:

(8) CNN has discussed the latest developments in the Israeli-Palestine dialogue.

(9) It is quite important that the governments of Israel and Palestine resume communication.

In the general case, topic of opinions will not be as clear-cut and easy to judge as the opinions above. There are at least a few issues, of which we are currently aware and which will make the process more difficult:

Multiple opinions in a sentence

In many cases, the sentence that you will be annotating contains more than one opinion. In the following sentence, for example:

(10) In her view, Tsai said, both sides have been endeavoring to prevent the results of Taiwan's recent elections from affecting the stability of cross-strait relations.

There are multiple opinions marked shown by their opinion triggers. Depending on which of the opinions in the sentence are being annotated, the topics of the opinions differ:

(10a) In her view, Tsai said, both sides have been endeavoring to prevent the results of Taiwan's recent elections from affecting the stability of cross-strait relations. (Topic: cross-strait relations)

(10b) In her view, Tsai said, both sides have been endeavoring to prevent the results of Taiwan's recent elections from affecting the stability of cross-strait relations. (Topic: both sides (Taiwan and China))

When annotating, you will need to carefully understand which of the multiple opinions are annotated and assign the appropriate topic.

Non-sentiment vs. sentiment opinions

As already mentioned and perhaps hinted in the last example, often topics of opinions that carry sentiment are easier to judge than general, non-sentimental opinions. This is due to the fact that sentiment is typically clearly stated and directed toward a specific entity or event. Non-sentiment opinions can be generally vague and concern multiple entities, events, or ideas (multiple topics of opinions are discussed in more detail in the next subsection). When you judge the topic of non-sentimental opinions, in many cases you will have to carefully read the span of the opinion and make a conjecture about the topic, possibly based on the surrounding context. For example, the following opinion can be judged only from the context:

(11) John has repeatedly stated that although the defense needs some upgrades, unless the coaching staff spends more resources on the offensive skill positions,

it will be another disappointing season. (Topic: the Philadelphia Eagles, but not mentioned anywhere in the sentence)

The significance of judging the problem in context is intimately connected with the problem of having more than one topic in a single opinion. For instance, the above example could be considered topic coreferent with opinions regarding offensive skill positions, if it was a part of say a paragraph on offensive skill positions around the NFL. For example if the next sentence was:

(12) Sue agrees that her beloved Broncos due most of their recent success on the depth that they possess on **the offensive skill positions**.

Then the two opinions would be considered coreferent in a cluster concerned with offensive skill positions. This problem of more than one possible topic is discussed in more detail in the next subsection.

The topic annotation software attempts to facilitate the annotation task by distinguishing the opinions that were judged to carry sentiment from the non-sentimental opinions. Opinions sentiments are color-coded in the software with gray highlighting for the opinion words signaling non-sentimental opinions. Additionally, the software distinguishes between direct opinions and expressive-subjectivity, since the latter more often than not carry a sentiment.

Multiple topics per opinion

Often one opinion can be considered to be concerning multiple topics. In many documents in the corpus, the exact topic of opinion may be hard to judge and depend on the context. Consider the following examples:

(13) Tsai Ing-wen said Tuesday she foresees the possibility of **the two sides of the Taiwan Strait resume dialogue next year**.

...

(14) "It all depends on how mainland China interprets President Chen's latest remarks on cross-strait relations and how the two sides cultivate an environment favorable for resumption of their long-stalled dialogue," Tsai explained.

The question is, are the opinions of (13) and (14) topic coreferent. The topic of the first sentence is expressed quite clearly. The opinion in the second sentence, however, concerns more than a single entity or event and can be judged only from the context. Tsai's statement in the sentence discusses the resumption of the dialogue (as does the opinion in the first sentence), but also talks about President Chen's remarks and the cultivation of environment of the two sides.

When annotating opinions with more than one potential topic, we consider the topic of the opinion to be the part of the opinion which is being emphasized. To make this judgment, we have to ask ourselves the question what is the purpose of the information

that is being conveyed in the opinion. For example, consider the original context of the sentence. It was after the following sentence:

(15) Tsai said that there should be opportunities for the two sides to resume talks.

In this context, the opinion concerns the resumption of dialogue and is coreferent with statement. In the context, the information that the statement gives is predominantly concerning the talks and Chen's statements and the environment can be considered secondary issues.

Let's assume that the preceding sentence was slightly different:

(16) Tsai concurred with Zhen on the importance of the remarks.

In this context, the topic of the opinion is Chen's remark as the statement serves the purpose to elaborate on the remark. In this case the opinions from and should not be considered coreferent.

Topic hierarchies

Yet another problem with the topic annotations is that opinions might be concerning different aspects or parts of the same topic. Remember the previous example:

(17) Tsai Ing-wen said Tuesday she foresees the possibility of the two sides of the Taiwan Strait resume dialogue next year.

The topic of the opinion here is the resumption of the dialogue by the two sides of the Taiwan Strait next year. Notice, however, that this topic can have different level of specificity:

- The resumption of the dialogue next year.
- The resumption of the dialogue.
- The dialogue.

The task of our annotation is to judge whether opinions are topic coreferent. So the question is, given another opinion, the topic of which may differ in its level of specificity, are the two opinions topic coreferent. Say we were given the sentence:

(18) In Tsai's opinion, the dialogue in the Taiwan Strait is very important.

We have to judge whether it is topic coreferent with example. For the purpose of the annotation, we will assume the following definition concerning the specificity of opinion topics: Two opinions are topic coreferent if they discuss different part or aspects of the same general topic; the label of the topic is the most general common topic.

Following this definition, the opinions from our examples should be considered coreferent, with the label for the topic being "the Taiwan Strait dialogue." Note that the topic label is "the least common divisor" of the topics. For example, if the two topics

were the resumption of the talks and the history of the talks the label of the opinion cluster would be “the talks” although it might not appear as a separate topic of an opinion.

Note that the above definition talks about different parts or aspects of the same topic. In the example we talk about the resumption and the history of the talks (or the dialogue). Note that these opinions differ only on the level of specificity, but not in the general topic. If, on the other hand, we had an opinion concerning Iraq and another one concerning the war in Iraq, the two should not be considered coreferent although they both discuss Iraq. The difference is that the war in Iraq is not a mere specification or attribute of Iraq, but rather a separate concept.

Topic Spans

In addition to the topic coreference annotations, our ML approaches can benefit from knowing the part of text which provides the evidence for the decision of the topic. That is, we would like to manually add the span of text which indicates the expression of the opinion topic (the text with yellow highlights in the examples). For this purpose, the task should be fairly straightforward – add to each annotation the part of the opinion sentence which indicated the topic of the opinion.

When performing this annotation, please select the minimal span of text which addresses the topic. Recall the example:

(17a) Tsai Ing-wen said Tuesday she foresees the possibility of the two sides of the Taiwan Strait resume dialogue next year. (Topic: dialogue in the Taiwan Strait)

Here the topic was the dialogue between the sides of the Taiwan Strait and this is exactly the part that we will annotate. If on the other hand, the topic was “the resumption of the dialogue in the Taiwan Strait next year,” we would annotate the example as follows:

(17b) Tsai Ing-wen said Tuesday she foresees the possibility of the two sides of the Taiwan Strait resume dialogue next year. (Topic: resumption of the dialogue in the Taiwan Strait next year)

In addition, note that in some cases the expression of the topic might not be in the same sentence as the opinion. Recall example (11):

(11) John has repeatedly stated that although the defense needs some upgrades, unless the coaching staff spends more resources on the offensive skill positions, it will be another disappointing season. (Topic: the Philadelphia Eagles, but not mentioned anywhere in the sentence)

In this example, the topic is not explicitly mentioned in the opinion sentence. If you encounter such an example, please annotate the part of the document outside of the sentence which mentions the topic and on which you based your topic coreference decisions.

Topic labels

So far we have mentioned, but have not discussed the labels of topic clusters. Once opinions are separated into clusters, we face the task of assigning a label representing the opinions in the cluster. In the previous section we mention that the label should be the “least common denominator” of the opinion topics in the cluster; In addition, we would like the label of the opinion to be a text segment that occurs somewhere in the text¹. Furthermore, we would like for the text segment to come from one of the parts that you have annotated as spans expressing the opinions (the text with yellow highlights).

Thus, the task of labeling can be viewed as the task of finding the most general and representative span of text that describes the topic for the cluster from the text spans which express the topics of opinions. Going back to the previous examples:

(19) Tsai Ing-wen said Tuesday she foresees the possibility of **the two sides of the Taiwan Strait resume dialogue next year**.

(20) In Tsai’s opinion, **the dialogue in the Taiwan Strait** is very important.

A good label for this cluster would be “dialogue in the Taiwan Strait”, coming from the part of the second sentence highlighted in yellow. If for example, the title of the document from which these opinions came was “Taiwan Strait talks”, this could also be a good characterization of the topic of the cluster, but we would not desire this label since it did not come from a span expressing the topic of any of the opinions in the cluster.

In some cases, the label of the opinion cluster may not be expressed in any of the opinion sentences. In this case it is acceptable to create a topic label that is not found in any of the documents text.

¹ The need for labels to be actual text segments is motivated by the fact that we will attempt to recover topics and labels using machine learning techniques. For machine learning techniques it is generally harder to infer text that did not occur in the same form anywhere in the document.