

Supplementary Notes to *Opinion Holder and Target Extraction based on the Induction of Verbal Categories (CoNLL 2015)*

Michael Wiegand Josef Ruppenhofer

June 21, 2015

On the usage of this document: *This document provides further information regarding the annotation of the resources introduced this paper. It mostly concerns (high-level) conceptual definitions/annotation guidelines and their motivation. This content could not be included in the actual paper because of size restrictions. Please note that each resource introduced in this paper also comes with a README file which provides more technical information regarding its format.*

1 Introduction

Sentiment analysis is a very actively researched area of computational linguistics and natural language processing. While the detection of so-called subjective expressions such as *good*, *lovely*, *disgust*, *love*, *hateful*, *lousy* has received a lot of attention, so far there has been less focus and fewer resources available for the extraction of opinion holders and targets, which we jointly refer to as opinion roles.

Opinion holders are the entities whose point of view is represented by the subjective expression and opinion targets are what the opinions are about. Consider example 1: the holder of the positive sentiment that is expressed is *Peter*, the target towards which this positive sentiment is directed is *it*.

- (1) Speaking on BBC Breakfast, Abbey said: “[Peter *Holder*] **loves** [it *Target*]!”

Note that from the text it is clear that we learn about Peter’s sentiment only by way of Abbey and whoever composed the sentence in which Abbey’s speech is cited. However, for our purposes of lexical analysis, we are interested only in the *immediate* source of a subjective expression and not in any other so-called *nested sources* (cf. Wiebe et al. [2005]).

- (2) [“The Last Airbender” movie *Target*] **sucked!**

We focus exclusively on *verbal* subjective expressions because they are central to the expression of opinion roles: among those subjective expressions that

take both the opinion holder and the opinion target as syntactic dependents, many more are verbs than adjectives.

In this contribution, we introduce two lexicons, a fine-grained one and a more coarse-grained version (*we produce two lexicons for English, however only a coarse-grained version for German, cf. Section 4*), containing information on how subjective expressions realize opinion holders and opinion targets in the syntax.¹ The lexicons do *not* contain any information about the polarity and intensity of subjective expressions.

A **high-level summary of the features of the lexicons** is the following:

- The verbal entries are assigned to one of three types, which are defined by how the opinion holder relates to the semantic roles of the verb (cf. Section 2.1).
- The scope of the lexicon is to support local, inherent sentiment only; it is **not** meant to support opinion inferences (cf. Section 2.2).
- The same verb may have multiple entries because several opinions may be associated with it (cf. Section 2.3).
- The entries are on the lemma-level rather than at the level of word senses (cf. Section 2.4).

Please see Section 2 and its subsections for details.

Accompanying, the lexicons are annotations of textual instances of the subjective expressions contained in the lexicon. These resources are intended to support further theoretical and applied research on opinion role extraction.

2 Lexicon

*We introduce two types of lexicons in this paper, a **fine-grained** and a **coarse-grained** lexicon. Since the coarse-grained lexicon is (automatically) derived from the fine-grained lexicon, our remarks in this section mostly concern the fine-grained lexicon.*

The lexicons we constructed record how opinion holders and opinion targets map to the semantic roles of verbal predicates. If we consider, for instance, a verb such as *abhor* in example 3, we notice that its opinion holder corresponds to its semantic role agent and its target to the semantic role patient.

- (3) They are people [who *Holder*] **abhor** [violence *Target*], people who seek peace.

Consider now the verb *affront*, illustrated in example 4.

- (4) [These flaws *Target*] **affronted** [the liberal imagination *Holder*].

With *affront* in example 4, the situation is the opposite from *abhor*: the opinion holder corresponds to the verb's semantic role patient and the target to the verb's semantic role agent.

Finally, consider the verb *adulterate*, exemplified in example 5.

¹We use the terms opinion holder and opinion target here. Another commonly used term for opinion holder in the literature is source, while what we call targets are often referred to as opinion topics.

- (5) It is genius, and not the want of it, [that *Target*] **adulterates** philosophy, and fills it with error and false theory.

The use of the verb *adulterate* expresses an opinion by an unexpressed source – one that is either a nested source or the speaker/writer of the relevant utterance, as the case may be – about the cause of the adulteration.

2.1 Views

We chose the three verbs discussed above (i.e. *abhor*, *affront*, *adulterate*) as examples because they represent the most common types of verbal predicates in terms of how holders and targets map to high-level semantic roles. We formulate this typology in terms of views: who is the source of the opinion.

agent view *abhor* represents the case where the agent role corresponds to the opinion holder.

patient view *affront* represents the case where the patient role corresponds to the opinion holder.

speaker view *adulterate* represents the case where a nested source, rather than any semantic role of the predicate, corresponds to the opinion holder. (Maks and Vossen [2012] call this type *speaker subjectivity* which is why we name the corresponding view: *speaker view*. In-context annotation of opinion verbs with these views may not be as straightforward as the annotation of opinion verbs having an agent view or patient view. Section 5.3 describes our methodology.)

2.2 Scope: Sentiment Inherent in Subjective Expressions

Implicit in our project of mapping opinion holders to semantic roles of individual predicates is the idea that in the analysis of opinions we need to distinguish two things: Opinions that are conveyed by inherently subjective predicates, and inferred opinions that are calculated using, for instance, knowledge about attitudes towards participants as well as knowledge about positive or negative effects of events on participants (cf. the work of Deng and Wiebe [2014], Wiebe and Deng [2014]). Consider the following concrete example:

- (6) People seem to relish that he lost a Super Bowl.

While we might be interested in extracting “people’s” opinion about “him”, as we understand the task, we need to do so in two steps. First, we recognize the opinion that is lexically inherent in *relish* that people have a positive attitude towards the overall loss event. Second, as proposed in several recent publications by Wiebe and colleagues, we use inference rules and knowledge about the negative effect on the agent of *lose* to *infer* that “people” likely have a negative attitude towards “him”.

The purpose of the lexicon and the annotations we provide is to aid in the first of these two steps, the sentiment step, only. Providing information about effects and inferences rules is **not** a task we address here. At the same time, we need to of course admit that delineating what is lexically inherent meaning versus pragmatic inference is not always easy. Just as ‘yesterday’s

syntax is today’s morphology’, so is ‘yesterday’s pragmatics today’s (lexical) semantics’. Accordingly, if we have included opinions in our entries that other people consider due to inference, then it is due to different perceptions only. We deliberately chose the set of verbs to originate from the *Subjectivity Lexicon* – one of the most frequently used sentiment lexicons – since those entries are commonly thought to convey *inherent* sentiment.

2.3 Multiple Views

It has long been known that there are predicates that express more than one opinion. A classical example is *brag*. It conveys a positive sentiment by its agent (holder) about its patient (target) but it also conveys negative sentiment towards the agent by an external source. Accordingly, the verb *brag* has two entries in our lexicon.

2.4 Lemma-based Analysis

It is clear to us that ultimately one wants to use word-sense based resources for opinion analysis. However, like others, we have to deal with limitations of present resources and tools and therefore decided to provide lemma-level information.

2.5 Annotation with Online Dictionaries

For the annotation of the fine-grained lexicon, the annotators largely relied on information provided by online dictionaries, such as *Macmillan Dictionary*.². Such dictionaries provide two useful types of information:

- word definitions
- example sentences

Word definitions may help to identify which *views* some verb actually evokes. We found this particularly helpful when it comes to decide whether a speaker view is evoked. For example, the verb *complain* was found to exclusively convey an agent view, while *carp* was found to convey both an agent and a speaker view. This can be read off from the corresponding lexicon definitions (from *Macmillan Dictionary*):

- complain: *to say that you are not satisfied with something*
- carp: *to complain a lot, especially about things that are not important*

carp does not only convey an agent view similar to *complain*. It additionally conveys that the speaker thinks that the complaint is disproportionate (*things that are not important*). While the above definitions for *complain* and *carp* allow a clear derivation of opinion views, there may, of course, be cases in which a distinction is more difficult. By allowing our annotators to consult more than one online dictionary, these situations should, however, only rarely occur.

Example sentences are helpful in acquiring information concerning the argument realization of opinion holders and targets. In contrast to the semantic-role

²www.macmillandictionary.com

labeling representation of *PropBank* [Kingsbury and Palmer, 2002], we decided to be more specific with regard to the type of patient. That is, rather than defining a patient as *A1* (similar to PropBank), we explicitly state which type of phrase the patient is realized. For example, the patient of *abscond* comes as a prepositional phrase headed by *with*, while for *ignore* it is a direct object³. Such information can be read off from example sentences.

In order to reduce inaccuracies caused by our lemma-based annotation, we asked annotators to focus on *prototypical* senses of a word. This means that senses ascribed to only very specialized contexts (e.g. archaic usage) should be skipped.

3 Predominant Views and the Coarse-Grained Lexicon

While the fine-grained lexicon lists all possible views that are evoked by a particular opinion verb, the coarse grained lexicon just lists for each opinion verb a predominant view. The assignment of opinion holder and opinion target is derived from the predominant view:

- **agent view:** holder=agent; target=patient
- **patient view:** holder=patient; target=agent
- **speaker view:** holder=*N/A*; target= agent

We automatically convert the fine-grained lexicon into a coarse-grained lexicon with the help of the following rules:

- speaker view + agent view + patient view → speaker view
- speaker view + agent view → speaker view
- speaker view + patient view → speaker view
- patient view + agent view → patient view

The result of this automatic conversion was manually corrected, however, it was only necessary to alter the predominant view for 3% of the opinion verbs. This shows that the automatic conversion produces the right categorization for the vast majority.

The conversion rules, themselves, are motivated in the paper. They mainly rely on linguistic observation. However, our in-context evaluation shows that the proposed coarse-grained representation (produced by this conversion) produces almost as good results as the fine-grained representation.

³As mentioned in Section 2.1, we assume that passive sentences are normalized to their active-voice counterpart. The *subject* in a passive-voice construction would correspond to a *direct object* in an active-voice construction.

4 Some Notes on the German Lexicon

For this paper, we **only** reproduced a German version of the English coarse-grained lexicon, i.e. a lexicon assigning the predominant view (agent view, patient view, or speaker view) to each opinion verb. The basis of this lexicon are the opinion verbs contained in the German sentiment lexicon of the PolArt-system [Klenner et al., 2009]. It comprises 1416 (verb) entries. The lexicon was manually annotated. Similar to the English lexicon, the annotators used verb definitions and example sentences of online dictionaries, such as *Duden*⁴, as a guide for their annotation.

We only produced a coarse-grained lexicon in order to demonstrate the applicability of the three opinion views to another language. Furthermore, we wanted to show that induction methods effective for English are equally effective on German data.

5 In-Context Annotation

In order to validate our induction approach and also the manually-compiled opinion-role lexicons, we created a corpus comprising the opinion verbs under investigation in context. The sentences were sampled from the North American News Text Corpus (LDC95T21).

A **high-level summary of the features of the in-context annotation** is the following:

- Following the annotation scheme of the lexicon (cf. Section 2) we annotate local, inherent sentiment only; it is **not** meant to support opinion inferences.
- The annotation is carried out on a parsed text (cf. Section 5.1).
- We only annotate mentions of the opinion verbs (cf. Section 5.2).
- Opinion holders that are not realized as dependents of our opinion verbs are annotated but specially marked (cf. Section 5.3).

5.1 Input to Annotation

The input to the annotation are sentences that have been parsed with a constituency parser. (We used the Berkeley parser [Petrov and Klein, 2007].) We decided against an annotation on plain text since both opinion holders and opinion targets coincide to syntactic constituents. While annotators may be able to identify constituents spanning over a few words from plain text as it is the case with opinion holders (which are typically noun phrases), they may be less able to do so when it comes to large constituents. Opinion targets fall among the latter category. Since opinion targets can cover entire propositions, their span is typically significantly larger than those of opinion holders. We hope that by giving annotators pre-parsed text, we will come up with a gold standard that is also largely compatible with automatic analyses w.r.t. the spans of opinion holders and opinion targets.

⁴www.duden.de

For annotation, the annotators are only provided as context the sentence in which an opinion verb occurs. There is no further context, such as preceding or following sentences, available to them.

5.2 Ambiguous Opinion Verbs

Annotators were asked to annotate mentions of the pre-specified opinion verbs only. This means that mentions of opinion words other than those opinion verbs were ignored.

Not every mention of our opinion verbs conveys subjectivity (some opinion verbs are known to be ambiguous [Akkaya et al., 2009]). Those cases were retained in our corpus but do not contain any annotation. As a consequence, classifiers may mistakenly interpret these mentions as opinions. This will decrease the performance of all classifiers. However, we deliberately retain these mentions in our corpus since there do not yet exist any robust methods for (subjectivity) word-sense disambiguation. We pursue a realistic scenario: we only have a prior notion about which verbs can convey subjectivity and which can not. We are not able to predict in which context an ambiguous opinion verb conveys subjectivity and in which context it does not.

5.3 Opinion Holders not Realized as Dependent of Opinion Verbs

Mentions of an opinion verb conveying a speaker view often come without any expression of an opinion holder and when this is so, the speaker of the utterance is usually the holder. Consider example 7 below, where there is no opinion holder. Still, it is possible that mentions of SP-verbs are accompanied by mentions of the relevant opinion holder somewhere in the discourse context. However, such mentions are never syntactic arguments of the SP-verb.

For instance, in example 8, the nested source *He* of the main clause predicate *believe* is the holder of the opinion expressed by *suck*. (Note that *He* is not the speaker of the utterance.) Similarly, in example 9 the preposition *according to* introduces the relevant opinion holder for *stink* in an optional prepositional phrase.

Following the annotation scheme of the MPQA corpus, we annotated explicit mentions of opinion holders in the context of SP-verbs. However, we do specially mark them with a flag indicating that they are no syntactic dependents of the opinion verb.⁵

- (7) ["The Last Airbender" movie *Target*] **sucked!**
- (8) [He *Holder*] believes that ["The Last Airbender" movie *Target*] **sucked!**
- (9) According to [Tanya *Holder*] [it *Target*] "just **stank**".

⁵In all our experiments, we retained those mentions of opinion holders, even though this means that our alignment rules between syntactic arguments and opinion roles cannot reach those entities. We retain these cases since we seek to have a realistic gold standard.

References

- Cem Akkaya, Janyce Wiebe, and Rada Mihalcea. Subjectivity Word Sense Disambiguation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 190–199, Singapore, 2009.
- Lingjia Deng and Janyce Wiebe. Sentiment Propagation via Implicature Constraints. In *Proceedings of the Conference on European Chapter of the Association for Computational Linguistics (EACL)*, pages 377–385, Gothenburg, Sweden, 2014.
- Paul Kingsbury and Martha Palmer. From TreeBank to PropBank. In *Proceedings of the Conference on Language Resources and Evaluation (LREC)*, pages 1989–1993, Las Palmas, Spain, 2002.
- Manfred Klenner, Angela Fahrni, and Stefanos Petrakis. PolArt: A Robust Tool for Sentiment Analysis. In *Proceedings of the Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 235–238, Odense, Denmark, 2009.
- Isa Maks and Piek Vossen. A lexicon model for deep sentiment analysis and opinion mining applications. *Decision Support Systems*, 53:680–688, 2012.
- Slav Petrov and Dan Klein. Improved inference for unlexicalized parsing. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the ACL (HLT/NAACL)*, pages 404–411, Rochester, NY, USA, 2007.
- Janyce Wiebe and Lingjia Deng. An account of opinion implicatures. *CoRR*, abs/1404.6491, 2014. URL <http://arxiv.org/abs/1404.6491>.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. Annotating Expressions of Opinions and Emotions in Language. *Language Resources and Evaluation*, 39(2/3):164–210, 2005.