

Verb-centered Sentiment Inference with Description Logics

Manfred Klenner
Computational Linguistics
University of Zurich
Switzerland
klenner@cl.uzh.ch

Abstract

We introduce description logics as a means to carry out sentiment inferences triggered by some verbs on their semantic roles. Verbs might impose polar effects on some roles, but also have polar expectations on other roles. For instance, an entity inherits a positive polarity just by being actor of some verb (“she succeeds”). More complicated scenarios arise if we take subclause embeddings, negation and polarity conflicts into consideration. Polarity propagation and effect inversion need to be coped with, then. We have implemented a prototype in OWL covering a substantial subset of our verb lexicon covering about 140 German verbs.

1 Introduction

Verbs and their role in sentiment analysis have raised some interest only recently, cf. Neviarouskaya et al. (2009), Reschke and Anand (2011), Maks and Vossen (2012), Hollenstein et al. (2014), Deng and Wiebe (2014). More or less common to these approaches is the notion of sentiment inferences triggered by verbs. For instance, given the sentence “Greece blames EU and IMF for ‘obstacles’ in talks”, we expect that the PP (“for”) must be something negative and we understand that the direct object (“EU”) receives a negative effect and that the subject (“Greece”) is the source of this. The aforementioned approaches differ much in the details, but they all strive to cope with this kind of implicit information.

More complicated cases such as “U.N. Refugee Chief criticizes Europe for not rescuing migrants” raise the need to cope with subclause embedding and negation. Even polarity conflicts might arise as in “Palestinian students admire terrorist Dalal Mughrabi” where a negative direct object seems to produce an odd scenario. However,

what about “I hugely admire refugees” where we have a negative direct object as well (“refugee” is negative in the way “ill” is), but where no polarity clash seems to be arising. We argue that fine-grained distinctions are needed to distinguish these cases, namely to distinguish factual from emotional and moral polarities. In this article, we simplify matters and introduce a concept called *sympathy entity* which covers factual negative entities such as refugee etc.

In order to draw sentiment inferences, rules are needed and their application must not interfere with each other nor should the set of rules be inconsistent (which is a problem for complex rule-based systems). We believe description logics (Baader, 2009) is well suited as a framework to model such an inference task. The big advantage is that we do not need to care about the order in which rules are applied; also consistency checks are part of the machinery. Our model is a competence model that - at least in principle - should be able to properly cope with arbitrarily complex embeddings and negation.

2 Related Work

We only discuss the most prominent approaches. A rule-based approach based on semantic verb classes was introduced by Neviarouskaya et al. (2009). Also fine-grained polarity labels are used there, namely the one from Appraisal Theory (Martin and White, 2005); positive or negative polarities are either related to appreciation (object inherent properties like “beautiful”), affect or judgement (the moral side, say “cheating”). Each verb instantiation is described from an internal and an external perspective. For example, “to admire a mafia leader” is classified as affective positive (the subjects attitude) given the internal perspective while it is judgement negative externally. The authors do not give any details about how they carry out rule application. Also, compared to our frame-

work, they are not able to tell “admire refugees” from “admire a mafia leader” as discussed above.

Reschke and Anand (2011) capture the polarity of a verb frame instantiation as a function of the polarity of the verb’s roles. For instance, if a terrorist loses something positive, then this is positive as a whole. No rules are specified, just evaluativity classes are defined. It is hard to see how less drastic cases are to be treated (e.g. “the thief who loses all his friends” - is this positive?).

Recently, Deng and Wiebe (2014) have introduced an ambitious conceptual framework for inferring (sentiment) implications. Here, the private state of the author of a sentence is in question, what his attitudes towards various targets in the sentence are (and also what he believes the private states of the agents of the sentence would be). Rules within a graph-based model are used, propagation continues until no rules can be applied anymore. The model builds (in part) on the ideas described in Reschke and Anand (2011) (to commit a crime is good for the crime since it comes into existence), in contrast to Neviarouskaya et al. (2009) no external perspective is envisaged. “to admire” is a good-for situation, it is unclear how this influences the critical cases: “admire refugees” compared to “admire a mafia leader. Does the last example produces a polarity conflict? We would say that it should and that the subject receives a negative effect as a consequence.

3 Verb Polarity Frames

Some verbs do impose polar restrictions and cast polar perspectives on their complements, see Klenner et al. (2014a) and Klenner et al. (2014b). Take “to blame someone for something”. As a consequence of the instantiation of “to blame” the direct object receives a negative effect and we expect the PP to realize something negative (or something neutral that contextually is perceived as being negative). We distinguish *verb role expectations* from *verb role effects*. *Effects* are propagated as the result of the verb frame instantiation (given a concrete sentence), while *expectations* are a kind of presuppositional aspects of the meaning of a verb - they are constant under negation: “not to blame for X” and “to blame for X”, both presuppose that X is negative.

Verb polarity frames are used to capture this kind of information: “to blame” is a “direct object negative effect verb” and it also is a “PP com-

plement negative expectation verb”. Our German verb model comprises about 140 verbs (about 300 verb polarity frames) and their effects and expectations. Although there are cases, where a particular syntactic frame of a verb allows for more than one verb sense, this is not the rule. Often sortal restrictions might help in these cases (“sorgen für” in German might denote “care for”, if a person is involved and “to organize”, if e.g. a non-animate object takes the role: “sorgen für Verpflegung”, “to organize food”). Word sense disambiguation (of verbs) is thus not so much a problem and we ignore the rare cases, where we would need it.

We are not interested in the private state of the author of a text, but in the polar conceptualization the text imposes on potential targets. Our goal is a target-specific sentiment analysis, which means that we focus on positive and negative contexts the target is in. Verb polarity frames capture the unmarked condition of verb usage. That is, deviations (e.g. expectation violations) are possible, but they lead to sanctions. For instance, if someone criticizes something positive or if someone supports something negative then he receives a negative effect (stemming from - what we call - a polarity clash between verb polarity frame and the polarity of the actual filler object). A *polarity clash* or violation arises, if a polar expectation or a polar effect casted from the verb meets an inverse entity polarity that comes either from the polarity lexicon or from noun phrase composition (“sick colleague”).

Our polarity lexicon comprises about 6,700 nouns and adjectives (see Clematide and Klenner (2010)). Nouns that denote factual negative concepts form a new class of objects, namely those that deserve (or have) our sympathy (refugee, victim, hostage, depressed and poor people etc.). This helps to overcome problems with sentiment composition: “to admire something negative” is negative, except if we talk about negative entities that have our sympathy: “to admire a sick colleague” (no polarity clash) compared to “to admire a cheating colleague” (polarity clash).

The distinction between verb effects and verb expectations is crucial - we need both. Effects might result in positive or negative polarity attributions to entities, which is not true for expectations. Both are needed in order to identify unexpected and in a sense deviating situations as in “to blame for success”.

4 A Case Study

In order to clarify the kind of sentiment inferences we envisage, we work through the sentence: “Die ganze Welt kritisiert Europa dafür, dass es den Flüchtlingen nicht hilft.” A translation which preserves the German subclause construction is given in example 1. Examples 2-4 are variants of it regarding negation, i.e. ”not” or ”refuse to” (POS = positive):

1) criticize that NOT POS: The whole world criticizes Europe for the fact that it does not help the migrants.

→ In our verb model, “to criticize” imposes a negative effect on both, the direct object and the subclause. “to help” has a positive effect on the indirect object, which gives a positive situation, while the “criticize” frame is negative. Here, no (positive) effect is propagated to “migrants” since the subclause is negated, an “ordinary” negative effect is given to “Europe” (stemming from “criticize”). Moreover, “not to help someone who needs our help” is an odd situation (we call it a *Neg_Clash_Situation*) and the subject of such a situation (Europe) is penalized (it is per definitionem a *Neg_Clash_Entity*).

2) criticize that POS: Russia has criticized Europe for the fact that it helps the Ukraine.

→ “Ukraine” receives a positive effect; since a violation (“to criticize”: a negative effect meets a positive sentence) is encountered, “Russia”, as the actor of such a situation clash, is a *Neg_Clash_Entity*, while “Europe” receives a negative effect from being criticized (but since “criticize” is here a *Neg_Clash_Situation*, this might be regarded as irrelevant or neutralized - not yet implemented).

3) NOT (refuse to) criticize that NOT POS: China has refused to criticize Europe for the fact that it does not help the migrants.

→ “Europe” and “China” are of class *Neg_Clash_Entity*

4) NOT (refuse to) criticize that POS: USA has refused to criticize Europe for the fact that it helps the Ukraine. → “Ukraine” gets a positive effect

Deeper embeddings are possible, e.g.: “The USA forces the UN not to criticize Europe for the fact that it helps the Ukraine.” Eventually, our model should be able to handle this (and all permutations of “not” among the clauses) as well.

5 A Prototype Implementation

We use the description logic OWL (Horrocks and Patel-Schneider, 2011) to represent the verb model, the effect propagation, the polarity clash classes, and we let the reasoner carry out all inferences coupled with sentiment propagation. Our model is meant to be a competence mode: we model cases even if they hardly are to be found in real texts. We are interested in the underlying principles (the correct predictions), not so much in the empirical impact (in terms of recall) given a sentiment analysis application.

Description logics (DL) were first introduced in the 1980s – so called KL-ONE languages. The big advantage of this (subset of predicate) logic is that the consistency of a knowledge base and also decidability are guaranteed. The reasoner can be used to relate instances from the so-called A-Box (assertions) to the concept hierarchy, the terminological part (the T-Box). The system cares for the proper (truth) maintenance of derived facts, i.e. if an assertion is ceased, any inferences made on the basis of it are ceased as well. We use the Protégé¹ editor for engineering and Hermit (Glimm et al., 2014) as an inference machinery. The following OWL specifications are given in Manchester Syntax (Horridge et al., 2006).

5.1 A-Box Representation

In order to carry out inferences, individuals (class instantiations) need to be specified (then reasoning or realization, the DL term, might be triggered). Verbs are referred to by a constant representing a verb event instantiation. Technically, their base form followed by a digit is used (e.g. blame-1 is an instance of a blame event). Binary relations link the participants to their event. We stick with grammatical functions (subject, etc.) for readability, but they are meant to represent arguments (AO, A1, etc.). The labels subj, obja (direct object), objd (indirect object) and objc (complement clause) are modelled as OWL properties, we also define a property class “participant” that covers subj, obja and objd and we define inverse properties, e.g. subj-of.

So “Russland kritisiert, dass Europa der Ukraine hilft” (take the translation given in section 4, example 2) would be represented as given in Fig. 1. The class “asserted” is used to indicate

¹available from <http://protege.stanford.edu/>

that the verb instance is not negated (non_asserted, otherwise).

Description: kritisieren-5	Description: helfen-5
Types +	Types +
● asserted	● asserted
● kritisieren	● helfen
Property assertions: kritisieren-5	Property assertions: helfen-5
Object property assertions +	Object property assertions +
■ subj Russland-5	■ subj Europa-5
■ objc helfen-5	■ objd Ukraine-5
■ obja Europa-5	

Figure 1: Instance Representation

5.2 Concepts for Sentiment Inference

Verbs have effects and expectations (see section 3 for their definition). Effects are directly applied if a verb polarity frame instantiation is feasible without any violation (e.g. of expectations). A target then is classified as a Pos_Eff_Entity or Neg_Eff_Entity depending on the polarity frame. If everything is as expected, the instantiation as a whole is classified as a Pos_Situation or a Neg_Situation according to the verb class. A violation occurs if a polar effect or polar expectation is not met. Then a Neg_Clash_Situation is found. Effects might get inverted or canceled and the subject of the situation is classified as a Neg_Clash_Entity. A situation that embeds a Neg_Clash_Situation is also a Neg_Clash_Situation and its subject then becomes a Neg_Clash_Entity and so on.

5.3 T-Box Representation

Fig. 2 shows a simplified version of the taxonomy. We distinguish attributes (Attribute, polar and non-polar adjectives) from entities (Entity) and situations (Situation). Entities comprise non-polar entities (e.g. nations) and polar entities, which divide into Neg_Entity and Pos_Entity (the polarity lexicon entries) and Eff_Polar_Entity (meant to capture the inferred verb effects), Neg_Clash_Entity (expectation violation entities) and Composed_Polar_Entity – the class that realizes noun phrase composition (e.g. that “cheating colleague” is negative). Sympathy_Entity is meant to capture entities that have our sympathy (e.g. “refugee”).

Situations are divided into polar situations (Polar_Situation) which capture positive and negative clause level denotations and Neg_Clash_Situations

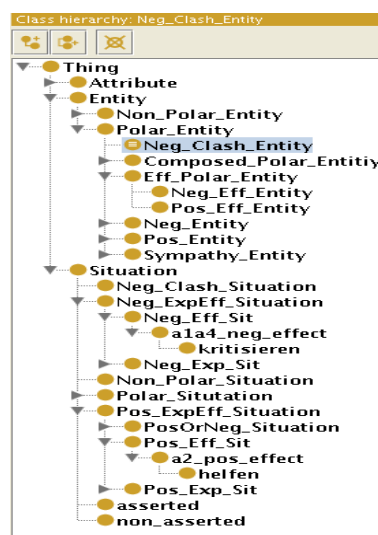


Figure 2: OWL Classes (partial)

(violation of subclause expectations or effects, e.g. “criticizes something positive”) and the various verb classes.

Verb classes are represented as primitive (undefined) concepts, their class name indicates effect and expectation patterns. For instance, “helfen” (“to help”) is, among others, a a2_pos_effect verb, since the indirect object receives a positive effect. We use the following shortcuts: a0, a1, a2, a3, a4 for subj, direct object, indirect object, PP complement, clausal complement, respectively. A primitive definition prevents the reasoner from automatically determining class membership. There is no need to let the reasoner classify verbs, since the parse tree provides all information needed in order to identify the verb class (the verb lemma, its grammatical roles).

The class a2_pos_effect is a subclass of Pos_Eff_Situation. In order to assign the effect that comes with the verb classes, the concept Pos_Eff_Entity is defined as a non-primitive subclass (equivalence class in Protégé) of Eff_Polar_Entity²:

objd-of some (a2_pos_effect and asserted)

That is, the indirect object (objd) of a a2_pos_effect verb that is not negated (i.e. asserted) is automatically classified as a Pos_Eff_Entity (i.e. it gets a positive effect). The equivalent in terms of predicate logic is:

$$\forall x : (\exists y : \text{objd-of}(x, y) \wedge \text{a2_pos_effect}(y) \wedge \text{asserted}(y)) \rightarrow \text{Pos_Eff_Entity}(x)$$

A situation (denoted by the verb and its instan-

²All definitions given here are simplified.

tiated roles) is positive situation (a subclass of Polar_Situation from Fig. 2), if (among others):

a2_pos_effect and asserted and (objd some (Sympathy_Entity or Comp_Neut_Entity or Comp_Pos_Entity))

That is: the verb is a a2_pos_effect and asserted (not negated) and the indirect object of that verb is an entity that has our sympathy or is a composed positive or neutral entity (nouns without further modification like “Ukraine” are treated as simply composed, neutral entities).

The verb “to criticize” is a a1a4_neg_effect verb, the direct object as well as the subclause inherit a negative effect. For our example (“Russia has criticized Europe for the fact that it helps the Ukraine”) we get (see Fig. 3): “helfen-5” is a positive situation³, thus, a polarity clash occurs (“criticize something positive”).



Figure 3: Sentiment Inferences

The class Neg_Clash_Situation captures this (cf. Fig. 4). The first line states that a situation with a participant (a property class that subsumes subj, obja and objd) that is a M_Neg_Comp_Entity and is at the same time a Pos_Effect_Entity, is a Neg_Clash_Situation. That is, any situation, where a (morally) negative entity has received a positive effect, is a situation clash.




Figure 4: Neg_Clash_Situation

According to line 2 (and 3), a non-negated a1a4_neg_effect verb that embeds a Pos_Situation (or, to capture recursive cases, a Neg_Clash_Situation) is classified as a Neg_Clash_Situation. “kritisieren-5” is classified accordingly (see Fig. 3) and the next step

³Lines marked in yellow are inferred by the reasoner.

is to propagate this situation polarity clash to the subject. We define a Neg_Clash_Entity to be the subject of a Neg_Clash_Situation: (subj-of some Neg_Clash_Situation). In our case, this holds for “Russland-5”. Russia thus is penalized for criticizing a positive situation.

5.4 The Current State of the Model

A dependency parse of a sentence is easily converted into the input format of our model. One just has to take the verb from the tree, take the lemma as its class, create a verb instance, find its grammatical roles and create constants in order to provide individual referents. If the verb is negated, its OWL instance is set to non_asserted.

We have created an interface to a dependency parser (Sennrich et al., 2009), but just worked with hand-crafted sample sentences (that are inspired by real sentences, anyway). We expect that an empirical evaluation will reveal gaps in the model and parts that need to be refined. It is already clear that we need to deal more explicitly with implication signatures of verbs in the sense of Nairn et al. (2006). For instance, German “zwingen” turns out to be a one-way implicative (which is in line with its English translation “to force”). Thus, negated “zwingen” does not entail anything regarding the factuality (truth) of its complement clause.

Currently, we have 12 verb classes in the model, covering about 80 different verbs.

6 Conclusions

This paper introduces description logics as a framework for verb-centered sentiment inferences. We have sketched a model and implemented a prototype of it in order to demonstrate the power of such an approach. Subclause embedding, negation and polarity conflicts can be handled in such a framework in a very concise way. We have not yet fully explored the various model variants that appear to be interesting. Further experiments and an empirical evaluation are needed.

Acknowledgements: I am indebted to my colleague Michael Amsler for the specification of the German verb model, which I used as a basis for my explorations. I would like to thank Simon Clematide for his suggestions. This work has been partially supported by the Swiss Commission for Technology and Innovation (CTI) under the grant 17384.1 PFES-ES.

References

- Franz Baader. 2009. Description logics. In Franconi E. Eiter Th. Gutierrez C. Handschuh S. Rousset M.-C. Schmidt R. Tessaris, S., editor, *Reasoning Web: Semantic Technologies for Information Systems*, pages 1–39. Springer.
- Simon Clematide and Manfred Klenner. 2010. Evaluation and extension of a polarity lexicon for German. In *Proceedings of the First Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA)*, pages 7–13.
- Lingjia Deng and Janyce Wiebe. 2014. Sentiment propagation via implicature constraints. *Meeting of the European Chapter of the Association for Computational Linguistics (EACL-2014)*.
- Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. 2014. HermiT: An OWL 2 reasoner. *Journal of Automated Reasoning*, 53(3):245–269.
- Nora Hollenstein, Michael Amsler, Martina Bachmann, and Manfred Klenner. 2014. Sa-uzh: Verb-based sentiment analysis. In *Proceed. of the 8th Intern. Workshop on Semantic Evaluation (SemEval 2014)*, pages 503–507.
- Matthew Horridge, Nick Drummond, John Goodwin, Alan Rector, Robert Stevens, and Hai H Wang. 2006. The Manchester OWL syntax. In *OWL: Experiences and Directions (OWLED)*.
- Ian Horrocks and Peter F. Patel-Schneider. 2011. KR and reasoning on the Semantic Web: OWL. In John Domingue, Dieter Fensel, and James A. Hendler, editors, *Handbook of Semantic Web Technologies*, chapter 9, pages 365–398. Springer.
- Manfred Klenner, Michael Amsler, and Nora Hollenstein. 2014a. Verb polarity frames: a new resource and its application in target-specific polarity classification. In *Proceedings of KONVENS 2014*, pages 106–115.
- Manfred Klenner, Susanna Tron, Michael Amsler, and Nora Hollenstein. 2014b. The detection and analysis of bi-polar phrases and polarity conflicts. In *Proceedings of 11th International Workshop on Natural Language Processing and Cognitive Science*. Venice, Italy.
- Isa Maks and Piek Vossen. 2012. A lexicon model for deep sentiment analysis and opinion mining applications. *Decision Support Systems*, 53(4):680–688.
- J. R. Martin and P. R. R. White. 2005. *Appraisal in English*. Palgrave, London.
- Rowan Nairn, Cleo Condoravdi, and Lauri Karttunen. 2006. Computing relative polarity for textual inference. In *Proceed. of Inference in Computational Semantics (ICoS 5)*, Buxton, England, pages 67–75.
- Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. 2009. Semantically distinct verb classes involved in sentiment analysis. In Hans Weghorn and Pedro T. Isaías, editors, *IADIS AC (1)*, pages 27–35. IADIS Press.
- Kevin Reschke and Pranav Anand. 2011. Extracting contextual evaluativity. In *Proceedings of the Ninth International Conference on Computational Semantics*, pages 370–374.
- Rico Sennrich, Gerold Schneider, Martin Volk, and Martin Warin. 2009. A new hybrid dependency parser for German. In *Proc. of the German Society for Computational Linguistics and Language Technology*, pages 115–124.