A model for multi-perspective opinion inferences

Klenner, Manfred

IJCAI 2016 workshop

Natural Language Processing meets Journalism

Proceedings of the Workshop

Larry Birnbaum, Octavian Popescu and Carlo Strapparava (eds.)

10 July 2016
New York City

A Model for Multi-Perspective Opinion Inferences

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

Abstract

A text might give rise to various projections: a writer, a text and a reader projection. Given the (proclaimed) factuality of a text, the overt or hidden attitudes between the various referents can be inferred, as well as the writers opinion and - given the reader's preferences - his or her perception of the whole. Moreover, some sentences might even indicate controversial topics if viewed from a common sense perspective. We introduce an approach based on Description Logics that integrates these various perspectives into a joint model.

1 Introduction

Sentences might express a positive or negative relationship between people, organizations, nations etc. For instance, in the sentence "EU supports Greece" a positive attitude of the EU towards Greece is expressed. At the same time, a positive effect that is meant to be true is asserted. That is, Greece benefits from the situation described. If the reader has a positive attitude towards the beneficiary (Greece), he might regard the initiator (EU) as a benefactor and, thus, takes a positive attitude towards him as well (he is a proponent of his). If he does not like the beneficiary for some reasons, he might, as a consequence, regard the seemingly benefactor as his opponent. If the sentence is negated or embedded into a non-factive verb like "to pretend" ("EU pretends to support Greece") neither the positive relationship between the referents nor the positive effect on Greece do hold any longer. Instead, the matrix verb "to pretend" casts a negative effect on EU. If a positive effect in such a sentence is casted on an entity that from a common sense perspective is negative, then the actor of the described situation might be regarded as a common sense disturber (e.g. "The minister supports terrorism").

This is the kind of reasoning we have in mind. We would like to be able to answer the following questions: Given a text, what is good or bad *for* the entities mentioned in the text, what is good or bad *of* these entities, what are the attitudes of the entities towards each other and what follows from the reader's stance, i.e. his prior attitudes towards some entities, for his attitudes towards the entities mentioned in the sentence. The user of our system then could mine texts for proponents and opponents of his, in the sense that entities that do things (or like others that) he likes are proponents and entities that act in the opposite way (or like others he dislikes) are opponents. Also, controversial topics can be identified on the basis of a common sense perspective.

In contrast to existing work, we stress the point that verb signatures in the sense of [Karttunen, 2012] capturing (non-)factuality information regarding complement clauses need to be taken into account in order to properly draw such inferences. We focus on complex sentences where a matrix verb restricts its subclauses with respect to factuality depending on its affirmative status (i.e. whether the matrix clause is asserted or negated).

We have realized a joint model with Description Logics (DL), namely OWL [Horrocks and Patel-Schneider, 2011] and SWRL [Horrocks and Patel-Schneider, 2004]. The OWL model is language-independent, however, the parser and the lexicon resources are not. We rely on English examples, however our pipeline (and the empirical evaluation) is for German.

2 Related Work

An early rule-based approach to sentiment inference is [Neviarouskaya *et al.*, 2009]. 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 subject's attitude towards the direct object) given the internal perspective while it is (as a whole) negative externally. Factuality and subclause embedding do not play any role in their work. The same is true for [Reschke and Anand, 2011]. They capture the polarity of a verb frame instantiation as a function of the polarity of the verb's roles - we, instead, do not know in advance, but intend to infer the (contextual) polarity of the roles. Recently, [Deng and Wiebe, 2015] have introduced an advanced conceptual framework for inferring (sentiment) implicatures. Their work is most similar to our approach. Various model versions exist, the most recent one [Deng and Wiebe, 2015] also copes with event-level sentiment inference, which brings it even closer to our model. Probabilistic Soft Logic is used for the definition of the model and for drawing inferences. The goal of the systems is to detect pairs of entities that are in a PosPair or NegPair relation. However, factuality is not taken into account in their framework, while we believe it is crucial for certain inference steps.

How Description Logics can be used to identify so-called polarity clashes is described in [Klenner, 2015]. However, attitudes and the factuality of situations are not part of that model.

3 The Verb Model: Polarity Frames

The basis of our approach is a verb resource that we call polarity frames, cf. [Klenner and Amsler, 2016]. The current lexicon comprises 330 German verbs which gives 690 polarity frames. We are particularly interested in those verbs that subcategorize for complement clauses (78 verbs), since especially they are crucial for reasoning.

For each argument (agent, patient etc.) of a polarity frame we specify whether it casts a polar effect on its argument filler, e.g. the patient of "to help" gets a positive effect. We distinguish between effect roles that indicate that something is good/bad *of* or *for* someone. The agent role is an *of-role* - it is good *of* A to help B. The patient, but depending on the verb also the theme or recipient roles are *for-role*s, it is good *for* B if A helps him.

Take the verb "to help". There are at least two polarity frames, the transitive use (A helps B) and the one with an embedded (infinitival) subclause (A helps to XCOMP). In the first frame, both argument fillers receive a positive effect (in an affirmative, factual use of the verb). The agent is a positive *of-role*, which we call the *pos-of* role (a sub role of *of-role*). The patient is, accordingly, a *pos-for* role. Both roles are generalizations of the traditional roles (agent, ..). They ease the development of general inference rules and they have a particular function in the reasoning process. We would like to be able to state that something is good or bad *of or for* someone.

In the second frame (help to XCOMP), the agent again is the bearer of the *pos-of* role. But now it is the subclause that receives a positive effect, i.e. it is good for the situation denoted by the subclause that it receives help. Thus, not only entities but also situations are affected by the polarity a verb casts on its arguments. In order to distinguish roles for situations from roles for entities, we call the role for positively and negatively affected situations *poseff* and *negeff*, respectively.

3.1 Verb Signatures

Verbs that subcategorize for a clausal complement are further specified for (non-)factuality of the clausal complement. Factuality means that the situation described in the subclause is meant (by the writer) to be true (to hold). We follow the work of [Karttunen, 2012], who distinguishes factive, non-factive and implicative verbs. Factuality of the subclause depends on the (matrix) verb signature and the presence or absence of negation (in the matrix clause). Factive verbs such as "to regret" cast factuality on their subclause, whether the main clause is negated or not. If A regrets that COMP, then COMP is true in the sense that the speaker believes (or a least asserts) COMP to be true. The same holds for A does NOT regret that COMP. Sublcauses of non-factive verbs, on the other hand, are never meant to be true (e.g. "to pretend", "to hope").

Then there are verbs called implicatives that cast a mixture of factuality and non-factuality. Two-way implicatives like "to forget to" have non-factual subclauses in an affirmative use, but factual subclauses if negated. One-way implicatives only give rise to factuality in either the affirmative ("to force") or negated matrix verb contexts ("to refuse"). Table 1 summarizes the signatures, introduces the concept labels (e.g. clAF) we use to represent it and gives example verbs.

concept	explanation	matrix verb
clF	factual	to regret
clAF	factual, if affirmative	to force
clNaNF	non-factual, if non-affirmative	to manage
clANF	non-factual, if affirmative	to forget
clNaF	factual, if non-affirmative	to forget
clNaO	true or false, if non-affirmative	to help

Table 1: (Non-)Factuality of Subclauses

In Table 2 we give the polarity frames of some verbs.

		of	for	sc	aff	neg
1	to criticize	none	-	neg	clAF	clNaF
2	to refuse	none	-	neg	clANF	clNaO
5	to help	pos	pos	-		
7	to survive	-	pos	-		

Table 2: Polarity Frames

A hyphen indicates that the role is not part of the verb frame in question, *pos* and *neg* stand for positive and negative effect, respectively and *none* states that although the argument role exists, there is no (i.e. a neutral) effect attached to it (*sc* means subclause effect). The last two columns relate to the verb signatures as introduced in Table 1, the forelast column reports the restriction if the matrix verb is aff(irmative) and the last column if it is neg(ated). For example, the sublcause of "to refuse" (row 2) is non-factual if the refuse sentence is affirmative (clANF), but its truth value is unspecified (clNaO) if negated.

4 Description Logics Model

We strive to be able to combine different perspectives in a joint model. Firstly, there is the question of who actually profits (or has a disadvantage) from the described situation. We call this the layer of *effect projection*. Then there is the relational level that determines the attitudes of the participants towards each other, this is called the *attitude projection*. Both are derived from the input text, they represent, so to speak, the way the text puts the world (the *text perspective*). There is also the perspective of the reader, the *reader projection* and the perspective of the author (not coped with in this paper). Finally, we also deal with what we call the common sense perspective. Here we focus on the detection of controversial topics where a polarity conflict occurs given the sentence.

Inferences are based on the text perspective, i.e. the view of the world that the author of the sentence intends to establish with his text. From the *text perspective* the attitudes of the author sometimes are evident, but in the kind of sentence that we envisage, this is normally not the case. We focus on sentences that report the view of the subject of the matrix clause ("A criticizes that ..").

Effect	Attitude	Reader
beneficiary	pro	MyOpponent
benefactor	contra	MyProponent
victim	cs_disturber	SympathyEntity
villain		NonSympathyEntity

Table 3: Projections: Concepts and Properties

Description Logics seemed to be well suited for such intermingled inference tasks. One must not care about the concrete sequence the inferences are drawn and there is the notion of global consistency that might help to identify and get rid of unwanted side effects. It turned out to be convenient to use SWRL rules [Horrocks and Patel-Schneider, 2004] instead of pure OWL concepts [Horrocks and Patel-Schneider, 2011] to define the relational inference layer. Our system was developed on the basis of the Protégé editor, HermiT [Glimm *et al.*, 2014] was used as a SWRL and OWL reasoner.

4.1 Overview: Concepts and Properties

Table 3 shows the concepts and properties of the various projection layers. We give a brief description of the overall system in order to instantiate the OWL constructs from Table 3. We use a dependency parse tree as input. A simple rule-based component (see [Klenner and Amsler, 2016]) extracts the grammatical roles (subject etc.) of each verb from the parse trees (thereby normalizing passive voice and making implicit arguments explicit, i.e. given control or raising verbs). The output of this component are the instantiated verb frames, i.e. the filler objects of the grammatical roles of the verbs given the sentence. Each grammatical role then is mapped to a polar role (*pos-of*, *neg-of* etc.). If we know the grammatical role of a referent then we know his polar role, that is the core functionality of our polarity frames. The next step is to produce the OWL representation of the sentence (see section 4.3). For every verb its affirmative status is given by the parse tree (this becomes also part of the OWL representation). Whether the main clause is factual or not is determined by a simple heuristics: if no modal verbs or modifiers are present, then the sentence is factual. Factuality of subclauses are predicted by OWL definitions.

In a nutshell this is how the various layers from Table 3) interact. Take “EU no longer supports Greece”. Here, “Greece” is victim – it suffers from the situation. From the parse tree we know that it is the direct object of “support”, from the polarity lexicon we known that the direct object is a *pos-for* role. Since the sentence negated, the *pos-for* gets inverted and becomes a *neg-for* role. Now the OWL definition of a victim is met (see section 4.4). This is an example of an effect projection. Furthermore since “EU” is responsible for a negative effect on “Greece”, it must have a negative attitude towards “Greece”, a *contra* relation if found (an attitude projection). Finally, if “Greece” is a *SympathyEntity* of the reader (the concept representing the reader’s prior attitudes), then “EU” becomes an instance of *MyOpponent* of the reader. If the sentence was “EU supports neoliberal greed”, then “EU” becomes a common sense disturber (cs_disturber) since a polarity conflict occurs. A positive effect on a negative denotation (“neoliberal greed”) is found, which is from a common sense perspective not desirable, it represents a conflict, thus.

4.2 Properties

OWL properties represent two-placed relations between concepts, they have domain and range restrictions (we do not specify the concrete restrictions here). We have a property *for-role* with sub properties *pos-for* and *neg-for* and a property *of-role* with *pos-of*, *neg-of* as sub properties. These are roles for entities, for situations we use a general role *cl-role* denoting a non-polar subclause (e.g. the verb ”to remember” (that) would have it) and *negeff* and *poseff* for positive and negative effects, the matrix verb casts on its complement clause. These roles also have inverse roles, indicated by an preceding initial I (e.g. *I-pos-of*). Table 4 summarizes these properties. They are use to represent an input sentence, i.e. the instantiated verb frames. We now turn to this part of model. Please note that, in contrast, the properties of the attitude projection (cf. Table 3, second column) are subject to SWRL inference rules (see section 5).

of-role	the agent
(pos|neg)-of	the filler gets a positive (negative) effect
for-role	the patient,recipient, beneficiary or theme
(pos|neg)-for	a positive (negative) for-role
cl-role	the subclause
(pos|neg)eff	subclause receives a positive (negative) effect

Table 4: Properties Representing Verb Argument Roles

4.3 A-Box Representation

We represent verb instantiations in a manner that is inspired by Davidson’s approach [Davidson, 1967]. Our example sentence, ”The minister has criticized that the EU has helped Greece to survive” is represented by the assertions from Table 5 (the specifications are given in a slightly simplified Manchester syntax, cf. [Horridge *et al.*, 2006]).

criticize-1 : (aff AND clAF)	help-1 : (aff AND clAF)
criticize-1 of-role minister-1	help-1 pos-of EU
criticize-1 negeff help-1	help-1 pos-for Greece
survive-1 : affirmative	help-1 poseff survive-1
survive-1 pos-for Greece	

Table 5: A-Box Representation

$criticize$-1 is a instance of both, the classes *aff* irmative and *clAF* (and, not shown here, *clNaF*), it has e.g. the role *negeff* with $help$-1 as its filler. The concepts *aff* irmative and *non-aff* irmative are used to represent the affirmative or negated use of a verb predicate in a sentence. The individuals $minister$-1, EU and $Greece$ are all instances of a general concept called *RealWorldEntity*.

4.4 T-Box

As mentioned, we distinguish between the perspective of the reader, *MyView*, and the perspective of the text, *TextView*, see Fig.1. *TextView* tells us, what the author believes to be true. One task of the reader as part of the understanding of a text is to find out what the text entails (class *Implication*) about the described situation (class *Situation*). A situation is either affirmative (class *affirmative*) or negated (class *non-affirmative*), which is known given the sentence (thus, both are primitive concepts).

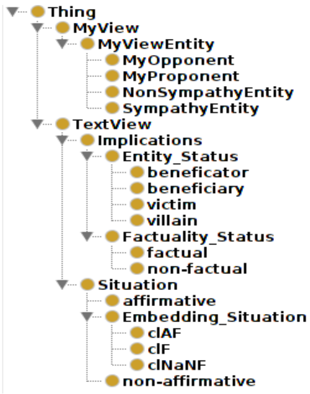

Figure 1: T-Box

The factuality of an embedd*ed* situation depends on the factuality class (e.g. clAF) of the embedd*ing* situation (class *Embedding_Situation*) given by the verb signature of the (embedding) verb (see Table 1 for the subclasses of *Embedding_Situation* not shown in Fig.1). For instance, according to Table 5 *criticize*-1 is an instance of *clAF* since the verb "to criticize" bears that signature: whatever affirmative *criticize* embeds, it is factual[1]. Thus, all subclasses of *Embedding_Situation* are primitive concepts. Whether an embedded situation is factual or non-factual (its *Factuality_Status*) depends on the factuality class of the embedding verb and whether the embedding verb is affirmative or non-affirmative: *factual* and *non-factual* are defined classes. The definition of *factual* is (in Manchester syntax):

(I-cl-role some (clF or (affirmative and clAF) or (non-affirmative and clNaF)))

I-cl-role is the inverse of *cl-role*.

A situation is *factual* if it is embedded (*I-cl-role*) into a situation that is described by a factive verb (class *clF* from Table 1), or is *affirmative* and has the signature *clAF* or is *non-affirmative* and of type *clNaF*. Given this (together with the definition of *non-factual*), we are able to determine the factuality status of an embedded situation of any depth of embedding.

We now turn to the concept *Entity_Status*. We distinguish four classes and call them programmatically *benefactor*, *beneficiary*, *victim* and *villain*. We just give the definition of *beneficiary*. The idea behind our definition is that the beneficiary of a situation is somebody who benefits from it independently of any attitude that somebody might have towards him. So if A wins, A is the beneficiary, if A is liked by someone or not. What must be the case is that A occupies the *pos-for* role of a situation that is *factual* (not just imagined) and *affirmative*. Here is the definition of *beneficiary*:

(I-pos-for some (affirmative and factual))

5 SWRL Model: Attitude Projection

The main goal is to find out, whether A is for B, which we model with the property *pro*; or whether A is against B, here *contra* is used.

Firstly, a verb might directly reveal the relation between the participants within the same clause: if A supports B, then A is pro B. Provided, of course, the situation is *factual*. In our SWRL rules the following class abbreviations are used: fact=*factual*,aff=*affirmative*,neg=*non-affirmative*,pfor=*pos-for*,nfor=*neg-for*.

```
r1 fact(?s),aff(?s),pfor(?s,?y),of_role(?s,?x)
     -> pro(?x,?y)
```

The first rule (variables are indicated by a leading question mark, e.g.?x) (r1) states: An actor ?x (the *of-role*) is pro ?y if in a factual, affirmative sentence ?s, ?y is the filler of the *pfor* role (e.g. "A supports B" gives pro(A,B)).

If a sentence ?s embeds a sentence ?s2, then rules like the following are in charge:

```
r5 aff(?s),fact(?s),aff(?s2),negeff(?s,?s2),
   of_role(?s,?x),nfor(?s2,?y)
      -> pro(?x, ?y)
r7 aff(?s),fact(?s),neg(?s2),negeff(?s,?s2),
   of_role(?s,?x),nfor(?s2,?y)
      -> contra(?x, ?y)
```

According to r5 an affirmative and factual clause ?s that embeds an affirmative subclause ?s2 bearing a negative effect (*negeff*) gives rise to a *pro* relation between the *of-role* of the matrix clause and the *nfor* role of the subclause. If A criticizes (clAF) or fears (clANF) that B punishes C, then A is pro C.

The agent-patient relation of rule r5 only holds if both, the matrix ?s and the subclause ?s2 are affirmative. If ?s2 is negated (cf. rule r7), then *pro* turns into *contra* (A criticizes that B does not punish C gives contra(A,C)).

More complicated scenarios arise in the case of multiple embeddings. We discuss this given the two example sentences: 1) A criticizes that B refuses to help C and 2) A criticizes that B not refuses to help C. The task here is to fix the attitude of the subject of the matrix clause wrt. to any role at any level of subclause embedding. According to Table 2, both "to criticize" and "to refuse" put a negative effect on their complement clauses. We could say then that A (the matrix subject of example sentence 1) disapproves the negative effect of refuse, and thus *approves* the help situation (all this provided that the matrix situation is affirmative; the intermediate subclause must be affirmative as well; no information is needed wrt. to the innermost subclause).

```
r8 aff(?s),fact(?s),aff(?s2),of_role(?s,?x),
   negeff(?s,?s2),negeff(?s2,?s3)
     -> approve(?x,?s3)
r9 aff(?s),fact(?s),neg(?s2),of_role(?s,?x),
   negeff(?s,?s2),negeff(?s2,?s3)
     -> disapprove(?x,?s3)
```

[1] Clearly in: "A criticizes that B intends to lie", the intention is factual, not the lying.

```
r10 aff(?s),fact(?s),aff(?s2),of_role(?s,?x),
    negeff(?s,?s2),poseff(?s2,?s3)
     -> disapprove(?x,?s3)
```

That is: A *negeff* on a *negeff* gives (if ?s and ?s2 are affirmative) *approve* (see r8, sentence 1). If ?s2 is negated (?s is affirmative), a *negeff* on a *negeff* gives *disapprove* (see r9, sentence 2). The next rules describe how *approve* and *disapprove* propagate to pro or $contra$ properties.

```
r11  approve(?x,?s),aff(?s),pfor(?s,?y)
       -> pro(?x,?y)
r13 disapprove(?x,?s),aff(?s),pfor(?s,?y)
       -> contra(?x,?y)
```

If someone approves an affirmative situation that is positive (*pos-for*) for someone, then he is for this person (rule r11). Sentence 1: A is pro C (rule r8 and r11). According to rule r9 and r13, A is contra C (sentence 2).

Finally, some rules are used to propagate $contra$ and pro to derived $contra$ and pro properties. According to rule r15, if A is against B and B is against C then A is for C.

```
r15 contra(?x, ?y),contra(?y,?z)-> pro(?x,?z)
r18 pro(?x,?z),contra(?y,?z) -> contra(?x,?y)
```

Take: "A hopes that B does not offend C". Here, A is for C, but there is no inference regarding A's attitude towards B. However, if we know (e.g. from world knowledge) that B is against C, i.e. contra(B,C), then we can derive that A (presumably) is against B (rule r18: pro(A,C),contra(B,C) thus contra(A,B)).

6 Example

Take the following (hypothetical) sentence with the A-Box representation given in Table 5: "The minister has criticized that the EU has helped Greece to survive". The following inferences take place.

```
Beneficiary(Greece)          by OWL definition
pro(EU,Greece)                     by rule r1
contra(minister-1,EU)              by rule r8
disapprove(minister-1,survive-1)   by rule r10
contra(minister-1,Greece)          by rule r13
```

If Greece is an instance of *SympathyEntity*, it follows (by OWL definitions and the derived pro and $contra$) that

```
MyProponent(EU)
MyOpponent(minister-1)
```

7 Empirical Evaluation of the Core Model

We have implemented a prototype system for German: a verb frame extractor and converter to A-Box representations. There is no annotated German corpus available, so we created a gold standard of 50 sentences from newspaper texts. A sentence in order to get selected was required to have at least two verbs from our lexicon and two named-entities as role fillers of these verbs. From the first 1000 sentences we get from the WaCky corpus [Baroni M., 2009], we randomly selected 50 and annotated them for pro, $contra$, *beneficiary*,*victim*. A f-measure of 84.39% was achieved, the precision was 82.94%. and the recall was 85.88%. A error analysis revealed that parsing errors (subject mistaken as object etc.), missing polarity frames (especially prepositional phrases) and verb ambiguity are the main causes for the errors.

8 Common Sense Conflicts

In the core model, every actor who has according to the attitude projection a positive attitude (a pro relation) towards an instance of the reader's *NonSympathyEntity* is an *Opponent* of the reader. The class *NonSympathyEntity* is meant to capture those real-word entities the reader does not like - particular political parties, politicians, etc. These are personal preferences. But what about entities whose polar value is culture specific? We use the concepts *CommonSensePositiveEntity* and *CommonSenseNegativeEntity* in order to represent this kind of information. For instance, a terrorist would belong to *CommonSenseNegativeEntity* while freedom is an instance of *CommonSensePositiveEntity*. Such knowledge is captured normally by a polarity lexicon. We could merge it into *(Non)SympathyEntity*, but this would confuse personal preferences with broader accepted shared preferences. We keep it separate in order to design a common sense conflict detector and to predict the class of common sense disturbers.

A *common sense conflict* is any situation, where an entity benefits (or suffers) from a situation, but that entity is not worth (does not deserves) it. For instance, if A support terrorism or if A disapproves freedom, a common sense conflict occurs. The actor of it is the *common sense disturber*. We believe that a system that is able to find such sentences could be of great interest for text exploration purposes. Sentences with common sense conflicts might indicate controversial topics, pronounced stance or unusual opinions or at least something that is not desirable from a common sense perspective.

Take these sentences (English translations) found by our system: "Moscow's half-baked attempt to solve the problem only strengthens the radical tendencies" and "These authoritarian forms of dealing with homosexuality are definitely accepted by those conservatives". According to our polarity lexicon, the direct object of "strengthen" ("radical tendencies") receives a positive effect. What makes it a common sense conflict is the bottom-up polarity of "radical tendencies" which is negative. Thus a positive effect (top-down) on a negative entity (bottom-up) establishes a polarity conflict. The same is true for the second sentence where something negative ("authoritarian forms") receives a positive effect (object of "accept") which is a conflict. Articles that contain such conflicts might contain controversial material, highlighting such sentences in a single article might help to focus on the most interesting parts of it.

Currently, we have 8 SWRL rules that establish this inference layer. The goal property is common sense disturber, *cs_disturber*. The first rule is:

```
c1  of_role(?s,?x),aff(?s),factual(?s),
    pfor(?s, ?y),cs_neg(?y)
     -> cs_disturber(?x)
```

If A acts in way that a positive effect (*pfor*) on B takes place, but B is a *CommonSenseNegativeEntity*, *cs_neg* for short, then A is a common sense disturber, *cs_disturber* (e.g. A supports terrorism). However, the degree of negativity of the direct object might play a role. If "The minister supports the rather poor argument", then this might be unwise, but does not touch his common sense integrity. Our polarity lexicon ([Klenner *et al.*, 2014] is designed along the principles of the Appraisal

Theory [Martin and White, 2005]. That is, we distinguish between factually, emotionally or morally positive or negative words. A conflict occurs only if the moral (e.g. crime) or emotional (e.g. fear) dimension is violated, not the factual (i.e. poor decision) one.

Also negated sentences are relevant (rule c2):

```
c2  of_role(?s,?x),neg(?s),factual(?s),
    nfor(?s,?y),cs_neg(?y)
     -> cs_disturber(?x)
```

If A acts in a way that a negative effect on a negative entity does not occur, then A is a common sense disturber.

More complex cases arise if subclause embedding is involved (rule c3):

```
c3  neff(?s,?s2),aff(?s),aff(?s2),
    factual(s2),pof(s2,?y),of_role(?s,?x)
    ->cs_disturber(?x)
```

If a negative effect (*neff*) on a subclause ?s2 is present and the actor of ?s2 receives a positive effect (*pof*), then A, the actor of the matrix verb is a common sense disturber. So if A criticizes that B has helped C to survive, then A is a common sense disturber.

We have carried out a first empirical test on the basis of the German newspaper treebank TüBa-D/Z [Telljohann *et al.*, 2009] comprising 95'500 sentences. Clearly, we cannot expect a huge number of such conflicts, our small lexicon, parsing and verb frame extraction errors are part of the problem. However, the system predicted 64 conflicts from which 31 were - after manual inspection - real conflicts (cf. the two examples from above). So precision is about 50%. Every second sentence proposed by the system does actually point out some interesting charged constellation.

9 Summary

Our model strives to answer the following questions, given a parsed text and the personal preferences of a single user: who benefits (or suffers) from the situations described, what does the text (implicitly) tells about the relationship of the actors involved, which topics does an actor like or dislike and - given all this - what does this implies for the user: who are proponents or opponents of his or hers. Our system is also able to predict situations that – from a common sense perspective – bear controversial or charged content. This could be useful as a new service in the area of media monitoring.

References

[Baroni M., 2009] Ferraresi A. Zanchetta E. Baroni M., Bernardini S. The WaCky Wide Web: A collection of very large linguistically processed Web-crawled corpora. *Language Resources and Evaluation*, 43(3):209–226, 2009.

[Davidson, 1967] Donald Davidson. The logical form of action sentences. In Nicholas Rescher, editor, *The Logic of Decision and Action*. 1967.

[Deng and Wiebe, 2015] Lingjia Deng and Janyce Wiebe. Joint prediction for entity/event-level sentiment analysis using probabilistic soft logic models. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 179–189, 2015.

[Glimm *et al.*, 2014] Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. HermiT: An OWL 2 reasoner. *Journal of Automated Reasoning*, 53(3):245–269, 2014.

[Horridge *et al.*, 2006] Matthew Horridge, Nick Drummond, John Goodwin, Alan Rector, Robert Stevens, and Hai H Wang. The Manchester OWL syntax. In *OWL: Experiences and Directions (OWLED)*, 2006.

[Horrocks and Patel-Schneider, 2004] Ian Horrocks and Peter F. Patel-Schneider. A proposal for an OWL rules language. In *Proc. of the Thirteenth International World Wide Web Conference*, pages 723–731. ACM, 2004.

[Horrocks and Patel-Schneider, 2011] Ian Horrocks and Peter F. Patel-Schneider. 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, 2011.

[Karttunen, 2012] Lauri Karttunen. Simple and phrasal implicatives. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics*, SemEval '12, pages 124–131, Stroudsburg, PA, USA, 2012. Association for Computational Linguistics.

[Klenner and Amsler, 2016] Manfred Klenner and Michael Amsler. Sentiframes: A resource for verb-centered german sentiment inference. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France, may 2016. European Language Resources Association (ELRA).

[Klenner *et al.*, 2014] Manfred Klenner, Michael Amsler, and Nora Hollenstein. Verb polarity frames: a new resource and its application in target-specific polarity classification. In *Proceedings of KONVENS 2014*, pages 106–115, 2014.

[Klenner, 2015] Manfred Klenner. Verb-centered sentiment inference with Description Logics. In *6th Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis*, pages 134–140, September 2015.

[Martin and White, 2005] J. R. Martin and P. R. R. White. *Appraisal in English*. Palgrave, London, 2005.

[Neviarouskaya *et al.*, 2009] Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Semantically distinct verb classes involved in sentiment analysis. In Hans Weghorn and Pedro T. Isaías, editors, *IADIS AC (1)*, pages 27–35. IADIS Press, 2009.

[Reschke and Anand, 2011] Kevin Reschke and Pranav Anand. Extracting contextual evaluativity. In *Proceedings of the Ninth International Conference on Computational Semantics*, pages 370–374, 2011.

[Telljohann *et al.*, 2009] Heike Telljohann, Erhard W. Hinrichs, Sandra Kübler, Heike Zinsmeister, and Kathrin Beck. Stylebook for the Tübingen treebank of written German (TüBa-D/Z). Technical report, Universität Tübingen, Seminar für Sprachwissenschaft, 2009.