Sentiframes: A Resource for Verb-centered German Sentiment Inference

Manfred Klenner, Michael Amsler
Institute of Computational Linguistics
University of Zurich
Binzmühlestrasse 14, 8050 Zurich, Switzerland
klenner@cl.uzh.ch, mamsler@cl.uzh.ch

Abstract

In this paper, a German verb resource for verb-centered sentiment inference is introduced and evaluated. Our model specifies verb polarity frames that capture the polarity effects on the fillers of the verb's arguments given a sentence with that verb frame. Verb signatures and selectional restrictions are also part of the model. An algorithm to apply the verb resource to treebank sentences and the results of our first evaluation are discussed.

Keywords: German Sentiment Analysis, verb-centered, Sentiframes

1. Introduction

Some verbs impose a positive or negative polarity on the fillers of some of their semantic roles. Often, an opinion holder is present in such a sentence as well. Then, the opinion holder is understood to have a positive or negative attitude towards the opinion target - as in *Europe* [holder] *criticizes Russia* [target]. Sometimes, we do not know the opinion holder, e.g., in passive constructions as *he was criticized*. Moreover, it could be argued that some verbs such *to win* or *to survive* not even implicitly refer to an opinion holder and nevertheless imply a target polarity. The relation between holder and target, thus, is not the most general perspective on these verbs - it is a special case. We have created Sentiframes[1], a resource of verb effects that does not specify holder-target relations, but verb-specific polar connotations (from which such a relational view could easily be derived).

We have created 471 so-called *verb polarity frames* for 251 German verbs. A verb polarity frame is a subcategorization frame that specifies for each verb role whether a polarity effect is triggered given that the appropriate filler object occupies the role. For instance, the direct object of *to criticize* receives a negative effect, while the direct object of *to admire* gets a positive effect. Also, subclauses might get an effect. Again, *to criticize that* assigns a negative effect to the subclause. This (in principle) supports inferences in the style of Deng and Wiebe (2014), where the relation between an opinion holder of the matrix verb and an actor of a subclause can be induced. We have not created such a system, yet. But the resource is meant to serve as the lexical basis of such an approach. A matrix verb might form an intensional context, e.g., in *I hope that my paper gets accepted* nothing positive or negative can be derived: *my paper* should not receive a positive effect in such a context. In order to prevent such invalid inferences, we have specified for each verb that subcategorizes for a subclause its verb signature in the sense from Nairn et al. (2006). This includes the behaviour under negation as well. Finally, we found that verb sense disambiguation - although in principle needed - often could be avoided by means of shallow selectional restrictions. So the German *sorgen für* is ambiguous between *to provide/to organize* (e.g., *für Nahrung sorgen* (Engl. *to organize food*)) and *to care for* (e.g, *für deine Mutter sorgen* (Engl. *to care for your mother*)). The crucial difference is the type of the PP object: either a human being or a non-animate object.

In this paper, we give an empirical evaluation of our resource. In order to avoid preprocessing (parsing) errors that obscure the picture, we use the dependency version of the Tübinger treebank, TüBa-D/Z (Telljohann et al., 2009), as a basis.

2. Description of the Resource

Previous versions of our approach including the verb resource are described in Klenner et al. (2014a), Klenner et al. (2014b) and Klenner (2015). Recently, we added another 90 polarity frames; and we specified coarse-grained selectional restrictions and verb signatures for the whole resource. Altogether, we have modelled 251 verbs, which provides us with 471 verb polarity frames. A single frame consists of the subcategorization frame, and the polarity effects associated with the roles.

Besides polarity effects, polarity expectations and polarity restrictions need to be specified in order to restrict the instantiation of a frame accordingly. For instance, the verb *verhindern* (Engl. *to prevent*) assigns a positive effect on the subject, but only if the direct object is negative (e.g., *to prevent an accident*), otherwise the effect is negative (e.g., *to prevent the solution of a problem*). These phrase-level, bottom-up *polarity restrictions* are determined on the basis of our polarity lexicon (comprising roughly 6,000 nouns and adjectives) and a plain sentiment composition (cf. (Klenner et al., 2014a)).

Polarity expectations are a weaker form of restrictions. While polarity restrictions exclude the case of neutral fillers as well (since the effect could not be reliably assigned), expectations only block inverse polarities but allow for neutral fillers. Take *to regret*: if someone regrets that the cat entered the house - a neutral subclause - he is in negative mood, obviously. On the other hand, if someone prevents that the cat enters the house, nothing of that sort could be safely induced.

[1] It was our colleague Simon Clematide who named it that way.

As an example for an annotated sentiframe, two entries for *verhindern* (Engl. *to prevent*) are given in Fig. 1. The first polarity case (the upper part of Fig. 1) describes the following restrictions and the instantiated frame: the subject role filler is animate and the direct object is non-animate. If we find in addition the direct object as carrying a negative effect (on the basis of a bottom-up (bup) calculated compositional polarity) the subject gets a positive effect through the instantiation of the frame. *To prevent an accident* would be an instantiation of such a polarity frame.

Restriction	Subject Role	Direct Object Role
effect	subj=positive	obja=bup-negative
selectional	subj=animate	obja=non-animate
effect	subj=positive	objc=bup-negative
selectional	subj=animate	objc=situation
verb signature	non-factual if unnegated	factual if negated

Figure 1: Two frames with respective roles and the verb signature for *to prevent*

The second polarity frame given (the lower part of Fig. 1) also specifies the verb signature in the sense of Nairn et al. (2006). The verb signature tells us about the factuality of the subclause given that the matrix verb is negated or unnegated. In this case, the subclause of *verhindern* is non-factual if *verhindern* is unnegated, if *verhindern* is negated, the subclause is factual. Of course, tense, modality etc. must also be taken into account. In: *She has prevented that he criticized the project manager* the project manager was not criticized while in the negated version (i.e., *She has not prevented*) he was criticized. Only in the second case the effect associated with *to criticize* should be assigned.
The combination of the polarity restrictions, the polarity expectations, and the verb signature allows us to infer the polarity for the involved roles (if there is any), according to the instantiated frames and via the produced effects.

3. Polarity Frame Instantiation

The actual role fillers of a verb, its predicate argument structure (PAS), are not necessarily easy to identify given a parse tree. Passive voice, the use of anaphors, verb ellipsis, verb (and noun) coordination and all cases of implicit verb arguments given raising and control verbs need to be considered. Since our goal was to evaluate the verb resource, we needed a way to safely extract the predicate argument structure from the trees. In order not to propagate parsing errors to our evaluation, we rely on the gold standard parse trees from the treebank.
On the basis of 120 manually created sentences that are meant to mirror the main cases of syntactic variation mentioned above, a set of extraction rules were automatically extracted from parse trees produced by a dependency parser ParZu (Sennrich et al., 2009) for German (cf. the development set, DevSet, from Fig.2). Only if a tree was correct wrt. the argument structure of the verb, we annotated the underlying argument structure, and also coordination, negation and auxiliary verb constructions. Also rules

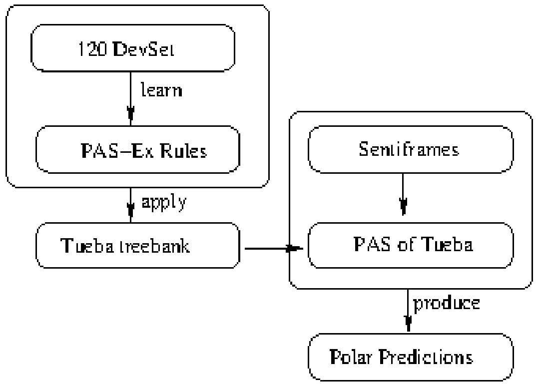

Figure 2: Evaluation architecture

for intrasentential coreference resolution are enabled this way. For instance, all control verbs that also allow for *that* subclauses (e.g., *to promise*) are given resolution preference according to their control type (subject or object control). For instance, the pronoun *he* gets resolved to *minister* in *The minister promised the chancellor that he helps the refugees*, since *to promise* is a subject control verb.
From correct parse trees and its predicate argument structures rules were extracted. Fig. 3 shows an example of an extraction rule derived from the sentence *Die EU profitiert davon, dass sie andere Länder zwingt, auf Wachstum zu verzichten.* (Engl. *The EU profits from the fact that it forces other countries to abandon growth*) and its predicate argument structures, e.g., *profitieren(EU,davon,zwingen)*. The entities in blue form the predicate (*profitieren*) and its arguments (*EU, davon, zwingen*), the labels in red indicate the paths (dependency labels). We derived 92 verb extraction rules, 22 rules for coreference, 5 rules for coordination, 6 rules for negation and 7 rules for auxiliary verbs (see Fig. 2: PAS-Ex Rules). The rationale behind this strategy, i.e., not manually specifying extraction rules but annotating syntax trees and extracting the rules, is that we can easily switch to another parser without the need to re-engineer the rules; we just have to run the extraction algorithm again.
We applied the learned extraction rules to the TüBa-D/Z, the Tübinger treebank (Telljohann et al., 2009), in order to extract correct frame instantiations from correct parses (PAS of Tueba from Fig. 2). Of course, the rule component is not perfect since there are rule conflicts and the applied conflict resolution is not capable of catching all errors. However, we believe that we have minimized noise stemming from preprocessing this way.
In order to extract polarity effects (Polar Predictions from Fig. 2) given the predicate argument structures of a sentence (one for each verb) and their corresponding polarity frames (Sentiframes from Fig. 2), an inference algorithm is needed. In this paper, we are only interested in the effects stemming from the verb frame instantiations. For example, our system predicts that the entity EU receives a positive effect (cf. Fig. 3) from *profits* and that *countries* gets a negative effect from *verzichten* - *zwingen* is a object control verb, thus the subject of *verzichten* is *countries*. The attitudes between the entities is not in the focus of our evaluation. Hence we do not try to predict that *countries*

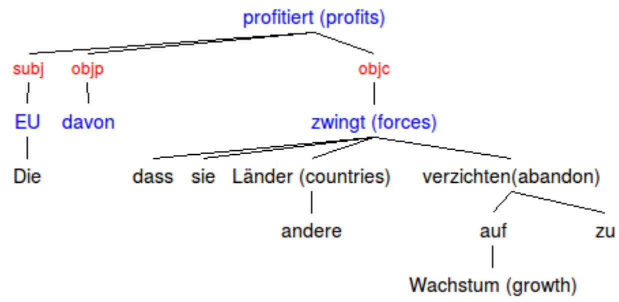

Figure 3: Extraction rule

and *EU* are opponents, so to speak. As we discussed previously, such a perspective could be derived from our resource. However, it would require additional rules of the form: if someone (here *EU*) receives a positive effect from an action that leads to a negative effect for someone else (here *countries*) then both are opponents (at least wrt. the situation at hand).
Given the predicate argument structures of a sentence, the algorithm proceeds outside-in. First, it determines whether the outmost verb assigns its effects or not. This depends on the verb signature and verb negation. Take *to force*: it is factual if it is not negated, and non-factual otherwise. For instance in *The minister does not force the president to cheat* nothing happens at all, neither *minister* nor *president* receive an effect. That is, the main clause (predicate) is non-factual as well as the subclause (predicate). On the other hand, in *The minister forces the president to cheat* both are factual. Thus, if the factuality of the subclause is licenced by the matrix verb, then the algorithm proceeds to determine the status of the subclause. Here again negation has to be coped with, and even further subclause embeddings are possible, in principle. Given *to force him to pretend to tell the truth*. Here, *pretend* is factual, the subject of it gets a negative effect. However, unnegated *pretend* casts non-factuality on its subclause, so *he tells the truth* would not lead to a positive effect (on *he*).

4. Evaluation

The TüBa-D/Z consists of 95,595 sentences. 11,742 sentences contain one or more verbs that are in our model, with a total of 12,662 verb occurrences. This shows that the verbs of our resource have a good coverage given newspaper texts. Not all of these sentences also lead to a verb frame instantiation. There are various reasons: selectional restrictions are not met, no extraction rule was found. Moreover, some verbs behave quite different given the various prepositions that they allow for (*sorgen für*, *sorgen wegen*, *sorgen um*). We have not modelled each and every variant. Neither could we expect to have a complete set of extraction rules, given that we started with 120 sentences forming their basis. So, we ended up with 6,603 instantiated verb frames which are the basis of our evaluation. A stratified sample was generated (500 sentences), two annotators checked whether the assigned effects were correct.
For the evaluation we treated all calculated polarity effects - of which more than one can occur given a specific sentiframe - for the 500 sentences as single instances. The accuracy was 0.76 and 0.80 for the respective annotators. For the annotation, we report a Cohen's κ of 0.87 for the 762 single instances annotated.
A closer error analysis revealed that verb ambiguity is, as expected, an important source of mistakes, but also modality stemming from modal verbs (indicating non-factuality) and negation (there are only a few rules for negation, currently) contribute to the incorrect cases.
We have evaluated our resource on the basis of a plain algorithm (described in section 3.) carrying out sentiment inferences. Our results thus are setting a baseline for further comparison, e.g., with a machine learning approach. At the moment, however, there is not enough training material to reliably learn such a classifier. The classifier would need to decide for each instantiated verb role whether it is positive, negative or neutral (does not get an effect). From the total of 1,023 single verb roles (i.e., the sum of the number of roles of every verb frame that we annotated) 495 are effect roles. That is, 528 roles are not affected, they are neutral, so to speak. A simple baseline that would just assign the polarity of the verb (in this case *to blame* is negative) to every role, thus would produce about 50% false positives. (given *to blame*, the subject should not get a negative effect, only the direct object). In order not to over-generate effects, the sentiframes have proven to be a useful resource that makes use of the selectional restrictions on the one hand. On the other hand, it takes into account what could be inferred on the phrase-level from a bottom-up calculated polarity based on compositionality.
The most frequent verbs for which we have annotated the sentiframes are (English translation in brackets): *unterstützen* (*to support*), *gewinnen* (*to win*), *verlieren* (*to lose*), *ablehnen* (*to decline*), *bedrohen* (*to threaten*), *vermissen* (*to miss*), *schlagen* (*to beat*), *töten* (*to kill*), *feiern* (*to celebrate*), *scheitern* (*to fail*), *verbieten* (*to forbid*), *vorwerfen* (*to reproach*), *leiden* (*to suffer*), *kritisieren* (*to criticize*), *helfen* (*to help*), *empfehlen* (*to promote*), *weigern* (*to refuse*), *zwingen* (*to force*), *stören* (*to disturb*).

5. The Role of Factuality

A major claim of this paper is that factuality plays a crucial role in the determination of polar effects posed by particular verbs. Consider for illustration the simple case of a positive effect that the subject receives given the sentence *The president wins the election*. Various approaches from the literature rather focus on the attitudes between the participants or even those of the author of a text towards the participants (e.g. (Deng and Wiebe, 2014)). They do not care for factuality. But in these scenarios as well, factuality (i.e. truth) plays a role, although not in each and every constellation. Attitudes that stem from intra-clausal dependencies require factuality. For instance, in *The president criticizes the minister* a negative attitude of the president towards the minister holds if the sentence is factual. If the sentence is embedded in a verb casting non-factuality like *to hope*, this no longer is true (cf. *The vice president hopes that the president criticizes the minister*). This sentence, however, also shows that predictions between an actor of a (factual) matrix clause and those of a (non-factual) subclause are possible. Here the vice president has negative attitude towards the minister although the subclause is non-factual. Factuality is crucial, but there are also constellations, where it neither licenses nor prevents attitude predictions.

6. 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 given the internal perspective (the subject's attitude towards the direct object), while for the sentence (as a whole) the judgement is externally negative (here the concepts introduced by the Appraisal theory are used, cf. Martin and White (2005)). The authors do not give any details about how they carry out rule application. Factuality, negation 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. In their approach, if a murderer loses something positive, then this is positive as a whole. While turning to less drastic cases, it becomes less clear how they should be treated (e.g., *the thief who loses all his friends* - should the overall evaluation still be positive? We would say: it is negative for the thief). Recently, Deng and Wiebe (2014) have introduced an ambitious conceptual framework for inferring (sentiment) implicatures. Here, the private state of the author of a sentence is in question. His attitudes towards various targets in the sentence are of interest (and also what he believes the private states of the agents of the sentence would be). Again, factuality and negation do not influence the inference process. How Description Logics can be used in order to identify so-called polarity conflicts is described in Klenner (2015). However, the relations between the referents and again, the factuality of situations are not part of this model.

The focus of this paper lies on the resource we created, not so much on the algorithm to carry out inferences. Except of Choi et al. (2014), no lexical resources have been made available so far. Our resource is the first one (freely available[2]) for German. The only work similar to our work is Ruppenhofer and Brandes (2015). However, this resource is not yet released and, moreover, their annotations are at the sense-level, leaving one with the burden of verb sense disambiguation. We argue that verb ambiguity often disappears if the verbs' subcategorization frames are taken into account together with coarse-grained selectional restrictions (cf. the discussion in section 1 wrt. *care for*).

7. Conclusion

We have introduced a new resource for verb-centered sentiment inference for German. It is the first resource that also takes verb signatures into account. A first evaluation has shown that our word-level verb polarity frames have a good coverage in the domain of newspaper texts (about 12% of the sentences carry at least one verb). An empirical evaluation based on a straightforward inference procedure yielded good results which leaves us confident to apply the resource in scenarios for automated (media) content analysis.

[2]https://pub.cl.uzh.ch/projects/sentiframes

8. Acknowledgements

The authors would like to thank the three reviewers for their comments and suggestions. This work has been carried out in cooperation with Eurospider (http://www.eurospider.com) in a project that was supported by the Swiss Commission for Technology and Innovation (CTI) under the grant 17384.1 PFES-ES.

9. Bibliografical References

Choi, Y., Deng, L., and Wiebe, J. (2014). Lexical acquisition for opinion inference: A sense-level lexicon of benefactive and malefactive events. In *Proc. of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, Baltimore, Maryland. Association for Computational Linguistics.

Deng, L. and Wiebe, J. (2014). Sentiment propagation via implicature constraints. *Meeting of the European Chapter of the Association for Computational Linguistics (EACL-2014).*

Klenner, M., Amsler, M., and Hollenstein, N. (2014a). Verb polarity frames: a new resource and its application in target-specific polarity classification. In *KONVENS*, pages 106–115.

Klenner, M., Tron, S., Amsler, M., and Hollenstein, N. (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.

Klenner, M. (2015). Verb-centered sentiment inference with description logics. In *Proceed. of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 134–139, Lisboa, Portugal. Association for Computational Linguistics.

Martin, J. R. and White, P. R. R. (2005). *Appraisal in English*. Palgrave, London.

Nairn, R., Condoravdi, C., and Karttunen, L. (2006). Computing relative polarity for textual inference. In *Proceed. of Inference in Computational Semantics (ICoS 5), Buxton, England*, pages 67–75.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. (2009). Semantically distinct verb classes involved in sentiment analysis. In *IADIS AC (1)*, pages 27–35.

Reschke, K. and Anand, P. (2011). Extracting contextual evaluativity. In *Proc. of the Ninth International Conference on Computational Semantics*, pages 370–374.

Ruppenhofer, J. and Brandes, J. (2015). Extending effect annotation with lexical decomposition. In *Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 67–76, Lisboa, Portugal, September. Association for Computational Linguistics.

Sennrich, R., Schneider, G., Volk, M., and Warin, M. (2009). A new hybrid dependency parser for German. In *Proc. of the German Society for Computational Linguistics and Language Technology*, pages 115–124.

Telljohann, H., Hinrichs, E. W., Kübler, S., Zinsmeister, H., and Beck, K. (2009). Stylebook for the Tübingen treebank of written German (TüBa-D/Z). Technical report, Universität Tübingen, Seminar für Sprachwissenschaft.