



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2017

An Unification-based Model for Attitude Prediction

Klenner, Manfred

DOI: <https://doi.org/10.1007/978-3-319-64206-2>

Posted at the Zurich Open Repository and Archive, University of Zurich
ZORA URL: <https://doi.org/10.5167/uzh-139437>
Conference or Workshop Item

Originally published at:

Klenner, Manfred (2017). An Unification-based Model for Attitude Prediction. In: 20th International Conference on Text, Speech, and Dialogue (TSD), Prague, Czechia, 27 August 2017 - 31 August 2017, Springer.
DOI: <https://doi.org/10.1007/978-3-319-64206-2>

An Unification-based Model for Attitude Prediction

Manfred Klenner

Computational Linguistics, University of Zurich

www.cl.uzh.ch

klenner@cl.uzh.ch

Abstract. Attitude prediction strives to determine whether an opinion holder is positive or negative towards a given target. We cast this problem as a lexicon engineering task in the context of deep linguistic grammar formalisms such as LFG or HPSG. Moreover, we demonstrate that attitude prediction can be accomplished solely through unification of lexical feature structures. It is thus possible to use our model without altering existing grammars, only the lexicon needs to be adapted. In this paper, we also show how our model can be combined with dependency parsers. This makes our model independent of the availability of deep grammars, only unification as a processing mean is needed.

Keywords: Sentiment, Opinion Inference, Lexical Functional Grammar

1 Introduction

Attitude prediction comprises the identification of an opinion holder, an opinion target and the positive or negative attitude of the holder towards the target. It is a variant of stance detection where the targets are not known in advance and the writer is not necessarily the only opinion holder. Take *Peres accused Syria to support Hezbollah*. Here, the writer claims that Peres (an opinion holder) has a negative attitude towards Syria (the target). However, the sentence also implies that Peres is against the Hezbollah (another target). Moreover, Hezbollah is the (potential) beneficiary of a support event which - as the reason of the accusation - is (contextually) perceived as being negative. Being against a event means that one is also against any beneficiaries of that event. Correspondingly, one has a positive attitude towards victims of any disapproved event (e.g. in *A0 complained that A1 was hurt*, A0 is a proponent of the victim A1). As discussed in [7], the truth commitment that comes with particular verbs, its interaction with negation in the course of the determination of event factuality [9] are crucial components of any attitude prediction. So far, attitude prediction has been cast as logic-based inference, either with machine learning (Probabilistic Soft Logic, see [2]) or Description Logics (cf. [7]). In this paper, we introduce a lean model for attitude prediction. Instead of deduction, unification is used and instead of rules, verb classes are specified on the basis of attribute value pairs (features). Our model can easily be combined with existing deep grammars such as the Pargram LFG grammars [1]. Only the lexicon needs to be adapted leaving the grammar rules as they are. However, our model is independent of the existence of such grammars. We only need unification, not unification grammars. Actually, we show how to combine our model with dependency parsers. This makes our approach widely applicable. Our current model is for German.

2 Unification-based Grammars

The core of unification or constraint-based grammars are feature structures (f-structure, henceforth) and unification. A feature structure is a list of attribute value pairs, where the attributes are, among others, names of grammatical functions or morphological categories. Values are atomic or complex, i.e. feature structures again. Unification of two feature structures is successful if the unification of corresponding attributes is successful. Attribute values unify if both are atomic and identical or, in the case of complex values, both recursively unify. We work with LFG [3] and the Xerox Linguistic Environment (XLE) grammar engineering tool [4]. In addition to the standard unification principle, LFG requires feature structures to be coherent (a governable function must have a governor) and complete (any governed function must be realized). In XLE so-called templates are available. A template is a means to build classes and to save specification effort (which is interesting from an engineering perspective).

3 Attitudes and Polar Effects

If we assume that the writer is committed to the truth of his text, we can infer the attitudes among the entities referred to, but we also are able to find out whether the scenario is good or bad *of* or *for* a referent. If *A0 injures A1* this is bad *of* A0 and bad *for* A1. We call these resulting states a predicate produces the *polar effects* and use *polar roles* like *PFOR* (positive for) to denote that a referent occupies such a role. Polar effects are crucial, they allow to understand the role a referent plays in a text. They allow to understand how an entity is perceived (or cast) by the text author. For instance, we found that the German Chancellor, Angela Merkel, takes quite often the role of a negative actor (a *NOF* (negative of) role) given the texts of the AfD, a German right-wing party.

Obviously verbs are most crucial for these tasks. For instance, we know that *cheat* means that A0, the actor, acts in a negative way and thus should be regarded as a villain, while A1, the patient, is a victim. If we know the polar effect of a verb (on the target), we know the attitude (of the opinion holder towards the target). The (situation-specific) attitude of A0 expressed by *cheat* towards A1 is negative. The point is that these polar effects and - to a certain degree also the attitudes - depend on the factuality status of the sentence.

We refer to the concept of event factuality as discussed in [9]. If an event is factual, polar effects take place, if an event is counterfactual, the inverted effects might (depending on the verb) take place. If the event is nonfactual, no effect is cast. The truth of *Italy has helped the migrants to survive* makes *Italy* a benefactor and the *migrants* a beneficiary. The negated form, *Italy not helped the migrants to survive*, turns *Italy* into a villain and the *migrants* into a victim. Finally the modal, nonfactuality indicating version *Italy might help the migrants to survive* blocks any inferences.

Such verbs do have a truth commitment wrt. their complement clause. An affirmative and truth committed subclause gives rise to event factuality. Thus *cheat* denotes a true event given that *A0 regrets that he cheated A1* is true. If the subclause is negated, the truth commitment makes the event counterfactual: *A0 regrets that he not has told*

the truth. In this case, the inverse event is true (*he lied*). We also have to take the affirmative status of the matrix verb into account, which might alter the truth commitment. According to [7], we need to distinguish for each verb a signature for the affirmative and for the negated usage (see also [5]). We propose to have three categories: 'T' (truth committing), 'F' (falsehood committing) and 'N' (no commitment). For instance, *regret* as a factive verb has 'T'-'T' (affirmative-negated), while *force* has 'T'-'N' and *refuse* has 'F'-'N'. *A0 refuses to cheat* means that *cheat* is counterfactual, while in *A0 not refuses to cheat* we cannot tell whether there will be a cheating event. Finally, non-factive verb like *hope* have 'N'-'N'. It crucial not to confuse truth commitment and factuality. The first one holds for the whole (even negated) subclause, while the second relates to the event expressed by the verb.

4 Unification-based Attitude Prediction and Role Assignment

We have specified a reasoning scheme on the basis of unification. Constraining equality and existential constraints are building blocks for this. Constraining equality (notation =c) is a global constraint mechanism that demands that a particular value is introduced by a defining equation (notation =) somewhere in the lexicon or grammar. Existential constraints pose the restriction that a particular attribute must or must not be present.

Figure 1 shows the f-structures of two sentences: *Peres accuses Syria to support Hezbollah* (left-hand side) and *Peres regrets that Syria supports Hezbollah* (right-hand side). The predicted attitudes are captured under the attribute REL (either *pro* or *con*), the polar roles here are *PFOR* and *POF* (in the *regret* sentence). For instance, the *accuse* version gives rise to the negative attitude of Peres towards Syria, represented by (dropping the grammatical functions) *con*<*Peres*,*Syria*> in Fig.1. Only in the *regret* version, we find a *pro* relation, *pro*<*Syria*,*Hezbollah*> (since *support* then is factual) and a *POF* filled by *Syria* and a *PFOR* role filled by *Hezbollah*.

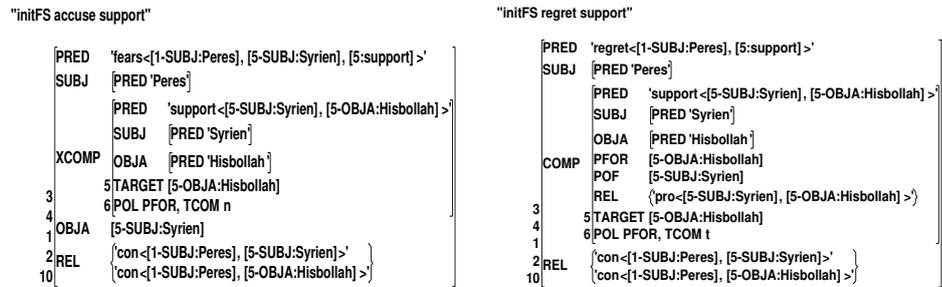


Fig. 1. F-Structures for Peres accuse/regret Syria support Hezbollah

A verb subcategorizing for a complement clause gets a verb signature ('T', 'F' or 'N') depending on the affirmative status (affirmative or negated) of the verb. We introduce an attribute *TCOM* and embed it under the subcategorized clausal complement

(either COMP or XCOMP). In Fig.1 *support* has *TCOM N* in the case of *accuse* while *TCOM T* with *regret*. Fig.2 shows the (partial) entry for *beschuldigen (accuse)*, a verb that has signature 'T-N'. We use XLE notation: verb form followed by part of speech,

```

beschuldigen (accuse) V *
1      (↑PRED)=’accuse<(↑SUBJ)(↑OBJA)(↑XCOMP)>’
2      (↑XCOMP SUBJ)=(↑OBJA)
3      {¬(↑NEG) (↑XCOMP TCOM ) = T | (↑NEG) (↑XCOMP TCOM ) = N}
4      @nforverb OBJA)
5      @(disapprove XCOMP)

```

Fig. 2. Lexical Entry for the Verb *beschuldigen (accuse)*

the morphology indicator * and the feature equations¹. Line 1 gives the semantic form of the verb (with the subcategorized grammatical functions). Line 2 is a control equation (*beschuldigen* has object control). Line 3 is the disjunctive condition for truth commitment (realization of 'T-N'). The attribute *NEG* exists if the verb is negated, otherwise the attribute is not present at all. The existential constraint (↑NEG) thus checks whether the clause is negated. It is a kind of precondition for the assignment of the subclause's truth commitment: 'T' if affirmative (¬(↑NEG)), 'N' if negated (↑NEG).

Lines 4 and 5 are template invocations that relate to the two main tasks, the prediction of polar effects and attitudes. Most of the time, the subject is the opinion source. The Target role is more flexible, either *OBJA*, *OBJD* or *OBJP* might occupy it (depending on the verb). We use the polar roles *PFOR* (positive for), *NFOR* (negative for), *POF* (positive of) and *NOF* (negative of) to model the verb-specific polar effects. For instance, *criticize* has a *PFOR* role (if factual) and the direct object *OBJA* as the target. Given the verb *enjoy*, it is the subject that takes the *PFOR* role. These polar roles are, thus, abstract semantic roles with a polarity load. The template *@nforverb* (line 4 from Fig.5) is invoked with the target role, *ROLE*, which for *accuse* is *OBJA*.

Before we introduce *@nforverb* we briefly discuss (the template for) factuality:

```

factual =
{¬(↑NEG) (↑TCOM)=c T | (↑NEG) (↑TCOM)=c F | ¬(COMP↑) ¬(XCOMP↑) ¬(↑MOD) ¬(↑NEG)}.

```

Factuality holds, if the verb is affirmative and has T as signature, or if negated and the signature is F, or if it is not embedded (inside-out determination) and there is neither modality nor negation present. ¬(COMP↑) prohibits (inside-out) embedding under COMP. This realises factuality determination of the outmost matrix verb.

Fig. 3 shows the definition of (the verb class) *@nforverb*. *@nforverb* and *@pforverb* carry out the assignment of polar roles under factuality and counterfactuality. They also establish the *inner* attitude prediction, i.e. in cases where source and target have the same verbal head (the templates *@direct_con* and *@direct_pro* from Fig.3). Line 1 from Fig.3 sets the target role. The rest of the definition depends on the factuality status of the verb. If the verb is factual (line 2) then the target role is set as *NFOR*. If

¹ An up arrow inserts a feature into the feature structure defined by the equation.

the verb is counterfactual (line 3) the polar role is inverted, e.g. set to *PFOR*. The verb

```

nforverb(ROLE) =
1      (↑TARGET)=(↑ROLE)
2      { @ (factual)      (↑NFOR)=(↑ROLE) @direct_con (↑NFORView)=+
3      | @ (counterfactual) (↑PFOR)=(↑ROLE) @direct_pro (↑PFORView)=+
4      | @nonfactual
5      | { (↑NEG)(↑PFORView)=+ | (↑NEG) (↑NFORView)=+ } }.

```

Fig. 3. Template Definition

criticize, e.g., is a *@nforverb* verb. If A0 criticizes A1, this is negative for A1 (*NFOR*) and the attitude of A0 towards A1 is negative (*@direct_con*). If *criticize* is nonfactual (line 4 and 5), e.g. embedded into *hope*, then no polar role is set. However, in order to determine the *outer* attitude, namely the one of the opinion source of the matrix clause towards the referents of the embedded verb in nonfactual cases, we need to know the polar role profile of the verb. Is the target someone who would benefit (*PFOR*) or suffer (*NFOR*) from a situation where the event denoted by the verb was true? To provide this information is the function of the attributes *NFORView* and *PFORView*: they define the role profile without instantiating roles.

Fig.4 shows the definition of the *@disapprove* template responsible for attitudes between the matrix clause and the subclause. Again, the factuality status is crucial, but this time it is the one of the matrix clause.

```

disapprove(C) =
1      { @factual @ (subj_con C)
2      | { (↑C PFORView) 'con<(↑SUBJ)(↑C TARGET)>' ∈ (↑REL)
3      | (↑C NFORView) 'pro<(↑SUBJ)(↑C TARGET)>' ∈ (↑REL) | @ (NOView C) }
4      | @counterfactual @ (subj_pro C)
5      | { (↑C PFORView) 'pro<(↑SUBJ)(↑C TARGET)>' ∈ (↑REL)
6      | (↑C NFORView) 'con<(↑SUBJ)(↑C TARGET)>' ∈ (↑REL) | @ (NOView C) }
7      | @nonfactual }.

```

Fig. 4. Template Definition

First of all, given a verb like *accuse* of type *@disapprove*, the opinion holder of the matrix clause is against the subject of the subclause if the matrix clause is factual: this is the function of the template call *@(subj_con C)*, line 1. The variable C indicates the subclause type (XCOMP, COMP). If the subclause obeys to a *PFORView* (*accuse that A0 has helped A1*) (see line 2) then a *con* relation is set. *REL* is an attribute that takes a set as its value (since the opinion source might have more than one attitude): \in is used, thus, instead of $(\uparrow=\downarrow)$. If the matrix is factual but the subclause turns out to have a *NFORView* (*accuse that A0 not has helped A1*) then the relation is *pro* (line 3).

The corresponding definition holds for counterfactuality (see line 4-6). In case that the subclause is a verb without a polar viewpoint, *@NOView* applies (it just verifies that the verb has no polar view).

In our model, direct inferences are restricted to a single level of embedding. We claim that this is sufficient. Inferences for deeper nested structures like in the sentence *X criticizes that A0 has not helped to free A1* where *X* has a positive attitude towards *A1* can be drawn with transitive rules like: $\text{con}\langle X, A0 \rangle \wedge \text{con}\langle A0, A1 \rangle \rightarrow \text{pro}\langle X, A1 \rangle$.

5 Model Initialisation from a Dependency Parse

We used the verb resource² of [6] in our implementation. It was automatically mapped to XLE specifications like the one shown in Fig.2.

Our model for attitude prediction is purely lexicalistic. A lexical entry of a verb fully specifies its behavior either as an embedding or embedded verb in a simple or complex sentence. It thus can be combined with any existing deep linguistic grammar. But we also can combine it with dependency parsers. We only need to determine the grammatical functions of the involved verbs at the right embedding level. This information is available from the dependency tree. Figure 5 shows an example. The right-hand side is

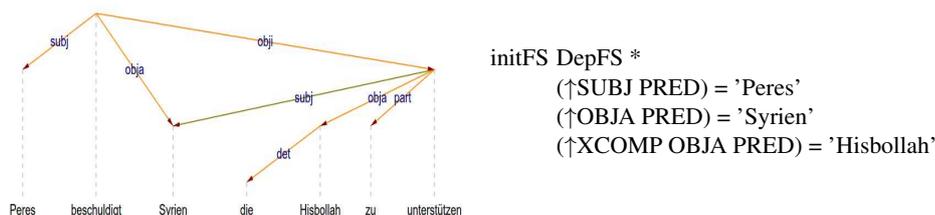


Fig. 5. Peres beschuldigt (accuses) Syrien die Hisbollah zu unterstützen (to support)

the result of the mapping: a feature structure of type *DepFS* with one embedding level. All information stems from the dependency tree. The label *obji* is mapped to *XCOMP*. We need a single XLE grammar rule (slightly simplified) in order to parse this :

$G \rightarrow \text{DepFS } V \ V: \{(\uparrow\text{COMP})=\downarrow \mid (\uparrow\text{XCOMP})=\downarrow\}$.

The expression to be parsed is: *parse {initFS accuse support}* (in general: *initFS matrix_verb subclause_verb1 subclause_verb2 . . .*). First *initFS* is identified as *DepFS*, its f-structure is unified with those of the matrix verb (*accuse*). The subclause verb is unified with the result, its f-structure gets embedded under *XCOMP*. See Fig.1 for the resulting f-structure.

6 Empirical Evaluation

It comes as a surprise that the interannotator agreement for the task at hand is low. In the empirical evaluation reported in [7] comprising 160 sentences it is 43% (Cohens κ

² Available from https://pub.cl.uzh.ch/projects/opinion/lrec_data.txt

= 0.19). One of the reasons is that humans seem to be “selective annotators and focus on the most striking attitudes more than on the more hidden ones” (see [7] page 83). They also report that the two annotators often produced different and - as it turned out - complementary annotations for a given sentence. Thus, reconciliation was unproblematic, i.e. both accepted their respective additional attitude annotations. This indicates that attentiveness might be a problem, but also that factors might play a role that are beyond the verb semantics, e.g. stemming from world knowledge.

We used the data from [7]. Note that our two models are similar³, but our model is leaner and realized with a totally different framework. Also, our notion of factuality differs (see Related Work). However, our attitude labels are easily mapped onto theirs.

Their data consists of a) 80 complex made-up sentence (their precision was 83.89%, recall is 93.72%) and b) 80 newspaper sentences (59.04% precision and 71.15% recall). Our system achieved 85.12%, recall 91.52% for a) and 65.24% precision and 75.13% recall for b). It is obvious that this is but a first evaluation. A larger data set is needed - which is not available yet. We also need to clarify how humans actually perform attitude prediction. This is future work.

7 Related Work

The goal of the rule-based approach of [2] is to detect entities that are in a positive (PosPair) or negative (NegPair) relation to each other. Rules are realized in the framework of Probabilistic Soft Logic, where the rule weights depend on the output of the preprocessing pipeline made out of two SVM classifiers and three existing sentiment analysis systems. The model of [2] also copes with event-level sentiment inference, however factuality is not taken into account. Also, polar roles do not play any role in their framework.

[7] stress the point that factuality determination is a crucial part of sentiment inferences. They introduce a rule-based system for German realized with Description Logic and SWRL. The rules also are taking the affirmative and factuality status of the sentence into account. The goal is to instantiate relations (con and pro) expressing the attitudes of entities towards each other. We agree that factuality is a crucial part of such a model. However, we use a tripartite distinction while their factuality labels are binary.

Recently, [8] have presented an elaborate model that is meant to explicate the relations between all involved entities: the reader, the writer, and the entities referred to by a sentence. Also, the internal states of the referents and their values are part of the model. The underlying resource, called connotation frames, was created in a crowd sourcing experiment, the model parameters (e.g. values for positive and negative scores) are average values. Our verb resource is, on the contrary, specified by experts. Again, factuality is not taken into account in their model.

8 Conclusion

In this paper, a purely unification-based approach for sentiment reasoning is introduced. The approach is independent of any existing deep linguistic grammar, but can be cou-

³ We also use the parser in [10].

pled easily with it. Only the verb lexicon needs to be augmented with additional verb-specific features. This would result in a system that carries out attitude prediction etc. while parsing instead of afterwards (like current systems do). In the current paper, we pursued another possibility. Namely, to couple the model with a dependency parser. Only the embedding skeleton needs to be derived from the dependency tree. Then unification with verb entries carries out the whole inference process. To the best of our knowledge, our work is the first that exploits the idea of feature unification on top of a dependency parse tree in order to solve a sophisticated problem. Feature structures and unification are an elegant representational scheme and provide powerful processing means. We have shown how to reap the benefits of this. This might stimulate other researcher to also pose their problems in terms of such a framework.

Our system realizes a linguistically informed approach to solve the problem of attitude prediction and the assignment of polar roles. Future work will focus on a broader evaluation in the context of stance detection.

References

1. Butt, M., King, T.H., Masuichi, H., Rohrer, C.: The parallel grammar project. In: J. Carroll, N., Sutcliffe, R. (eds.) *Proceedings of the Workshop on Grammar Engineering and Evaluation*. pp. 1–7 (2002), *cOLING02*
2. Deng, L., Wiebe, J.: 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)*. pp. 179–189. Lisbon, Portugal (2015)
3. Kaplan, R.M., Bresnan, J.: *Lexical-Functional Grammar: A formal system for grammatical representation*. In: *The Mental Representation of Grammatical Relations*. The MIT Press, Cambridge, MA (1982)
4. Kaplan, R.M., King, T.H., III, J.T.M.: Adapting existing grammars: The xle approach. In: J. Carroll, N., Sutcliffe, R. (eds.) *Proceedings of the Workshop on Grammar Engineering and Evaluation*. pp. 29–35 (2002), *cOLING02*
5. Karttunen, L.: Simple and phrasal implicatives. In: *Proceedings of the First Joint Conference on Lexical and Computational Semantics*. pp. 124–131. Association for Computational Linguistics, Stroudsburg, PA, USA (2012)
6. Klenner, M., Amsler, M.: Sentiframes: A resource for verb-centered German sentiment inference. In: Chair, N.C.C., Choukri, K., Declerck, T., Goggi, S., Grobelnik, M., Maegaard, B., Mariani, J., Mazo, H., Moreno, A., Odijk, J., Piperidis, S. (eds.) *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC)*. pp. 2888–2891. Portoro, Slovenia (2016)
7. Klenner, M., Cematide, S.: How factuality determines sentiment inferences. In: Claire Gardent, Raffaella Bernardi, I.T. (ed.) *Proceedings of *SEM 2016: The Fifth Joint Conference on Lexical and Computational Semantics*. pp. 75–84. Berlin, Germany (August 2016)
8. Rashkin, H., Singh, S., Choi, Y.: Connotation frames: A data-driven investigation. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*. pp. 311–321. Berlin, Germany (August 2016)
9. Saurí, R., Pustejovsky, J.: FactBank: a corpus annotated with event factuality. *Language Resources and Evaluation* 43(3), 227–268 (2009)
10. Sennrich, R., Schneider, G., Volk, M., Warin, M.: A new hybrid dependency parser for German. In: *Proceedings of the German Society for Computational Linguistics and Language Technology (GSCL)*. pp. 115–124. Potsdam, Germany (2009)