

Semantic frames as an anchor representation for sentiment analysis

Josef Ruppenhofer

Department of Information Science
and Natural Language Processing
University of Hildesheim, Germany
ruppenho@uni-hildesheim.de

Ines Rehbein

SFB 632: Information Structure
German Department
Potsdam University, Germany
irehbein@uni-potsdam.de

Abstract

Current work on sentiment analysis is characterized by approaches with a pragmatic focus, which use shallow techniques in the interest of robustness but often rely on ad-hoc creation of data sets and methods. We argue that progress towards deep analysis depends on a) enriching shallow representations with linguistically motivated, rich information, and b) focussing different branches of research and combining resources to create synergies with related work in NLP. In the paper, we propose SentiFrameNet, an extension to FrameNet, as a novel representation for sentiment analysis that is tailored to these aims.

1 Introduction

Sentiment analysis has made a lot of progress on more coarse-grained analysis levels using shallow techniques. However, recent years have seen a trend towards more fine-grained and ambitious analyses requiring more linguistic knowledge and more complex statistical models. Recent work has tried to produce relatively detailed summaries of opinions expressed in news texts (Stoyanov and Cardie, 2011); to assess the impact of quotations from business leaders on stock prices (Drury et al., 2011); to detect implicit sentiment (Balahur et al., 2011); etc. Accordingly, we can expect that greater demands will be made on the amount of linguistic knowledge, its representation, and the evaluation of systems.

Against this background, we argue that it is worthwhile to complement the existing shallow and pragmatic approaches with a deep, lexical-semantics based one in order to enable deeper analysis. We report on ongoing work in constructing Sen-

tiFrameNet, an extension of FrameNet (Baker et al., 1998) offering a novel representation for sentiment analysis based on frame semantics.

2 Shallow and pragmatic approaches

Current approaches to sentiment analysis are mainly pragmatically oriented, without giving equal weight to semantics. One aspect concerns the identification of sentiment-bearing expressions. The annotations in the MPQA corpus (Wiebe et al., 2005), for instance, were created without limiting what annotators can annotate in terms of syntax or lexicon. While this serves the spirit of discovering the variety of opinion expressions in actual contexts, it makes it difficult to match opinion expressions when using the corpus as an evaluation dataset as the same or similar structures may be treated differently. A similar challenge lies in distinguishing so-called polar facts from inherently sentiment-bearing expressions. For example, out of context, one would not associate any of the words in the sentence *Wages are high in Switzerland* with a particular evaluative meaning. In specific contexts, however, we may take the sentence as reason to either think positively or negatively of Switzerland: employees receiving wages may be drawn to Switzerland, while employers paying wages may view this state of affairs negatively. As shown by the inter-annotator agreement results reported by (Toprak et al., 2010), agreement on distinguishing polar facts from inherently evaluative language is low. Unsurprisingly, many efforts at automatically building up sentiment lexica simply harvest expressions that frequently occur as part of polar facts without resolving whether the subjectivity clues extracted are inherently evaluative or

merely associated with statements of polar fact.

Pragmatic considerations also lead to certain expressions of sentiment or opinion being excluded from analysis. (Seki, 2007), for instance, annotated sentences as “not opinionated” if they contain indirect hearsay evidence or widely held opinions.

In the case of targets, the work by (Stoyanov and Cardie, 2008) exhibits a pragmatic focus as well. These authors distinguish between (a) the **topic** of a fine-grained opinion, defined as the real-world object, event or abstract entity that is the subject of the opinion as *intended* by the opinion holder; (b) the **topic span** associated with an opinion expression is the closest, minimal span of text that mentions the topic; and (c) the **target span** defined as the span of text that covers the syntactic surface form comprising the contents of the opinion. As the definitions show, (Stoyanov and Cardie, 2008) focus on text-level, pragmatic relevance by paying attention to what the author intends, rather than concentrating on the explicit syntactic dependent (their target span) as the topic. This pragmatic focus is also in evidence in (Wilson, 2008)’s work on contextual polarity classification, which uses features in the classification that are syntactically independent of the opinion expression such as the number of subjectivity clues in adjoining sentences.

Among lexicon-driven approaches, we find that despite arguments that word sense distinctions are important to sentiment analysis (Wiebe and Mihalcea, 2006), often-used resources do not take them into account and new resources are still being created which operate on the more shallow lemma-level (e.g. (Neviarouskaya et al., 2009)). Further, most lexical resources do not adequately represent cases where multiple opinions are tied to one expression and where presuppositions and temporal structure come into play. An example is the verb *despoil*: there is a positive opinion by the reporter about the despoiled entity in its former state, a negative opinion about its present state, and (inferrable) negative sentiment towards the despoiler. In most resources, the positive opinion will not be represented.

The most common approach to the task is an information extraction-like pipeline. Expressions of opinion, sources and targets are often dealt with separately, possibly using separate resources. Some work such as (Kim and Hovy, 2006) has explored

the connection to role labeling. One reason not to pursue this is that “in many practical situations, the annotation beyond opinion holder labeling is too expensive” (Wiegand, 2010, p.121). (Shaikh et al., 2007) use semantic dependencies and composition rules for sentence-level sentiment scoring but do not deal with source and target extraction. The focus on robust partial solutions, however, prevents the creation of an integrated high-quality resource.

3 The extended frame-semantic approach

We now sketch a view of sentiment analysis on the basis of an appropriately extended model of frame semantic representation.¹

Link to semantic frames and roles Since the possible sources and targets of opinion are usually identical to a predicate’s semantic roles, we add *opinion frames* with slots for Source, Target, Polarity and Intensity to the FrameNet database. We map the Source and Target opinion roles to semantic roles as appropriate, which enables us to use semantic role labeling systems in the identification of opinion roles (Ruppenhofer et al., 2008).

In SentiFrameNet all lexical units (LUs) that are inherently evaluative are associated with *opinion frames*. The language of polar facts is not associated with opinion frames. However, we show in the longer version of this paper (cf. footnote 1) how we support certain types of inferred sentiment. With regard to targets, our representation selects as targets of opinion the target spans of (Stoyanov and Cardie, 2008) rather than their opinion topics (see Section 2). For us, opinion topics that do not coincide with target spans are inferential opinion targets.

Formal diversity of opinion expressions For fine-grained sentiment-analysis, handling the full variety of opinion expressions is indispensable. While adjectives in particular have often been found to be very useful cues for automatic sentiment analysis (Wiebe, 2000; Benamara et al., 2007), evaluative meaning pervades all major lexical classes. There are many subjective multi-words and idioms such as *give away the store* and evaluative meaning also attaches to grammatical constructions, even ones without obligatory lexical material. An exam-

¹We present a fuller account of our ideas in an unpublished longer version of this paper, available from the authors’ websites.

ple is the construction exemplified by *Him be a doctor?* The so-called *What, me worry?*-construction (Fillmore, 1989) consists only of an NP and an infinitive phrase. Its rhetorical effect is to express the speaker's surprise or incredulity about the proposition under consideration. The FrameNet database schema accommodates not only single and multi-words but also handles data for a construction (Fillmore et al., to appear) that pairs grammatical constructions with meanings.

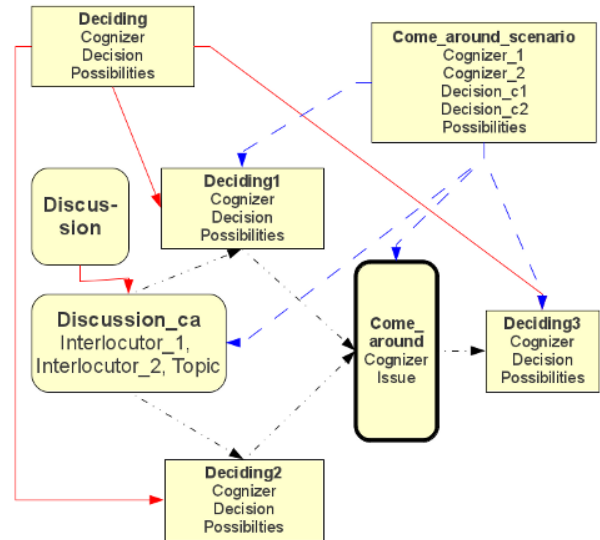
Multiple opinions We need to accommodate multiple opinions relating to the same predicate as in the case of *despoil* mentioned above. Predicates with multiple opinions are not uncommon: in a 100-item random sample taken from the Pittsburgh subjectivity clues, 17 involved multiple opinions.

The use of opinion frames as described above enables us to readily represent multiple opinions. For instance, the verb *brag* in the modified *Bragging* frame has two opinion frames. The first one has positive polarity and represents the frame-internal point of view. The *SPEAKER* is the Source relative to the *TOPIC* as the Target. The second opinion frame has negative polarity, representing the reporter's point of view. The *SPEAKER* is the Target but the Source is unspecified, indicating that it needs to be resolved to an embedded source. For a similar representation of multiple opinions in a Dutch lexical resource, see (Maks and Vossen, 2011).

Event structure and presuppositions A complete representation of subjectivity needs to include event and presuppositional structure. This is necessary, for instance, for predicates like *come around (on)* in (1), which involve changes of opinion relative to the *same* target by the *same* source. Without the possibility of distinguishing between attitudes held at different times, the sentiment associated with these predicates cannot be modeled adequately.

- (1) Newsom is still against extending weekday metering to evenings, but has COME AROUND on Sunday enforcement.

For *come around (on)*, we want to distinguish its semantics from that of predicates such as *ambivalent* and *conflicted*, where a *COGNIZER* *simultaneously* holds opposing valuations of (aspects of) a target. Following FrameNet's practice, we model presupposed knowledge explicitly in SentiFrameNet by



Constraints

- Come_around.Cognizer=Deciding1.Cognizer
- Come_around.Cognizer=Discussion.Interlocutor_1
- Deciding2.Cognizer=Discussion.Interlocutor_2
- Deciding1.Decision!=Deciding2.Decision
- Deciding1.Possibilities=Deciding2.Possibilities
- Discussion.Topic=Come_around.Issue
- Deciding3.Cognizer=Deciding1.Cognizer
- Deciding3.Decision=Deciding2.Decision
- Deciding3.Possibilities=Deciding1.Possibilities
- ...

Figure 1: Frame analysis for "Come around"

using additional frames and frame relations. A partial analysis of *come around* is sketched in Figure 1.

We use the newly added *Come around scenario* frame as a background frame that ties together all the information we have about instances of coming around. Indicated by the dashed lines are the SUB-FRAMES of the scenario. Among them are three instances of the *Deciding* frame (solid lines), all related temporally (dashed-dotted) and in terms of content to an ongoing *Discussion*. The initial difference of opinion is encoded by the fact that *Deciding1* and *Deciding2* share the same POSSIBILITIES but differ in the DECISION. The occurrence of *Come_around* leads to *Deciding3*, which has the same COGNIZER as *Deciding1* but its DECISION is now identical to that in *Deciding2*, which has been unchanged. The sentiment information we need is encoded by simply stating that there is a sentiment of positive polarity of the COGNIZER (as source) towards the DECISION (as target) in the *Deciding* frame. (This opinion frame is not displayed in the graphic.) The *Come around* frame itself is not as-

sociated with sentiment information, which seems right given that it does not include a DECISION as a frame element but only includes the ISSUE.

For a discussion of how SentiFrameNet captures factuality presuppositions by building on (Saurí, 2008)’s work on event factuality, we refer the interested reader to the longer version of the paper.

Modulation, coercion and composition Speakers can shift the valence or polarity of sentiment-bearing expressions through some kind of negation operator, or intensify or attenuate the impact of an expression. Despite these interacting influences, it is desirable to have at least a partial ordering among predicates related to the same semantic scale; we want to be able to find out from our resource that *good* is less positive than *excellent*, while there may be no ordering between *terrific* and *excellent*. In SentiFrameNet, an ordering between the polarity strength values of different lexical units is added on the level of frames.

The frame semantic approach also offers new perspectives on sentiment composition. We can, for instance, recognize cases of presupposed sentiment, as in the case of the noun *revenge*, which are not amenable to shifting by negation: *She did not take revenge* does not imply that there is no negative evaluation of some injury inflicted by an offender.

Further, many cases of what has been called valence shifting for us are cases where the evaluation is wholly contained in a predicate.

- (2) Just barely AVOIDED an accident today.
- (3) I had served the bank for 22 years and had AVOIDED a promotion since I feared that I would be transferred out of Chennai city.

If we viewed *avoid* as a polarity shifter and further treated nouns like *promotion* and *accident* as sentiment-bearing (rather than treating them as denoting events that affect somebody positively or negatively) we should expect that while (2) has positive sentiment, (3) has negative sentiment. But that is not so: accomplished intentional avoiding is always positive for the avoider. Also, the reversal analysis for *avoid* cannot deal with complements that have no inherent polarity. It readily follows from the coercion analysis that *I avoid running into her* is negative but that cannot be derived in e.g. (Moilanen and Pulman, 2007)’s compositional model which takes into account inherent lexical polarity, which *run (into)*

lacks. The fact that *avoid* imposes a negative evaluation by its subject on its object can easily be modeled using opinion frames.

4 Impact and Conclusions

Deep analysis Tying sentiment analysis to frame semantics enables immediate access to a deeper lexical semantics. Given particular application-interests, for instance, identifying statements of uncertainty, frames and lexical units relevant to the task can be pulled out easily from the general resource. A frame-based treatment also improves over resources such as SentiWordNet (Baccianella et al., 2008), which, while representing word meanings, lacks any representation of semantic roles.

Theoretical insights New research questions await, among them: whether predicates with multiple opinions can be distinguished automatically from ones with only one, and whether predicates carrying factivity or other sentiment-related presuppositions can be discovered automatically. Further, our approach lets us ask how contextual sentiment is, and how much of the analysis of pragmatic annotations can be derived from lexical and syntactic knowledge.

Evaluation With a frame-based representation, the units of annotation are pre-defined by a general frame semantic inventory and systems can readily know what kind of units to target as potential opinion-bearing expressions. Once inherent semantics and pragmatics are distinguished, the correctness of inferred (pragmatic) targets and the polarity towards them can be weighted differently from that of immediate (semantic) targets and their polarity.

Synergy On our approach, lexically inherent sentiment information need not be annotated, it can be imported automatically once the semantic frame’s roles are annotated. Only pragmatic information needs to be labeled manually. By expanding the FrameNet inventory and creating annotations, we improve a lexical resource and create role-semantic annotations as well as doing sentiment analysis.

We have proposed SentiFrameNet as a linguistically sound, deep representation for sentiment analysis, extending an existing resource. Our approach complements pragmatic approaches, allows us to join forces with related work in NLP (e.g. role labeling, event factuality) and enables new insights into the theoretical foundations of sentiment analysis.

References

- S. Baccianella, A. Esuli, and F. Sebastiani. 2008. SENTIWORDNET 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the Seventh conference on International Language Resources and Evaluation LREC10*, pages 2200–2204. European Language Resources Association (ELRA).
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley Framenet Project. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics-Volume 1*, pages 86–90. Association for Computational Linguistics.
- Alexandra Balahur, Jesús M. Hermida, and Andrés Montoyo. 2011. Detecting implicit expressions of sentiment in text based on commonsense knowledge. In *Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA 2.011)*, pages 53–60, Portland, Oregon, June. Association for Computational Linguistics.
- Farah Benamara, Sabatier Irit, Carmine Cesarano, Napoli Federico, and Diego Reforgiato. 2007. Sentiment analysis : Adjectives and adverbs are better than adjectives alone. In *Proc of Int Conf on Weblogs and Social Media*, pages 1–4.
- Brett Drury, Gaël Dias, and Luís Torgo. 2011. A contextual classification strategy for polarity analysis of direct quotations from financial news. In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 434–440, Hissar, Bulgaria, September. RANLP 2011 Organising Committee.
- Charles J. Fillmore, Russell Lee-Goldman, and Russell Rhodes, to appear. *Sign-based Construction Grammar*, chapter The FrameNet Constructicon. CSLI, Stanford, CA.
- Charles J. Fillmore. 1989. Grammatical construction theory and the familiar dichotomies. In R. Dietrich and C.F. Graumann, editors, *Language processing in social context*, pages 17–38. North-Holland/Elsevier, Amsterdam.
- S.M. Kim and E. Hovy. 2006. Extracting opinions, opinion holders, and topics expressed in online news media text. In *Proceedings of the Workshop on Sentiment and Subjectivity in Text*, pages 1–8. Association for Computational Linguistics.
- Isa Maks and Piek Vossen. 2011. A verb lexicon model for deep sentiment analysis and opinion mining applications. In *Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA 2.011)*, pages 10–18, Portland, Oregon, June. Association for Computational Linguistics.
- Karo Moilanen and Stephen Pulman. 2007. Sentiment composition. In *Proceedings of RANLP 2007*, Borovets, Bulgaria.
- A. Neviarouskaya, H. Prendinger, and M. Ishizuka. 2009. Sentiful: Generating a reliable lexicon for sentiment analysis. In *Affective Computing and Intelligent Interaction and Workshops, 2009. AII 2009. 3rd International Conference on*, pages 1–6. Ieee.
- J. Ruppenhofer, S. Somasundaran, and J. Wiebe. 2008. Finding the sources and targets of subjective expressions. In *LREC*, Marrakech, Morocco.
- Roser Saurí. 2008. *A Factuality Profiler for Eventualities in Text*. Ph.d., Brandeis University.
- Yohei Seki. 2007. Crosslingual opinion extraction from author and authority viewpoints at ntcir-6. In *Proceedings of NTCIR-6 Workshop Meeting*, Tokyo, Japan.
- Mostafa Shaikh, Helmut Prendinger, and Ishizuka Mitsuru. 2007. Assessing sentiment of text by semantic dependency and contextual valence analysis. *Affective Computing and Intelligent Interaction*, pages 191–202.
- Veselin Stoyanov and Claire Cardie. 2008. Topic identification for fine-grained opinion analysis. In *Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1, COLING '08*, pages 817–824, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Veselin Stoyanov and Claire Cardie. 2011. Automatically creating general-purpose opinion summaries from text. In *Proceedings of RANLP 2011*, pages 202–209, Hissar, Bulgaria, September.
- Cigdem Toprak, Niklas Jakob, and Iryna Gurevych. 2010. Sentence and expression level annotation of opinions in user-generated discourse. In *Proceedings of ACL-10, the 48th Annual Meeting of the Association for Computational Linguistics*, Portland. Association for Computational Linguistics.
- Janyce Wiebe and Rada Mihalcea. 2006. Word sense and subjectivity. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics, ACL-44*, pages 1065–1072, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2/3):164–210.
- Janyce Wiebe. 2000. Learning subjective adjectives from corpora. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*, pages 735–740, Austin, Texas.

Michael Wiegand. 2010. *Hybrid approaches to sentiment analysis*. Ph.D. thesis, Saarland University, Saarbrücken.

Theresa Ann Wilson. 2008. *Fine-grained Subjectivity and Sentiment Analysis: Recognizing the Intensity, Polarity, and Attitudes of Private States*. Ph.D. thesis, University of Pittsburgh.