How can NLP Tasks Mutually Benefit Sentiment Analysis?
A Holistic Approach to Sentiment Analysis

Lingjia Deng
Intelligent Systems Program

University of Pittsburgh
lid29@pitt.edu

Janyce Wiebe
Intelligent Systems Program
Department of Computer Science
University of Pittsburgh
wiebe@cs.pitt.edu

Abstract

Existing opinion analysis techniques rely on the clues within the sentence that focus on the sentiment analysis task itself. However, the sentiment analysis task is not isolated from other NLP tasks (co-reference resolution, entity linking, etc) but they can benefit each other. In this paper, we define dependencies between sentiment analysis and other tasks, and express the dependencies in first order logic rules regardless of the representations of different tasks. The conceptual framework proposed in this paper using such dependency rules as constraints aims at exploiting information outside the sentence and outside the document to improve sentiment analysis. Further, the framework allows exception to the rules.

1 Introduction

Opinions are ubiquitous in language. Existing opinion analysis techniques rely on the clues in the sentence that focus on the sentiment analysis task itself. Consider, for example,

(Ex1) Oh no, the voters defeated the bill.

The sentiment lexicons are used to recognize *Oh no* as a negative opinion, the semantic role labeling features are used to recognize the target is the *defeating* event (Yang and Cardie, 2013), and the implicatures are used to recognize the writer is positive toward the bill since the writer is negative toward the defeating event which harms the bill (Deng et al., 2014; Deng and Wiebe, 2015). These work mainly

rely on the clues that directly indicate opinions (e.g., recognizing *On no* as a negative opinion), or indicate components of opinions (e.g., recognizng the target being *defeating*), or indicate other opinions based on the information within the sentence (e.g., recognizing a positive opinion toward *the bill*). They do not exploit the vast amount of knowledge outside the sentence, which are outputs from many NLP tasks. But the task of sentiment analysis may benefit from those tasks. Consider (Ex2), for example,

(Ex2) President Obama proposed the healthcare reform. I support him.

we recognize in the second the sentence that the writer (*I*) is positive toward *him*. Further, we recognize the writer is positive toward *President Obama* since by co-reference resolution we know that *him* refers to *President Obama*.

Meanwhile, other NLP tasks may benefit from sentiment analysis. Consider, for example,

(Ex3) The allies successfully defeated Nazi. They are really brave.

The sentiment analysis system may infer that the writer is positive toward the allies and negative toward Nazi. Based on this information, we can infer that the word *they* in the second sentence refers to the allies instead of Nazi. Thus the sentiment analysis outputs help the co-reference resolution task.

The relation of sentiment analysis and other NLP tasks cannot be easily modelled as a pipeline. For example, in (Ex2) a co-reference resolution needs to be run first to infer the writer is positive toward Obama, while in (Ex3) the positive sentiments needs

53

to be recognized first to infer the word *they* refer to the allies. Previous work (Deng et al., 2014; Deng and Wiebe, 2015) develop joint models to infer sentiments based on the implicature rules (e.g, (Ex1)). They first develop independent systems to recognize sentiments and components of sentiments. Then joint approaches are used to take the outputs from independent systems as input and globally infer sentiments based on all the input information. The implicature rules are used as constraints in the joint approaches. Similar to their method, we can use rules introduced in this paper as constraints in the joint models, and jointly resolve sentiment analysis and other NLP tasks. Furthermore, though the representations of knowledge that different tasks generate are various, the dependencies in this paper are expressed in a unified way: first order logic rules.

In summary, this paper presents a conceptual framework using the newly defined dependency rules as constraints of joint models to exploit various kinds of knowledge to make progress toward a deeper interpretation of subjective language. The background of joint models is given in Section 2. The dependency rules and corresponding NLP tasks are given in Section 3. Furthermore, the framework allows exceptions to the rules, which will be discussed in Section 4. Finally we give the conclusion.

2 Background

The ultimate goal of this paper is to improve sentiment analysis by exploiting various knowledge, each of which corresponds to an NLP task. We define atoms corresponding to the tasks in Table 1.

The primary task of sentiment analysis is to assign scores to the atoms *pos(X,Y)* and *neg(X,Y)* (i.e., assigning true or false, or numeric scores to the atoms). Most of previous work directly assign scores to *pos(X,Y)* and *neg(X,Y)* without any dependency rule. Some recent work (Deng et al., 2014; Deng and Wiebe, 2015) take the scores as local scores and maximize the sum of scores of all the Primary Task and Implicature Knowledge atoms in Table 1 w.r.t. the constraints defined by a subset of implicature rules in (Wiebe and Deng, 2014a). Their experiments have shown that the joint models are able to choose a better assignment of the scores to all the atoms globally rather than make individual decisions

according to local scores only.

However, the previous conceptual framework (Wiebe and Deng, 2014b) only defines rules over the Implicature Knowledge atoms to only consider the information within the sentence. And they are limited to a particular type of event: +/-effect event. Instead, this paper introduces rules defined over the External Knowledge atoms to exploit knowledge outside the sentence and outside the document. Further, the atoms defined in this paper are general so that people can use these atoms to design more rules.

3 Dependency Rules and NLP Tasks

The dependency rules are expressed as first order logic rules. As a start, we represent one of the rules from (Wiebe and Deng, 2014a) in first order logic applied to (Ex1) in Section 1.

In (Ex1), we infer from the negative sentiment toward the defeated event that the writer is positive toward the bill. The defeated event is defined as a -effect event since it has negative effect on the theme, the bill (Deng et al., 2013). The instantiated rule is:

(R1) neg(writer, defeat) $\wedge$ -effect(defeat)
$\quad$ $\wedge$ theme(defeat, bill) $\Rightarrow$ pos(writer, bill)

Different from the rules defined in (Wiebe and Deng, 2014a), we define new rules depicting the dependencies between sentiments (e.g., *pos(X,Y)*) and external knowledge outside the sentence (e.g., *posExternal(X,Y)*). We focus on two types of knowledge. The first involves knowledge from elsewhere within the same document, and the second involves document-external knowledge such as that stored in a knowledge base (e.g., Freebase).

For ease of understanding, a rule is presented as an instantiated rule applied to an example (as (R1) above). There are variations of the rules listed in this paper according to different context.[1]

3.1 Rules of Intra-Document Knowledge

Co-reference Resolution. Recall (Ex2) in Section 1. The writer is positive toward Obama because the

---

[1] For example, a variation of (R1) is: pos(writer, defeat) $\wedge$ -effect(defeat) $\wedge$ theme(defeat, bill) $\Rightarrow$ neg(writer, bill)

Primary Task		
pos(X,Y)	X has positive sentiment toward Y, evoked in the current sentence	
neg(X,Y)	X has negative sentiment toward Y, evoked in the current sentence	
Implicature Knowledge		
+sentiment(S)	S is a positive sentiment	+effect(T)	T is a +effect event
-sentiment(S)	S is a negative sentiment	-effect(T)	T is a -effect event
source(S,X)	the source of S is X	agent(T,X)	the agent of T is X
target(S,Y)	the target of S is Y	theme(T,Y)	the theme of T is Y
External Knowledge		
posExternal(X,Y)	X has positive sentiment toward Y, external to the sentence	
negExternal(X,Y)	X has negative sentiment toward Y, external to the sentence	
sameEntity(X,Y)	X and Y refer to same entity	
altEntity(X,Y)	X and Y represent alternative entities	
agree(X,Y)	X and Y agree with each other	
reinforcing(X,Y)	X and Y are reinforcing sentiments	
non-reinforcing(X,Y)	X and Y are non-reinforcing sentiments	
ideology(X,I)	X holds ideology I	
aspect(X,Y)	an aspect (feature) of X is Y	

Table 1: Atoms in the Rules.

word *him* refers to Obama. The instantiated rule is:

(R2) posExternal(writer,him) ∧
    sameEntity(him,Obama)
    ⇒ pos(writer,Obama)

Agree. We may also infer that the writer has the same sentiments as sources with whom he or she agrees. While much previous work detects agreement at the turn level in conversation (Michel Galley, 2004; Wang et al., 2011), or identifies participants who agree with one another (Hassan et al., 2012; Abu-Jbara et al., 2012; Park et al., 2011), there is recent work on detecting agreement within documents (Wang and Cardie, 2014; Abbott et al., 2011; Misra and Walker, 2013). Consider, *I agree with Paul. ... The plan is a brilliant idea.* The writer (I) agree with Paul, and the writer is positive toward the plan. Then we infer that probably Paul is positive toward the plan.

(R3) agree(writer,Paul) ∧ posExternal(writer,plan)
    ⇒ pos(Paul,plan)

Opinion-oriented Discourse Models. Furthermore, previous work have developed opinion-oriented discourse models (*OODM*s) (Somasundaran, 2010). The OODM models recognize toward

which entities the writer's sentiments are the same (*sameEntity*), and toward which entities the writer's sentiments are opposite (*altEntity*). The discourse *sameEntity* relation covers not only identity, but also part-whole, synonymy, generalization, specialization, entity-attribute/aspect, instantiation, cause-effect, and implicit background topic, i.e., relations that have been studied by many researchers in the context of anaphora and co-reference (e.g. (Clark, 1975; Vieira and Poesio, 2000; Mueller and Strube, 2001)). Two entities are in an *altEntity* relation if they are mutually exclusive options in the context of the discourse. For example, in a debate about mobile phones, the iPhone and iOS are considered as *sameEntity*, while the Android and iPhone are considered as *altEntity*. In OODM models, same sentiments toward same entities express the same stance, and opposite sentiments toward alternative targets express the same overall stance (Somasundaran, 2010).

(R4) posExternal(writer,iOS) ∧
    sameEntity(iOS,iPhone) ⇒ pos(writer,iPhone)
(R5) posExternal(writer,iOS) ∧
    altEntity(iOS,Android) ⇒ neg(writer,Android)

However, the opinions throughout the documents

may not always be consistent. In the same document, a source may be both positive and negative toward a target. In this paper, we define **rules to explain conflicting opinions** in the document.

Aspect-Based Sentiment Analysis. In one case, the source has different opinions about different aspects of the same target. Consider *The iPhone display is beautiful. But it is too expensive.* The writer is positive toward the display while negative toward the price. Such case can be modelled via the rule:

(R6) posExternal(writer,iPhone) ∧
    sameEntity(iPhone,it) ∧ neg(writer,it) ⇔
    aspect(iPhone,display) ∧ aspect(it,price) ∧
    posExternal(writer,display) ∧ neg(writer,price)

Several researchers have focused on the task of mining data to discover aspects of products and sentiments toward different aspects (Liu, 2012).

(Non-)Reinforcing Sentiment Analysis. In the other case, people may be ambivalent, or change their minds in the course of a document. Two sentiments may be in *reinforcing* or *non-reinforcing* discourse scenarios. Reinforcing relations exist between opinions when they contribute to the same overall stance. Non-reinforcing relations exist between opinions that show ambivalence, which represents a discourse scenario in which inconsistent sentiments are expressed with respect to a stance (Somasundaran, 2010; Trivedi and Eisenstein, 2013; Bhatia et al., 2015). Consider, *It is expensive. ... However, I think it is worth a try if I loan to buy the phone.* Previous work (Somasundaran, 2010) may recognize that two non-reinforcing sentiments occur (indicated by the word *However*). *S1* represents the negative opinion in the first sentence expressed toward it, and *S2* represents the positive opinion in the second sentence expressed toward the phone.

(R7) non-reinforcing(S1,S2) ∧
    source(S1,writer) ∧ source(S2,writer) ∧
    target(S1,It) ∧ target(S2,the phone) ∧
    sameEntity(It, the phone)
    ∧ negExternal(writer,It) ⇒ pos(writer,the phone)

Two non-reinforcing opinions can also be expressed toward alternative entities. Consider, *The iPhone is too expensive. ... But the price of Android cannot guarantee a satisfactorily smooth operating system.* *S1* represents the negative opinion in the first sentence expressed toward iPhone, and *S2* represents the negative opinion in the second sentence expressed toward Android.

(R8) non-reinforcing(S1,S2) ∧
    source(S1,writer) ∧ source(S2,writer) ∧
    target(S1,iPhone) ∧ target(S2,Android) ∧
    altEntity(iPhone, Android)
    ∧ negExternal(writer,iPhone)
    ⇒ neg(writer,Android)

3.2 Rules of Extra-Document Knowledge

Entity Linking. Knowledge from outside the document is also important. For example, the work in entity linking maps entity mentions (e.g.,Obama, US President) in the text to entries in the knowledge base (e.g., BARACK OBAMA) (Ji and Grishman, 2011; Rao et al., 2013). Such information can be exploited to recognize *sameEntity*, as shown below.

(R9) sameEntity(Obama, BARACK OBAMA)
    ∧ sameEntity(US President, BARACK OBAMA)
    ⇒ sameEntity(Obama, US President)

Thus, we can use the knowledge base to enrich the recognition of *sameEntity* and help recognize more sentiments.

Ideology. Groups of people sharing the same ideology tend to have the same opinions about certain things. Suppose we have known that Donald Trump is conservative, and a conservative ideology is against the concept of gun control, then we probably infer that he is opposed to gun control in the context.

(R10) ideology(Donald Trump,CONSERVATIVE)
    ∧ negExternal(CONSERVATIVE,GUN CONTROL)
    ∧ sameEntity(GUN CONTROL, gun control)
    ⇒ neg(Donald Trump,gun control)

Rather than attempt to computationally define a general notion of ideology, people in NLP tend to use data for which specific ideologies have been defined. Previous work have studied recognizing ide-

ologies including political party affiliation (Iyyer et al., 2014), or labels such as *left*, *right*, and *center* (Sim et al., 2013), or use a proxy for ideology such as voting record (Gerrish and Blei, 2011).

4 Integrating Evidence Against Rules

The framework allows exceptions to the rules. The joint models implemented in the previous work (Deng et al., 2014; Deng and Wiebe, 2015) use implicature rules as soft constraints. However, previous work didn't investigate when the rules are blocked. In this section we introduce two types of evidences against the rules. [2] The first case is when the event is involuntarily conducted. Consider Ex(4A) below.

> (Ex4A) The insurance companies will increase their spending on health care improvement.
> (Ex4B) The insurance companies will **be required to** increase their spending on health care improvement.

Assuming that there is a positive sentiment toward health care improvement in (Ex4A), the implicature rules infer a positive sentiment toward the insurance companies since the companies are increasing the improvement. However, consider the variation (Ex4B). The implicature here is less strong and perhaps defeated. The reason is that the companies will be *forced* to increase their spending.

Another case is when an event is *accidental*. For example in (Ex5A),

> Ex(5A) John deleted the file I need.
> Ex(5B) John **accidentally** deleted the file I need.

the rules imply a negative sentiment toward John. However, this inference is weakened in the variation (Ex5B). To recognize these cases, lexical clues are important, such as unintentionally, involuntary. Given a list of seed words, resources such as Word-Net, word embeddings (Mikolov et al., 2013) and paraphrase databases (e.g., PPDB (Ganitkevitch et al., 2013)) can be utilized to find semantically similar words and phrases.

Further, we can integrate the evidences against the rules into the rules themselves. For example, (R11) is the rule applied to (Ex4A). If we want to integrate evidences against rules to model cases such as (Ex4B), we can revised rule by incorporating the atom *initiative* representing whether an event is conducted initiatively, as (R11*) shows. If initiative(increase) is false, this inference is blocked.

(R11) pos(writer, increase) $\wedge$
agent(company, increase)
$\Rightarrow$ pos(writer, company)
(R11*) pos(writer, increase) $\wedge$
agent(company, increase) $\wedge$ **initiative(increase)**
$\Rightarrow$ pos(writer, company)

This shows that the rule is flexible to add or remove an atom. Also, the framework is flexible to block a rule in the context. Such flexibility allows the framework to model various context and adapt to different genres.

5 Conclusion

Sentiment analysis is not isolated from other NLP tasks. The conceptual framework in this paper aims at improving sentiment analysis by introducing dependency rules between sentiments and various knowledge provided from various NLP tasks including co-reference resolution, opinion discourse analysis, entity linking and ideology, etc. The framework uses dependency rules as constraints in the joint models. Further, the framework can block a rule in context by recognizing evidences against the instantiated rule. Though it is a conceptual framework, it bridges different tasks of sentiment analysis and various tasks in NLP together to provide a holistic approach to sentiment analysis and the other tasks as well.

References

Rob Abbott, Marilyn Walker, Pranav Anand, Jean E. Fox Tree, Robeson Bowmani, and Joseph King. 2011.

---

[2] We realized these cases from the study of implicit sentiment in Greene and Resnik (Greene and Resnik, 2009).

How can you say such things?!?: Recognizing disagreement in informal political argument. In *Proceedings of the Workshop on Language in Social Media (LSM 2011)*, pages 2–11, Portland, Oregon, June. Association for Computational Linguistics.

Amjad Abu-Jbara, Pradeep Dasigi, Mona Diab, and Dragomir Radev. 2012. Subgroup detection in ideological discussions. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 399–409, Jeju Island, Korea, July. Association for Computational Linguistics.

Parminder Bhatia, Yangfeng Ji, and Jacob Eisenstein. 2015. Better document-level sentiment analysis from rst discourse parsing. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2212–2218, Lisbon, Portugal, September. Association for Computational Linguistics.

H. H. Clark. 1975. Bridging. *Theoretical issues in natural language processing . New York: Association for Computing Machinery*, page 6.

Lingjia Deng and Janyce Wiebe. 2015. 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*, pages 179–189, Lisbon, Portugal, September. Association for Computational Linguistics.

Lingjia Deng, Yoonjung Choi, and Janyce Wiebe. 2013. Benefactive/malefactive event and writer attitude annotation. In *ACL 2013 (short paper)*. Association for Computational Linguistics.

Lingjia Deng, Janyce Wiebe, and Yoonjung Choi. 2014. Joint inference and disambiguation of implicit sentiments via implicature constraints. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 79–88, Dublin, Ireland, August. Dublin City University and Association for Computational Linguistics.

Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. Ppdb: The paraphrase database. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 758–764, Atlanta, Georgia, June. Association for Computational Linguistics.

Sean Gerrish and David M. Blei. 2011. Predicting legislative roll calls from text. In *ICML*.

Stephan Greene and Philip Resnik. 2009. More than words: Syntactic packaging and implicit sentiment. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*,

pages 503–511, Boulder, Colorado, June. Association for Computational Linguistics.

Ahmed Hassan, Amjad Abu-Jbara, and Dragomir Radev. 2012. Detecting subgroups in online discussions by modeling positive and negative relations among participants. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 59–70, Jeju Island, Korea, July. Association for Computational Linguistics.

Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014. Political ideology detection using recursive neural networks. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1113–1122, Baltimore, Maryland, June. Association for Computational Linguistics.

Heng Ji and Ralph Grishman. 2011. Knowledge base population: Successful approaches and challenges. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 1148–1158. Association for Computational Linguistics.

Bing Liu. 2012. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.

Julia Hirschberg Elizabeth Shriberg Michel Galley, Kathleen McKeown. 2004. Identifying agreement and disagreement in conversational speech: Use of bayesian networks to model pragmatic dependencies. In *Proceedings of the 42th Annual Meeting of the Association for Computational Linguistics (ACL-2004)*.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.

Amita Misra and Marilyn Walker. 2013. Topic independent identification of agreement and disagreement in social media dialogue. In *Proceedings of the SIGDIAL 2013 Conference*, pages 41–50, Metz, France, August. Association for Computational Linguistics.

Christoph Mueller and Michael Strube. 2001. Annotating anaphoric and bridging relations with mmax. In *2nd SIGdial Workshop on Discourse and Dialogue*.

Souneil Park, Kyung Soon Lee, and Junehwa Song. 2011. Contrasting opposing views of news articles on contentious issues. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 340–349, Portland, Oregon, USA, June. Association for Computational Linguistics.

Delip Rao, Paul McNamee, and Mark Dredze. 2013. Entity linking: Finding extracted entities in a knowledge

base. In *Multi-source, multilingual information extraction and summarization*, pages 93–115. Springer.

Yanchuan Sim, Brice D. L. Acree, Justin H. Gross, and Noah A. Smith. 2013. Measuring ideological proportions in political speeches. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 91–101, Seattle, Washington, USA, October. Association for Computational Linguistics.

Swapna Somasundaran. 2010. *Discourse-Level Relations for Opinion Analysis*. Ph.D. thesis, Department of Computer Science, University of Pittsburgh.

Rakshit Trivedi and Jacob Eisenstein. 2013. Discourse connectors for latent subjectivity in sentiment analysis. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 808–813, Atlanta, Georgia, June. Association for Computational Linguistics.

Renata Vieira and Massimo Poesio. 2000. An empirically based system for processing definite descriptions. *Computational Linguistics*, 26(4):539–593.

Lu Wang and Claire Cardie. 2014. Improving agreement and disagreement identification in online discussions with a socially-tuned sentiment lexicon. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 97–106, Baltimore, Maryland, June. Association for Computational Linguistics.

Wen Wang, Sibel Yaman, Kristin Precoda, Colleen Richey, and Geoffrey Raymond. 2011. Detection of agreement and disagreement in broadcast conversations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 374–378, Portland, Oregon, USA, June. Association for Computational Linguistics.

Janyce Wiebe and Lingjia Deng. 2014a. An account of opinion implicatures. *arXiv*, 1404.6491[cs.CL].

Janyce Wiebe and Lingjia Deng. 2014b. A conceptual framework for inferring implicatures. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 154–159, Baltimore, Maryland, June. Association for Computational Linguistics.

Bishan Yang and Claire Cardie. 2013. Joint inference for fine-grained opinion extraction. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1640–1649, Sofia, Bulgaria, August. Association for Computational Linguistics.