

Chapter 4

Challenges in Sentiment Analysis

Saif M. Mohammad

Abstract A vast majority of the work in Sentiment Analysis has been on developing more accurate sentiment classifiers, usually involving supervised machine learning algorithms and a battery of features. Surveys by Pang and Lee (Found Trends Inf Retr 2(1–2):1–135, 2008), Liu and Zhang (A survey of opinion mining and sentiment analysis. In: Aggarwal CC, Zhai C (eds) In: Mining text data. Springer, New York, pp 415–463, 2012), and Mohammad (Mohammad Sentiment analysis: detecting valence, emotions, and other effectual states from text. In: Meiselman H (ed) Emotion measurement. Elsevier, Amsterdam, 2016b) give summaries of the many automatic classifiers, features, and datasets used to detect sentiment. In this chapter, we flesh out some of the challenges that still remain, questions that have not been explored sufficiently, and new issues emerging from taking on new sentiment analysis problems. We also discuss proposals to deal with these challenges. The goal of this chapter is to equip researchers and practitioners with pointers to the latest developments in sentiment analysis and encourage more work in the diverse landscape of problems, especially those areas that are relatively less explored.

Keywords Sentiment analysis tasks • Sentiment of the writer, reader, and other entities • Sentiment towards aspects of an entity • Stance detection • Sentiment lexicons • Sentiment annotation • Multilingual sentiment analysis

4.1 Introduction

There has been a large volume of work in sentiment analysis over the past decade and it continues to rapidly develop in new directions. However, much of it is on developing more accurate sentiment classifiers. In this chapter, we flesh out some of the challenges that still remain. We start by discussing different sentiment analysis

S.M. Mohammad (✉)
National Research Council Canada, 1200 Montreal Rd., Ottawa, ON, Canada
e-mail: Saif.Mohammad@nrc-cnrc.gc.ca

problems and how one of the challenges is to explore new sentiment analysis problems that go beyond simply determining whether a piece of text is positive, negative, or neutral (Sect. 4.2). Some of the more ambitious problems that need more work include detecting sentiment at various levels of text granularities (terms, sentences, paragraphs, etc); detecting sentiment of the reader or sentiment of entities mentioned in the text; detecting sentiment towards aspects of products; detecting stance towards pre-specified targets that may not be explicitly mentioned in the text and that may not be the targets of opinion in the text; and detecting semantic roles of sentiment. Since many sentiment analysis systems rely on sentiment lexicons, we discuss capabilities and limitations of existing manually and automatically created sentiment lexicons in Sect. 4.3. In Sect. 4.4, we discuss the difficult problem of sentiment composition—how to predict the sentiment of a combination of terms. More specifically, we discuss the determination of sentiment of phrases (that may include negators, degree adverbs, and intensifiers) and sentiment of sentences and tweets. In Sect. 4.5, we discuss challenges in annotation of data for sentiment. We provide categories of sentences that are particularly challenging for sentiment annotation. Section 4.6 presents challenges in multilingual sentiment analysis. This is followed by a discussion on the challenges of applying sentiment analysis to downstream applications, and finally, some concluding remarks (Sect. 4.7).

4.2 The Array of Sentiment Analysis Tasks

Sentiment analysis is a generic name for a large number of opinion and affect related tasks, each of which present their own unique challenges. The sub-sections below provide an overview.

4.2.1 *Sentiment at Different Text Granularities*

Sentiment can be determined at various levels: from sentiment associations of words and phrases; to sentiment of sentences, SMS messages, chat messages, and tweets; to sentiment in product reviews, blog posts, and whole documents. A word-sentiment (or valence) association lexicon may have entries such as:

delighted – positive
killed – negative
shout – negative
desk – neutral

These lexicons can be created either by manual annotation or through automatic means. Manually created lexicons tend to be in the order of a few thousand entries, but automatically generated lexicons can capture sentiment associations for hundreds of thousands unigrams (single word strings) and even for larger expres-

sions such as bigrams (two-word sequences) and trigrams (three-word sequences). Entries in an automatically generated lexicon often also include a real-valued score indicating the strength of association between the word and the valence category. These numbers are prior estimates of the sentiment of terms in an average usage of the term. While sentiment lexicons are often useful in sentence-level sentiment analysis,¹ the same terms may convey different sentiments in different contexts. The SemEval 2013 and 2014 *Sentiment Analysis in Twitter* shared tasks had a separate sub-task aimed at identifying sentiment of terms in context. Automatic systems have largely performed well in this task, obtaining F-scores close to 0.9. We discuss manually and automatically created sentiment lexicons in more detail in Sect. 4.3.

Sentence-level valence classification systems assign labels such as positive, negative, or neutral to whole sentences. It should be noted that the valence of a sentence is not simply the sum of the polarities of its constituent words. Automatic systems learn a model from labeled training data (instances that are already marked as positive, negative, or neutral) using a large number of features such as word and character ngrams, valence association lexicons, negation lists, word clusters, and even embeddings-based features. In recent years, there have been a number of shared task competitions on valence classification such as the 2013, 2014, and 2015 SemEval shared tasks titled *Sentiment Analysis in Twitter*, the 2014 and 2015 SemEval shared tasks on *Aspect Based Sentiment Analysis*, the 2015 SemEval shared task *Sentiment Analysis of Figurative Language in Twitter*, and the 2015 Kaggle competition *Sentiment Analysis on Movie Reviews*.² The NRC-Canada system (Mohammad et al. 2013a; Kiritchenko et al. 2014b), a supervised machine learning system, came first in the 2013 and 2014 competitions. Other sentiment analysis systems developed specifically for tweets include those by Pak and Paroubek (2010), Agarwal et al. (2011), Thelwall et al. (2011), Brody and Diakopoulos (2011), Aisopos et al. (2012), and Bakliwal et al. (2012). However, even the best systems currently obtain an F-score of only about 0.7.

Sentiment analysis involving many sentences is often broken down into the sentiment analysis of the component sentences. However, there is interesting work in sentiment analysis of documents to generate text summaries (Ku et al. 2006; Liu et al. 2007; Somprasertsri and Lalitrojwong 2010; Stoyanov and Cardie 2006; Lloret et al. 2009), as well as detecting the patterns of sentiment and detecting sentiment networks in novels and fairy tales (Nalisnick and Baird 2013a,b; Mohammad and Yang 2011).

¹The top systems in the SemEval-2013 and 2014 *Sentiment Analysis in Twitter* tasks used large sentiment lexicons (Wilson et al. 2013; Rosenthal et al. 2014a).

²<http://alt.qcri.org/semEval2015/task10/>

<http://alt.qcri.org/semEval2015/task12/>

<http://alt.qcri.org/semEval2015/task11/>

<http://www.kaggle.com/c/sentiment-analysis-on-movie-reviews>

4.2.2 *Detecting Sentiment of the Writer, Reader, and Other Entities*

On the surface, sentiment may seem unambiguous, but looking closer, it is easy to see how sentiment can be associated with any of the following: 1. the speaker or writer, 2. the listener or reader, or 3. one or more entities mentioned in the utterance. A large majority of research in sentiment analysis has focused on detecting the sentiment of the speaker, and this is often done by analyzing only the utterance. However, there are several instances where it is unclear whether the sentiment in the utterance is the same as the sentiment of the speaker. For example, consider:

James: *The pop star suffered a fatal overdose of heroine.*

The sentence describes a negative event (death of a person), but it is unclear whether to conclude that James (the speaker) is personally saddened by the event. It is possible that James is a news reader and merely communicating information about the event. Developers of sentiment systems have to decide before hand whether they wish to assign a negative or neutral sentiment to the speaker in such cases. More generally, they have to decide whether the speaker's sentiment will be chosen to be neutral in absence of clear signifiers of the speaker's own sentiment, or whether the speaker's sentiment will be chosen to be the same as the sentiment of events and topics mentioned in the utterance.

On the other hand, people can react differently to the same utterance, for example, people on opposite sides of a debate or rival sports fans. Thus modeling listener sentiment requires modeling listener profiles. This is an area of research not explored much by the community. Similarly, there is no work on modeling sentiment of entities mentioned in the text, for example, given:

Drew: *Jackson could not stop talking about the new Game of Thrones episode.*

It will be useful to develop automatic systems that can deduce that Jackson (not Drew) liked the new episode of *Game of Thrones* (a TV show).

4.2.3 *Sentiment Towards Aspects of an Entity*

A review of a product or service can express sentiment towards various aspects. For example, a restaurant review can speak positively about the service, but express a negative attitude towards the food. There is now a growing amount of work in detecting aspects of products in text and also in determining sentiment towards these aspects. In 2014, a shared task was organized for detecting aspect sentiment in restaurant and laptop reviews (Pontiki et al. 2014a). The best performing systems had a strong sentence-level sentiment analysis system to which they added localization features so that more weight was given to sentiment features close to the mention of the aspect. This task was repeated in 2015. It will be useful to develop

aspect-based sentiment systems for other domains such as blogs and news articles as well. (See proceeding of SemEval-2014 and 2015 for details about participating aspect sentiment systems.)

4.2.4 *Stance Detection*

Stance detection is the task of automatically determining from text whether the author of the text is in favor of, against, or neutral towards a proposition or target. For example, given the following target and text pair:

Target of interest: *women have the right to abortion*
Text: *A foetus has rights too!*

Humans can deduce from the text that the speaker is against the proposition. However, this is a challenging task for computers. To successfully detect stance, automatic systems often have to identify relevant bits of information that may not be present in the focus text. The systems also have to first identify the target of opinion in the text and then determine its implication on the target of interest. Note that the target of opinion need not be the same as the target of interest. For example, that if one is actively supporting foetus rights (target of opinion), then he or she is likely against the right to abortion (target of interest). Automatic systems can obtain such information from large amounts of domain text.

Automatically detecting stance has widespread applications in information retrieval, text summarization, and textual entailment. In fact, one can argue that stance detection can bring complementary information to sentiment analysis, because we often care about the authors evaluative outlook towards *specific targets* and propositions rather than simply about whether the speaker was angry or happy.

Mohammad et al. (2016b) created the first dataset of tweets labeled for both stance and sentiment. More than 4000 tweets are annotated for whether one can deduce favorable or unfavorable stance towards one of five targets ‘Atheism’, ‘Climate Change is a Real Concern’, ‘Feminist Movement’, ‘Hillary Clinton’, and ‘Legalization of Abortion’. Each of these tweets is also annotated for whether the target of opinion expressed in the tweet is the same as the given target of interest. Finally, each tweet is annotated for whether it conveys positive, negative, or neutral sentiment. Partitions of this stance-annotated data were used as training and test sets in the SemEval-2016 shared task competition, Task #6: Detecting Stance from Tweets Mohammad et al. (2016a). Participants were provided with 2,914 training instances labeled for stance for the five targets. The test data included 1,249 instances. All of the stance data is made freely available through the shared task website. The task received submissions from 19 teams. The best performing system obtained an overall average F-score of 67.8 in a three-way classification task: favour, against, or neither. They employed two recurrent neural network (RNN) classifiers: the first was trained to predict task-relevant hashtags on a large unlabeled Twitter corpus. This network was used to initialize a second RNN classifier, which

was trained with the provided training data (Zarrella and Marsh 2016). Mohammad et al. (2016b) developed a SVM system that only uses features drawn from word and character ngrams and word embeddings to obtain an even better F-score of 70.3 on the shared task’s test set. Yet, performance of systems is substantially lower on tweets where the target of opinion is an entity other than the target of interest.

Most of the earlier work focused on two-sided debates, for example on congressional debates (Thomas et al. 2006) or debates in online forums (Somasundaran and Wiebe 2009; Murakami and Raymond 2010; Anand et al. 2011; Walker et al. 2012; Hasan and Ng 2013; Sridhar et al. 2014). New research in domains such as social media texts, and approaches that combine traditional sentiment analysis with relation extraction can make a significant impact in improving the state-of-the-art in automatic stance detection.

4.2.5 Detecting Semantic Roles of Feeling

Past work in sentiment analysis has focused extensively on detecting polarity, and to a smaller extent on detecting the target of the sentiment (the stimulus) (Popescu and Etzioni 2005; Su et al. 2006; Xu et al. 2013; Qadir 2009; Zhang et al. 2010; Zhang and Liu 2011; Kessler and Nicolov 2009). However, there exist other aspects relevant to sentiment. Tables 4.1 and 4.2 show FrameNet (Baker et al. 1998) frames for ‘feelings’ and ‘emotions’, respectively. Observe that in addition to Evaluation, State, and Stimulus, several other roles such as Reason, Degree, Topic, and Circumstance are also of significance and beneficial to down-stream applications such as information retrieval, summarization, and textual entailment. Detecting these various roles is essentially a semantic role-labeling problem (Gildea and Jurafsky 2002; Màrquez et al. 2008; Palmer et al. 2010), and it is possible that they can be modeled jointly to improve detection accuracy. Li and Xu (2014) proposed a rule-based system to extract the event that was the cause of an emotional Weibo (Chinese microblogging service) message. Mohammad et al. (2015a) created a corpus of tweets from the run up to the 2012 US presidential elections, with annotations for sentiment, emotion, stimulus, and experiencer. The data also includes annotations for whether the tweet is sarcastic,

Table 4.1 The FrameNet frame for feeling

Role	Description
Core	
Emotion	The feeling that the experiencer experiences
State	The state the experiencer is in
Evaluation	A negative or positive assessment of the experiencer regarding his/her state
Experiencer	One who experiences the emotion and is in the state
Non-Core	
Explanation	The thing that leads to the experiencer feeling the emotion or state

Table 4.2 The FrameNet frame for emotions

Role	Description
Core	
Experiencer	The person that experiences or feels the emotion
State	The abstract noun that describes the experience
Stimulus	The person or event that evokes the emotional response in the experiencer.
Topic	The general area in which the emotion occurs
Non-Core	
Circumstances	The condition in which stimulus evokes response
Degree	The extent to which the experiencer's emotion deviates from the norm for the emotion
Empathy_target	The Empathy_target is the individual or individuals with which the experiencer identifies emotionally
Manner	Any description of the way in which the experiencer experiences the stimulus which is not covered by more specific frame elements
Reason	The explanation for why the stimulus evokes a certain emotional response

ironic, or hyperbolic. Diman Ghazi and Szpakowicz (2015) compiled FrameNet sentences that were tagged with the stimulus of certain emotions.

4.2.6 *Detecting Affect and Emotions*

Sentiment analysis is most commonly used to refer to the goal of determining the valence or polarity of a piece of text. However, it can refer more generally to determining one's attitude towards a particular target or topic. Here, attitude can even mean emotional or affectual attitude such as frustration, joy, anger, sadness, excitement, and so on. Russell (1980) developed a circumplex model of affect and showed that it can be characterized by two primary dimensions: valence (positive and negative dimension) and arousal (degree of reactivity to stimulus). Thus, it is not surprising that large amounts of work in sentiment analysis is focused on determining valence. However, there is barely any work on automatically detecting arousal and a relatively small amount of work on detecting emotions such as anger, frustration, sadness, and optimism (Strapparava and Mihalcea 2007; Aman and Szpakowicz 2007; Tokuhisa et al. 2008; Neviarouskaya et al. 2009; Bellegarda 2010; Mohammad 2012; Boucouvalas 2002; Zhe and Boucouvalas 2002; Holzman and Pottenger 2003; Ma et al. 2005; Mohammad 2012; John et al. 2006; Mihalcea and Liu 2006; Genereux and Evans 2006). Detecting these more subtle aspects of sentiment has wide-ranging applications, for example in developing customer relation models, public health, military intelligence, and the video games industry, where it is necessary to make distinctions between anger and sadness (both of which are negative), calm and excited (both of which are positive), and so on.

4.3 Sentiment of Words

Term–sentiment associations have been captured by manually created sentiment lexicons as well as automatically generated ones.

4.3.1 *Manually Generated Term-Sentiment Association Lexicons*

The General Inquirer (GI) has sentiment labels for about 3,600 terms (Stone et al. 1966). Hu and Liu (2004) manually labeled about 6,800 words and used them for detecting sentiment of customer reviews. The MPQA Subjectivity Lexicon, which draws from the General Inquirer and other sources, has sentiment labels for about 8,000 words (Wilson et al. 2005). The NRC Emotion Lexicon has sentiment and emotion labels for about 14,000 words (Mohammad and Turney 2010; Mohammad and Yang 2011). These labels were compiled through Mechanical Turk annotations.³

For people, assigning a score indicating the degree of sentiment is not natural. Different people may assign different scores to the same target item, and it is hard for even the same annotator to remain consistent when annotating a large number of items. In contrast, it is easier for annotators to determine whether one word is more positive (or more negative) than the other. However, the latter requires a much larger number of annotations than the former (in the order of N^2 , where N is the number of items to be annotated).

An annotation scheme that retains the comparative aspect of annotation while still requiring only a small number of annotations comes from survey analysis techniques and is called MaxDiff (Louviere 1991). The annotator is presented with four terms and asked which word is the most positive and which is the least positive. By answering just these two questions five out of the six inequalities are known. If the respondent says that A is most positive and D is least positive, then:

$$A > B, A > C, A > D, B > D, C > D$$

Each of these MaxDiff questions can be presented to multiple annotators. The responses to the MaxDiff questions can then be easily translated into a ranking of all the terms and also a real-valued score for all the terms (Orme 2009). If two words have very different degrees of association (for example, $A \gg D$), then A will be chosen as most positive much more often than D and D will be chosen as least positive much more often than A . This will eventually lead to a ranked list such that A and D are significantly farther apart, and their real-valued association scores are also significantly different. On the other hand, if two words have similar degrees

³<https://www.mturk.com/mturk/welcome>

of association with positive sentiment (for example, A and B), then it is possible that for MaxDiff questions having both A and B , some annotators will choose A as most positive, and some will choose B as most positive. Further, both A and B will be chosen as most positive (or most negative) a similar number of times. This will result in a list such that A and B are ranked close to each other and their real-valued association scores will also be close in value.

MaxDiff was used for obtaining annotations of relation similarity of pairs of items in a SemEval-2012 shared task (Jurgens et al. 2012). Kiritchenko and Mohamad (2016a) applied Best–Worst Scaling to obtain real-valued sentiment association scores for words and phrases in three different domains: general English, English Twitter, and Arabic Twitter. They showed that on all three domains the ranking of words by sentiment remains remarkably consistent even when the annotation process is repeated with a different set of annotators. They also determine the minimum difference in sentiment association that is perceptible to native speakers of a language.

4.3.2 *Automatically Generated Term-Sentiment Association Lexicons*

Semi-supervised and automatic methods have also been proposed to detect the polarity of words. Hatzivassiloglou and McKeown (1997) proposed an algorithm to determine the polarity of adjectives. SentiWordNet was created using supervised classifiers as well as manual annotation (Esuli and Sebastiani 2006). Turney and Littman (2003) proposed a minimally supervised algorithm to calculate the polarity of a word by determining if its tendency to co-occur with a small set of positive seed words is greater than its tendency to co-occur with a small set of negative seed words. Mohammad et al. (2013b) employed the Turney method to generate a lexicon (Hashtag Sentiment Lexicon) from tweets with certain sentiment-bearing seed-word hashtags such as (*#excellent*, *#good*, *#terrible*, and so on) and another lexicon (Hashtag Sentiment Lexicon) from tweets with emoticons.⁴ Since the lexicons themselves are generated from tweets, they even have entries for the creatively spelled words (e.g. *happpeee*), slang (e.g. *bling*), abbreviations (e.g. *lol*), and even hashtags and conjoined words (e.g. *#loveumom*). Cambria et al. (2016) created SenticNet that has sentiment entries for 30,000 words and multi-word expressions using information propagation to connect various parts of common-sense knowledge representations. Kiritchenko et al. (2014b) proposed a method to create separate lexicons for words found in negated context and those found in affirmative context; the idea being that the same word contributes to sentiment differently depending on whether it is negated or not. These lexicons contain sentiment associations for hundreds

⁴<http://www.purl.com/net/lexicons>

of thousands of unigrams and bigrams. However, they do not explicitly handle combinations of terms with modals, degree adverbs, and intensifiers.

4.4 Sentiment of Phrases, Sentences, and Tweets: Sentiment Composition

Semantic composition, which aims at determining a representation of the meaning of two words through manipulations of their individual representations, has gained substantial attention in recent years with work from Mitchell and Loapata (2010), Baroni and Zamparelli (2010), Rudolph and Giesbrecht (2010), Yessenalina and Cardie (2011), Grefenstette et al. (2013), Grefenstette and Sadrzadeh (2011), and Turney (2014). Socher et al. (2012) and Mikolov et al. (2013) introduced deep learning models and distributed word representations in vector space (word embeddings) to obtain substantial improvements over the state-of-the-art in semantic composition. Mikolov's word2vec tool for generating word embeddings is available publicly.⁵

Sentiment of a phrase or a sentence is often not simply the sum of the sentiments of its constituents. Sentiment composition is the determining of sentiment of a multi-word linguistic unit, such as a phrase or a sentence, based on its constituents. Lexicons that include sentiment associations for phrases as well as their constituent words are referred to as *sentiment composition lexicons (SCLs)*. Kiritchenko and Mohammad created sentiment composition lexicons for English and Arabic that included: (1) negated expressions Kiritchenko and Mohammad (2016a,b), (2) phrases with adverbs, modals, and intensifies Kiritchenko and Mohammad (2016a,b), and (3) opposing polarity phrases (where at least one word in the phrase is positive and at least one word is negative, for example, *happy accident* and *dark chocolate*) (Kiritchenko and Mohammad 2016c). Socher et al. (2013) took a dataset of movie review sentences that were annotated for sentiment and further annotated every word and phrasal constituent within those sentences for sentiment. Such datasets where sentences, phrases, and their constituent words are annotated for sentiment are helping foster further research on how sentiment is composed. We discuss specific types of sentiment composition, and challenges for automatic methods that address them, in the sub-sections below.

4.4.1 Negated Expressions

Morante and Sporleder (2012) define negation to be “a grammatical category that allows the changing of the truth value of a proposition”. Negation is often expressed

⁵<https://code.google.com/p/word2vec>

through the use of negative signals or negator words such as *not* and *never*, and it can significantly affect the sentiment of its scope. Understanding the impact of negation on sentiment improves automatic analysis of sentiment. Earlier works on negation handling employ simple heuristics such as flipping the polarity of the words in a negator's scope (Kennedy and Inkpen 2005; Choi and Cardie 2008) or changing the degree of sentiment of the modified word by a fixed constant (Taboada et al. 2011). Zhu et al. (2014) show that these simple heuristics fail to capture the true impact of negators on the words in their scope. They show that negators tend to often make positive words negative (albeit with lower intensity) and make negative words less negative (not positive). Zhu et al. also propose certain embeddings-based recursive neural network models to capture the impact of negators more precisely. As mentioned earlier, Kiritchenko et al. (2014b) capture the impact of negation by creating separate sentiment lexicons for words seen in affirmative context and those seen in negated contexts. They use a hand-chosen list of negators and determine scope to be starting from the negator and ending at the first punctuation (or end of sentence).

Several aspects about negation are still not understood though: for example, can negators be ranked in terms of their average impact on the sentiment of their scopes (which negators impact sentiment more and which impact sentiment less); in what contexts does the same negator impact the sentiment of its scope more and in what contexts is the impact less; how do people in different communities and cultures use negations differently; and how negations of sentiment expressions should be dealt with by paraphrase and textual entailment systems.

4.4.2 *Phrases with Degree Adverbs, Intensifiers, and Modals*

Degree adverbs such as *barely*, *moderately*, and *slightly* quantify the extent or amount of the predicate. Intensifiers such as *too* and *very* are modifiers that do not change the propositional content (or truth value) of the predicate they modify, but they add to the emotionality. However, even linguists are hard pressed to give out comprehensive lists of degree adverbs and intensifiers. Additionally, the boundaries between degree adverbs and intensifiers can sometimes be blurred, and so it is not surprising that the terms are occasionally used interchangeably. Impacting propositional content or not, both degree adverbs and intensifiers impact the sentiment of the predicate, and there is some work in exploring this interaction (Zhang et al. 2008; Wang and Wang 2012; Xu et al. 2008; Lu and Tsou 2010; Taboada et al. 2008). Most of this work focuses on identifying sentiment words by bootstrapping over patterns involving degree adverbs and intensifiers. Thus several areas remain unexplored, such as identifying patterns and regularities in how different kinds of degree adverbs and intensifiers impact sentiment, ranking degree adverbs and intensifiers in terms of how they impact sentiment, and determining when (in what contexts) the same modifier will impact sentiment differently than

its usual behavior. (See Kiritchenko and Mohammad (2016b) for some recent work exploring these questions in manually annotated sentiment composition lexicons.)

Modals (a kind of auxiliary verb) are used to convey the degree of confidence, permission, or obligation to the predicate. Thus, if the predicate is sentiment bearing, then the sentiment of the combination of the modal and the predicate can be different from the sentiment of the predicate alone. For example, *cannot work* seems less positive than *work* or *will work* (*cannot* and *will* are modals). There is little work on automatically determining the impact of modals on sentiment.

4.4.3 *Sentiment of Sentences, Tweets, and SMS messages*

Bag-of-words models such as the NRC-Canada system (Mohammad et al. 2013a; Kiritchenko et al. 2014a,b) and Unitn Severyn and Moschitti (2015) have been very successful in recent shared task competitions on determining sentiment of whole tweets, SMS messages, and sentences. However, approaches that apply systematic sentiment composition of smaller units to determine sentiment of sentences are growing in popularity. Socher et al. (2013) proposed a word-embeddings based model that learns the *sentiment* of term compositions. They obtain state-of-the-art results in determining both the overall sentiment and sentiment of constituent phrases in movie review sentences. This has inspired tremendous interest in more embeddings-based work for sentiment composition (Dong et al. 2014; Kalchbrenner et al. 2014). These recursive models do not require any hand-crafted features or semantic knowledge, such as a list of negation words or sentiment lexicons. However, they are computationally intensive and need substantial additional annotations (word and phrase-level sentiment labeling). Nonetheless, use of word-embeddings in sentiment composition is still in its infancy, and we will likely see much more work using these techniques in the future.

4.4.4 *Sentiment in Figurative Expressions*

Figurative expressions in text, by definition, are not compositional. That is, their meaning cannot fully be derived from the meaning of their components in isolation. There is growing interest in detecting figurative language, especially irony and sarcasm (Carvalho et al. 2009; Reyes et al. 2013; Veale and Hao 2010; Filatova 2012; González-Ibáñez et al. 2011). In 2015, a SemEval shared task was organized on detecting sentiment in tweets rich in metaphor and irony (Task 11).⁶ Participants were asked to determine the degree of sentiment for each tweet where the score is a real number in the range from -5 (most negative) to $+5$ (most positive). One of

⁶The proceedings will be released later in 2015.

the characteristics of the data is that a large majority is negative; thereby suggesting that ironic tweets are largely negative. The SemEval 2014 shared task Sentiment Analysis in Twitter Rosenthal et al. (2014a) had a separate test set involving sarcastic tweets. Participants were asked *not* to train their system on sarcastic tweets, but rather apply their regular sentiment system on this new test set; the goal was to determine performance of regular sentiment systems on sarcastic tweets. It was observed that the performances dropped by about 25% to 70%, thereby showing that systems must be adjusted if they are to be applied to sarcastic tweets. We found little to no work exploring automatic sentiment detection in hyperbole, understatement, rhetorical questions, and other creative uses of language.

4.5 Challenges in Annotating for Sentiment

Clear and simple instructions are crucial for obtaining high-quality annotations. This is true even for seemingly simple annotation tasks, such as sentiment annotation, where one is to label instances as positive, negative, or neutral. For word annotations, researchers have often framed the task as ‘is this word positive, negative, or neutral?’ Hu and Liu (2004), ‘does this word have associations with positive, negative, or neutral sentiment?’ Mohammad and Turney (2013), or ‘which word is more positive?’/‘which word has a greater association with positive sentiment’ (Kiritchenko et al. 2016; Kiritchenko and Mohammad 2016c). Similar instructions are also widely used for sentence-level sentiment annotations—‘is this sentence positive, negative, or neutral?’ (Rosenthal et al. 2015, 2014b; Mohammad et al. 2016a, 2015b). We will refer to such annotation schemes as *the simple sentiment questionnaires*. On the one hand, this characterization of the task is simple, terse, and reliant on the intuitions of native speakers of a language (rather than biasing the annotators by providing definitions of what it means to be positive, negative, and neutral). On the other hand, the lack of specification leaves the annotator in doubt over how to label certain kinds of instances—for example, sentences where one side wins against another, sarcastic sentences, or retweets.

A different approach to sentiment annotation is to ask respondents to identify the target of opinion, and the sentiment towards this target of opinion (Pontiki et al. 2014b; Mohammad et al. 2015b; Deng and Wiebe 2014). We will refer to such annotation schemes as *the semantic-role based sentiment questionnaires*. This approach of sentiment annotation is more specific, and more involved, than the simple sentiment questionnaire approach; however, it too is insufficient for handling several scenarios. Most notably, the emotional state of the speaker is not under the purview of this scheme. Many applications require that statements expressing positive or negative emotional state of the speaker should be marked as ‘positive’ or ‘negative’, respectively. Similarly, many applications require statements that describe positive or negative events or situations to be marked as ‘positive’ or ‘negative’, respectively. Instructions for annotating opinion towards targets do not

specify how such instances are to be annotated, and worse still, possibly imply that such instances are to be labeled as neutral.

Some sentence types that are especially challenging for sentiment annotation (using either the simple sentiment questionnaire or the semantic-role based sentiment questionnaire) are listed below:

- *Speaker's emotional state*: The speaker's emotional state may or may not have the same polarity as the opinion expressed by the speaker. For example, a politician's tweet can imply both a negative opinion about a rival's past indiscretion, and a joyous mental state as the news will impact the rival adversely.
- *Success or failure of one side w.r.t. another*: Often sentences describe the success or failure of one side w.r.t. another side—for example, 'Yay! France beat Germany 3-1', 'Supreme court judges in favor of gay marriage', and 'the coalition captured the rebels'. If one supports France, gay marriage, and the coalition, then these events are positive, but if one supports Germany, marriage as a union only between man and woman, and the rebels, then these events can be seen as negative.

Also note that the framing of an event as the success of one party (or as the failure of another party) does not automatically imply that the speaker is expressing positive (or negative) opinion towards the mentioned party. For example, when Finland beat Russia in ice hockey in the 2014 Sochi Winter Olympics, the event was tweeted around the world predominantly as "Russia lost to Finland" as opposed to "Finland beat Russia". This is not because the speakers were expressing negative opinion towards the Russian team, but rather simply because Russia, being the host nation, was the focus of attention and traditionally Russian hockey teams have been strong.

- *Neutral reporting of valenced information*: If the speaker does not give any indication of her own emotional state but describes valenced events or situations, then it is unclear whether to consider these statements as neutral unemotional reporting of developments or whether to assume that the speaker is in a negative emotional state (sad, angry, etc.). Example:

The war has created millions of refugees.

- *Sarcasm and ridicule*: Sarcasm and ridicule are tricky from the perspective of assigning a single label of sentiment because they can often indicate positive emotional state of the speaker (pleasure from mocking someone or something) even though they have a negative attitude towards someone or something.
- *Different sentiment towards different targets of opinion*: The speaker may express opinion about multiple targets, and sentiment towards the different targets might be different. The targets may be different people or objects (for example, an iPhone vs. an android phone), or they may be different aspects of the same entity (for example, quality of service vs. quality of food at a restaurant).
- *Precisely determining the target of opinion*: Sometimes it is difficult to precisely identify the target of opinion. For example, consider:

Glad to see Hillary's lies being exposed.

It is unclear whether the target of opinion is ‘Hillary’, ‘Hillary’s lies’, or ‘Hillary’s lies being exposed’. One reasonable interpretation is that positive sentiment is expressed about ‘Hillary’s lies being exposed’. However, one can also infer that the speaker has a negative attitude towards ‘Hillary’s lies’ and probably ‘Hillary’ in general. It is unclear whether annotators should be asked to provide all three opinion–target pairs or only one (in which case, which one?).

- *Supplications and requests*: Many tweets convey positive supplications to God or positive requests to people in the context of a (usually) negative situation. Examples include:

May god help those displaced by war.

Let us all come together and say no to fear mongering and divisive politics.

- *Rhetorical questions*: Rhetorical questions can be treated simply as queries (and thus neutral) or as utterances that give away the emotional state of the speaker. For example, consider:

Why do we have to quibble every time?

On the one hand, this tweet can be treated as a neutral question, but on the other hand, it can be seen as negative because the utterance betrays a sense of frustration on the part of the speaker.

- *Quoting somebody else or re-tweeting*: Quotes and retweets are difficult to annotate for sentiment because it is often unclear and not explicitly evident whether the one who quotes (or retweets) holds the same opinions as that expressed by the quotee.

The challenges listed above can be addressed to varying degrees by providing instructions to the annotators on how such instances are to be labeled. However, detailed and complicated instructions can be counter-productive as the annotators may not understand or may not have the inclination to understand the subtleties involved. See Mohammad (2016a) for annotation schemes that address some of these challenges.

4.6 Challenges in Multilingual Sentiment Analysis

Work on multilingual sentiment analysis has mainly addressed mapping sentiment resources from English into morphologically complex languages. Mihalcea et al. (2007) use English resources to automatically generate a Romanian subjectivity lexicon using an English–Romanian dictionary. The generated lexicon is then used to classify Romanian text. Wan (2008) translated Chinese customer reviews to English using a machine translation system. The translated reviews are then annotated using rule-based system that uses English lexicons. A higher accuracy is achieved when using ensemble methods and combining knowledge from Chinese and English resources. Balahur and Turchi (2014) conducted a study to assess the performance of statistical sentiment analysis techniques on machine-translated

texts. Opinion-bearing phrases from the New York Times Text (2002–2005) corpus were automatically translated using publicly available machine-translation engines (Google, Bing, and Moses). Then, the accuracy of a sentiment analysis system trained on original English texts was compared to the accuracy of the system trained on automatic translations to German, Spanish, and French. The authors conclude that the quality of machine translation is acceptable for sentiment analysis to be performed on automatically translated texts. Salameh et al. (2015) conducted experiments to determine loss in sentiment predictability when they translate Arabic social media posts into English, manually and automatically. As benchmarks, they use manually and automatically determined sentiment labels of the Arabic texts. They show that sentiment analysis of English translations of Arabic texts produces competitive results, w.r.t. Arabic sentiment analysis. They also claim that even though translation significantly reduces human ability to recover sentiment, automatic sentiment systems are affected relatively less by this.

Some of the areas less explored in the realm of multilingual sentiment analysis include: how to translate text so as to preserve the degree of sentiment in the source text; how sentiment modifiers such as negators and modals differ in function across languages; understanding how automatic translations differ from manual translations in terms of sentiment; and how to translate figurative language without losing its affectual gist.

4.7 Challenges in Applying Sentiment Analysis

Applications of sentiment analysis benefit from the fact that even though systems are not extremely accurate at determining sentiment of individual sentences, they can accurately capture significant changes in the proportion of instances that are positive (or negative). It is also worth noting that such sentiment tracking systems are more effective when incorporating carefully chosen baselines. For example, knowing the percentage of tweets that are negative towards Russian President, Vladimir Putin, is less useful than, for instance, knowing: the percentage of tweets that are negative towards Putin before vs. after the invasion of Crimea; or, the percentage of tweets that are negative towards Putin in Russia vs. the rest of the world; or, the percentage of tweets negative towards Putin vs. Barack Obama (US president).

Sentiment analysis is commonly applied in several areas including tracking sentiment towards products, movies, politicians, and companies (O'Connor et al. 2010; Pang and Lee 2008), improving customer relation models (Bougie et al. 2003), detecting happiness and well-being (Schwartz et al. 2013), tracking the stock market (Bollen et al. 2011), and improving automatic dialogue systems (Velásquez 1997; Ravaja et al. 2006). The sheer volume of work in this area precludes detailed summarization here. Nonetheless, it should be noted that often the desired application can help direct certain design choices in the sentiment analysis system. For example, the threshold between neutral and positive sentiment and the threshold between neutral and negative sentiment can be determined empirically by what

- Carvalho, P., L. Sarmento, M.J. Silva, and E. De Oliveira, 2009. Clues for detecting irony in user-generated contents: oh...!! it's so easy;-). In *Proceedings of the 1st International CIKM Workshop on Topic-Sentiment Analysis for Mass Opinion*, 53–56. ACM.
- Choi, Y., and C. Cardie. 2008. Learning with compositional semantics as structural inference for subsentential sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP'08, Honolulu, 793–801.
- Colton, S., J. Goodwin, and T. Veale. 2012. Full face poetry generation. In *Proceedings of the Third International Conference on Computational Creativity*, 95–102.
- Davis, H., and S. Mohammad. 2014. Generating music from literature. In *Proceedings of the 3rd Workshop on Computational Linguistics for Literature (CLFL)*, Gothenburg, 1–10.
- Deng, L., and J. Wiebe. 2014. Sentiment propagation via implicature constraints. In *EACL*, 377–385.
- Diman Ghazi, D.I., and S. Szpakowicz. 2015. Detecting emotion stimuli in emotion-bearing sentences. In: *Proceedings of the 2015 Conference on Intelligent Text Processing and Computational Linguistics*.
- Dong, L., F. Wei, M. Zhou, and K. Xu. 2014. Adaptive multi-compositionality for recursive neural models with applications to sentiment analysis. In *Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI)*.
- Eichstaedt, J.C., H.A. Schwartz, M.L. Kern, G. Park, D.R. Labarthe, R.M. Merchant, S. Jha, M. Agrawal, L.A. Dziurzynski, and M. Sap et al. 2015. Psychological language on twitter predicts county-level heart disease mortality. *Psychological Science* 2: 159–169.
- Esuli, A., and F. Sebastiani. 2006. SENTIWORDNET: A publicly available lexical resource for opinion mining. In *In Proceedings of the 5th Conference on Language Resources and Evaluation*, LREC'06, 417–422.
- Filatova, E. 2012. Irony and sarcasm: Corpus generation and analysis using crowdsourcing. In *LREC*, 392–398.
- Genereux, M., and R.P. Evans. 2006. Distinguishing affective states in weblogs. In *AAAI-2006 Spring Symposium on Computational Approaches to Analysing Weblogs*, Stanford, 27–29.
- Gildea, D., and D. Jurafsky. 2002. Automatic labeling of semantic roles. *Computational Linguistics* 28(3): 245–288.
- González-Ibáñez, R., S. Muresan, and N. Wacholder. 2011. Identifying sarcasm in twitter: A closer look. In *Proceedings of the ACL*, 581–586.
- Grefenstette, E., G. Dinu, Y.-Z. Zhang, M. Sadrzadeh, and M. Baroni. 2013. Multi-step regression learning for compositional distributional semantics. arXiv preprint arXiv:1301.6939.
- Grefenstette, E., and M. Sadrzadeh. 2011. Experimental support for a categorical compositional distributional model of meaning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 1394–1404.
- Hasan, Kazi Saidul, and Vincent Ng. 2013. Stance classification of ideological debates: Data, models, features, and constraints. In *The 6th international joint conference on natural language processing*, Nagoya.
- Hassan, A., A. Abu-Jbara, and D. Radev. 2012. Extracting signed social networks from text. In *Workshop Proceedings of TextGraphs-7 on Graph-Based Methods for Natural Language Processing*, 6–14.
- Hatzivassiloglou, V., and K.R. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the 8th Conference of European Chapter of the Association for Computational Linguistics*, Madrid, 174–181.
- Holzman, L.E., and W.M. Pottenger. 2003. Classification of emotions in internet chat: An application of machine learning using speech phonemes. Technical report, Leigh University.
- Hu, M., and B. Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD'04, New York, 168–177.
- John, D., A.C. Boucouvalas, and Z. Xu. 2006. Representing emotional momentum within expressive internet communication. In *Proceedings of the 24th IASTED International Conference on Internet and Multimedia Systems and Applications*, 183–188. Anaheim: ACTA Press.

- Jurgens, D., S.M. Mohammad, P. Turney, and K. Holyoak. 2012. Semeval-2012 task 2: Measuring degrees of relational similarity. In *Proceedings of the 6th International Workshop on Semantic Evaluation, SemEval'12*, Montréal, 356–364.
- Kalchbrenner, N., E. Grefenstette, and P. Blunsom. 2014. A convolutional neural network for modelling sentences. arXiv preprint arXiv:1404.2188.
- Kennedy, A., and D. Inkpen. 2005. Sentiment classification of movie and product reviews using contextual valence shifters. In *Proceedings of the Workshop on the Analysis of Informal and Formal Information Exchange During Negotiations*, Ottawa
- Kessler, J.S., and N. Nicolov. 2009. Targeting sentiment expressions through supervised ranking of linguistic configurations. In *3rd Int'l AAAI Conference on Weblogs and Social Media (ICWSM 2009)*.
- Kiritchenko, S., and S.M. Mohammad. 2016a. Capturing reliable fine-grained sentiment associations by crowdsourcing and best–worst scaling. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, San Diego.
- Kiritchenko, S., and S.M. Mohammad. 2016b. The effect of negators, modals, and degree adverbs on sentiment composition. In *Proceedings of the Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA)*.
- Kiritchenko, S., and S.M. Mohammad. 2016c. Sentiment composition of words with opposing polarities. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, San Diego.
- Kiritchenko, S., S.M. Mohammad, and M. Salameh. 2016. Semeval-2016 task 7: Determining sentiment intensity of English and arabic phrases. In *Proceedings of the International Workshop on Semantic Evaluation, SemEval-2016*, San Diego.
- Kiritchenko, S., X. Zhu, C. Cherry, and S. Mohammad. 2014a. Nrc-canada-2014: Detecting aspects and sentiment in customer reviews. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Dublin, 437–442.
- Kiritchenko, S., X. Zhu, and S.M. Mohammad. 2014b. Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research* 50: 723–762.
- Ku, L.-W., Y.-T. Liang, and H.-H. Chen. 2006. Opinion extraction, summarization and tracking in news and blog corpora. In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*, vol. 100107.
- Li, W., and H. Xu. 2014. Text-based emotion classification using emotion cause extraction. *Expert Systems with Applications* 41(4, Part 2): 1742–1749.
- Liu, B., and L. Zhang. 2012. A survey of opinion mining and sentiment analysis. In *Mining Text Data*, ed. C.C. Aggarwal and C. Zhai, 415–463. New York: Springer
- Liu, J., Y. Cao, C.-Y. Lin, Y. Huang, and M. Zhou. 2007. Low-quality product review detection in opinion summarization. In *EMNLP-CoNLL*, 334–342.
- Lloret, E., A. Balahur, M. Palomar, and A. Montoyo. 2009. Towards building a competitive opinion summarization system: challenges and keys. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, Companion Volume: S, 72–77.
- Louvière, and J.J. 1991. Best-worst scaling: A model for the largest difference judgments. Working Paper.
- Lu, B., and B.K. Tsou. 2010. Cityu-dac: Disambiguating sentiment-ambiguous adjectives within context. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, 292–295.
- Ma, C., H. Prendinger, and M. Ishizuka. 2005. Emotion estimation and reasoning based on affective textual interaction. In *First International Conference on Affective Computing and Intelligent Interaction (ACII-2005)*, ed. J. Tao, R.W. Picard, Beijing, 622–628.
- Màrquez, L., X. Carreras, K.C. Litkowski, and S. Stevenson. 2008. Semantic role labeling: An introduction to the special issue. *Computational Linguistics* 34(2): 145–159.

- Mihalcea, R., C. Banea, and J. Wiebe. 2007. Learning multilingual subjective language via cross-lingual projections. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*.
- Mihalcea, R., and H. Liu. 2006. A corpus-based approach to finding happiness. In *AAAI-2006 Spring Symposium on Computational Approaches to Analysing Weblogs*, 139–144. AAAI Press.
- Mikolov, T., I. Sutskever, K. Chen, G.S. Corrado, and J. Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, 3111–3119.
- Mitchell, J., and M. Loapata. 2010. Composition in distributional models of semantics. *Cognitive Science* 34(8): 1388–1429.
- Mohammad, S. 2012. Portable features for classifying emotional text. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Montréal, 587–591.
- Mohammad, S., S. Kiritchenko, and X. Zhu. 2013a. Nrc-canada: Building the state-of-the-art in sentiment analysis of tweets. In *Proceedings of the Seventh International Workshop on Semantic Evaluation Exercises (SemEval-2013)*, Atlanta.
- Mohammad, S., S. Kiritchenko, and X. Zhu. 2013b. NRC-Canada: Building the state-of-the-art in sentiment analysis of tweets. In *Proceedings of the International Workshop on Semantic Evaluation*, SemEval'13, Atlanta.
- Mohammad, S., and T. Yang. 2011. Tracking sentiment in mail: How genders differ on emotional axes. In *Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA 2011)*, Portland, 70–79.
- Mohammad, S.M. 2012. #emotional tweets. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation*, SemEval'12, Stroudsburg, 246–255.
- Mohammad, S.M. 2016a. A practical guide to sentiment annotation: Challenges and solutions. In *Proceedings of the Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*.
- Mohammad, S.M. 2016b. Sentiment analysis: Detecting valence, emotions, and other affectual states from text. In *Emotion Measurement*, ed. H. Meiselman. Amsterdam: Elsevier.
- Mohammad, S.M., S. Kiritchenko, P. Sobhani, X. Zhu, and C. Cherry. 2016a. Semeval-2016 task 6: Detecting stance in tweets. In *Proceedings of the International Workshop on Semantic Evaluation*, SemEval'16, San Diego.
- Mohammad, S.M., P. Sobhani, and S. Kiritchenko. 2016b, In Press. Stance and Sentiment in Tweets. *Special Section of the ACM Transactions on Internet Technology on Argumentation in Social Media*.
- Mohammad, S.M., and P.D. Turney. 2010. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL-HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, California.
- Mohammad, S.M., and P.D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence* 29(3): 436–465.
- Mohammad, S.M., X. Zhu, S. Kiritchenko, and J. Martin. 2015a. Sentiment, emotion, purpose, and style in electoral tweets. *Information Processing & Management* 51: 480–499.
- Mohammad, S.M., X. Zhu, S. Kiritchenko, and J. Martin. 2015b. Sentiment, emotion, purpose, and style in electoral tweets. *Information Processing and Management* 51(4): 480–499.
- Murakami, Akiko, and Rudy Raymond. 2010. Support or oppose?: Classifying positions in online debates from reply activities and opinion expressions. In *Proceedings of the 23rd international conference on computational linguistics*, Beijing.
- Nahar, V., S. Unanikard, X. Li, and C. Pang. 2012. Sentiment analysis for effective detection of cyber bullying. In *Web Technologies and Applications*, 767–774. Berlin/Heidelberg: Springer.

- Nalisnick, E.T., and H.S. Baird. 2013a. Character-to-character sentiment analysis in Shakespeare's plays. In *Proceedings of the 51st annual meeting of the association for computational linguistics (ACL)*, Short Paper, Sofia, 479–483, Aug 2013.
- Nalisnick, E.T., and H.S. Baird. 2013b. Extracting sentiment networks from Shakespeare's plays. In *2013 12th International Conference on Document Analysis and Recognition (ICDAR)*, 758–762. IEEE.
- Neviarouskaya, A., H. Prendinger, and M. Ishizuka. 2009. Compositionality principle in recognition of fine-grained emotions from text. In *Proceedings of the Proceedings of the Third International Conference on Weblogs and Social Media (ICWSM-09)*, 278–281, San Jose.
- O'Connor, B., R. Balasubramanyan, B.R. Routledge, and N.A. Smith. 2010. From tweets to polls: Linking text sentiment to public opinion time series. In *Proceedings of the International AAAI Conference on Weblogs and Social Media*.
- Orme, B. 2009. *Maxdiff analysis: Simple counting, individual-level logit, and HB*. Orem: Sawtooth Software, Inc.
- Pak, A., and P. Paroubek. 2010. Twitter as a corpus for sentiment analysis and opinion mining. In *Proceedings of the 7th Conference on International Language Resources and Evaluation, LREC'10*, Valletta.
- Palmer, M., D. Gildea, and N. Xue. 2010. Semantic role labeling. *Synthesis Lectures on Human Language Technologies* 3(1): 1–103.
- Pang, B., and L. Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* 2(1–2): 1–135.
- Pontiki, M., D. Galanis, J. Pavlopoulos, H. Papageorgiou, I. Androutsopoulos, and S. Manandhar. 2014a. SemEval-2014 Task 4: Aspect based sentiment analysis. In *Proceedings of the International Workshop on Semantic Evaluation, SemEval'14*, Dublin.
- Pontiki, M., H. Papageorgiou, D. Galanis, I. Androutsopoulos, J. Pavlopoulos, and S. Manandhar. 2014b. SemEval-2014 task 4: Aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation, SemEval'14*, Dublin.
- Popescu, A.-M., and O. Etzioni. 2005. Extracting product features and opinions from reviews. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing, HLT'05*, Stroudsburg, 339–346.
- Qadir, A. 2009. Detecting opinion sentences specific to product features in customer reviews using typed dependency relations. In *Proceedings of the Workshop on Events in Emerging Text Types, eETTs'09*, Stroudsburg, 38–43.
- Ravaja, N., T. Saari, M. Turpeinen, J. Laarni, M. Salminen, and M. Kivikangas. 2006. Spatial presence and emotions during video game playing: Does it matter with whom you play? *Presence: Teleoperators and Virtual Environments* 15(4): 381–392.
- Reyes, A., P. Rosso, and T. Veale. 2013. A multidimensional approach for detecting irony in twitter. *Language Resources and Evaluation* 47(1): 239–268.
- Rosenthal, S., P. Nakov, S. Kiritchenko, S. Mohammad, A. Ritter, and V. Stoyanov. 2015. SemEval-2015 task 10: Sentiment analysis in Twitter. In *Proceedings of the 9th International Workshop on Semantic Evaluation, SemEval'15*, Denver, 450–462.
- Rosenthal, S., P. Nakov, A. Ritter, and V. Stoyanov. 2014a. SemEval-2014 Task 9: Sentiment analysis in Twitter. In *Proceedings of the 8th International Workshop on Semantic Evaluation*, ed. P. Nakov, and T. Zesch, SemEval-2014, Dublin.
- Rosenthal, S., A. Ritter, P. Nakov, and V. Stoyanov. 2014b. SemEval-2014 Task 9: Sentiment analysis in Twitter. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Dublin, 73–80.
- Rudolph, S., and E. Giesbrecht. 2010. Compositional matrix-space models of language. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, 907–916.
- Russell, J.A. 1980. A circumplex model of affect. *Journal of Personality and Social Psychology* 39(6): 1161.

- Salameh, M., S.M. Mohammad, and S. Kiritchenko. 2015. Sentiment after translation: A case-study on arabic social media posts. In *Proceedings of the North American Chapter of Association of Computational Linguistics*, Denver.
- Schwartz, H., J. Eichstaedt, M. Kern, L. Dziurzynski, R. Lucas, M. Agrawal, G. Park, et al. 2013. Characterizing geographic variation in well-being using tweets. In *Proceedings of the International AAAI Conference on Weblogs and Social Media*.
- Severyn, A., and A. Moschitti. 2015. Unltn: Training deep convolutional neural network for twitter sentiment classification. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, 464–469. Denver: Association for Computational Linguistics.
- Socher, R., B. Huval, C.D. Manning, and A.Y. Ng. 2012. Semantic compositionality through recursive matrix-vector spaces. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP'12, Jeju.
- Socher, R., A. Perelygin, J.Y. Wu, J. Chuang, C.D. Manning, A.Y. Ng, and C. Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP'13, Seattle.
- Somasundaran, Swapna, and Janyce Wiebe. 2009. Recognizing stances in online debates. In *Proceedings of the joint conference of the 47th annual meeting of the ACL and the 4th international joint conference on natural language processing of the AFNLP*, Singapore.
- Somprasertsri, G., and P. Lalitrojwong. 2010. Mining feature-opinion in online customer reviews for opinion summarization. *Journal of Universal Computer Science* 16(6): 938–955.
- Sridhar, Dhanya, Lise Getoor, and Marilyn Walker. 2014. Collective stance classification of posts in online debate forums. In *Proceedings of the 52nd annual meeting of the association for computational linguistics*, Baltimore.
- Stone, P., D.C. Dunphy, M.S. Smith, D.M. Ogilvie, and associates. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. Cambridge, MA: The MIT Press.
- Stoyanov, V., and C. Cardie. 2006. Toward opinion summarization: Linking the sources. In *Proceedings of the Workshop on Sentiment and Subjectivity in Text*, 9–14.
- Strapparava, C., and R. Mihalcea. 2007. Semeval-2007 Task 14: Affective text. In *Proceedings of SemEval-2007*, Prague, 70–74.
- Su, Q., K. Xiang, H. Wang, B. Sun, and S. Yu. 2006. Using pointwise mutual information to identify implicit features in customer reviews. In *Proceedings of the 21st International Conference on Computer Processing of Oriental Languages: Beyond the Orient: The Research Challenges Ahead*, ICCPOL'06, 22–30. Berlin/Heidelberg: Springer.
- Taboada, M., J. Brooke, M. Tofiloski, K. Voll, and M. Stede. 2011. Lexicon-based methods for sentiment analysis. *Computational Linguistics* 37(2): 267–307.
- Taboada, M., K. Voll, and J. Brooke. 2008. Extracting sentiment as a function of discourse structure and topicality. *Simon Fraser University School of Computing Science Technical Report*.
- Thelwall, M., K. Buckley, and G. Paltoglou. 2011. Sentiment in Twitter events. *Journal of the American Society for Information Science and Technology* 62(2): 406–418.
- Thomas, Matt, Bo Pang, and Lillian Lee. 2006. Get out the vote: Determining support or opposition from congressional floor-debate transcripts. In *Proceedings of the 2006 conference on empirical methods in natural language processing*. Sydney: Association for Computational Linguistics.
- Tokuhisa, R., K. Inui, and Y. Matsumoto. 2008. Emotion classification using massive examples extracted from the web. In *Proceedings of the 22nd International Conference on Computational Linguistics – Volume 1*, COLING'08, 881–888.
- Turney, P., and M.L. Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems* 21(4): 315–346.
- Turney, P.D. 2014. Semantic composition and decomposition: From recognition to generation. arXiv preprint arXiv:1405.7908.
- Veale, T., and Y. Hao. 2010. Detecting ironic intent in creative comparisons. *ECAI* 215: 765–770.
- Velásquez, J.D. 1997. Modeling emotions and other motivations in synthetic agents. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Conference on Innovative Applications of Artificial Intelligence*, AAAI'97/IAAI'97, 10–15. AAAI Press.

- Walker, Marilyn A., et al. 2012. A corpus for research on deliberation and debate. In *proceedings of the eighth international conference on language resources and evaluation (LREC)*, Istanbul.
- Wan, X. 2008. Using bilingual knowledge and ensemble techniques for unsupervised Chinese sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP'08, 553–561.
- Wang, C., and F. Wang. 2012. A bootstrapping method for extracting sentiment words using degree adverb patterns. In *2012 International Conference on Computer Science and Service System (CSSS)*, 2173–2176. IEEE.
- Wilson, T., Z. Kozareva, P. Nakov, S. Rosenthal, V. Stoyanov, and A. Ritter. 2013. SemEval-2013 Task 2: Sentiment analysis in Twitter. In *Proceedings of the International Workshop on Semantic Evaluation*, SemEval'13, Atlanta.
- Wilson, T., J. Wiebe, and P. Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, HLT'05, Stroudsburg, 347–354.
- Xu, G., C.-R. Huang, and H. Wang. 2013. Extracting Chinese product features: Representing a sequence by a set of skip-bigrams. In *Proceedings of the 13th Chinese Conference on Chinese Lexical Semantics*, CLSW'12, 72–83. Berlin/Heidelberg: Springer.
- Xu, R., K.-F. Wong, Q. Lu, Y. Xia, and W. Li. 2008. Learning knowledge from relevant webpage for opinion analysis. In *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, 2008. WI-IAT'08.*, vol. 1, 307–313. IEEE.
- Yessenalina, A., and C. Cardie. 2011. Compositional matrix-space models for sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 172–182.
- Zarella, G., and A. Marsh. 2016. MITRE at SemEval-2016 Task 6: Transfer learning for Stance detection. In *Proceedings of the International Workshop on Semantic Evaluation*, SemEval'16, San Diego.
- Zhang, C., D. Zeng, Q. Xu, X. Xin, W. Mao, and F.-Y. Wang. 2008. Polarity classification of public health opinions in Chinese. In *Intelligence and Security Informatics*, 449–454. Berlin/Heidelberg: Springer.
- Zhang, L., and B. Liu. 2011. Identifying noun product features that imply opinions. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers – Volume 2*, HLT'11, 575–580.
- Zhang, L., B. Liu, S.H. Lim, and O'Brien-E. Strain. 2010. Extracting and ranking product features in opinion documents. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, COLING'10, Stroudsburg, 1462–1470.
- Zhe, X., and A. Boucouvalas. 2002. *Text-to-Emotion Engine for Real Time Internet Communication*. *Text-to-Emotion Engine for Real Time Internet Communication*, 164–168.
- Zhu, X., H. Guo, S. Mohammad, and S. Kiritchenko. 2014. An empirical study on the effect of negation words on sentiment. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Baltimore, 304–313.