CHAPTER

SEMANTIC ASPECTS IN SENTIMENT ANALYSIS

3

M. Nissim[a], V. Patti[b]

University of Groningen, Groningen, The Netherlands[a] University of Turin, Torino, Italy[b]

1 INTRODUCTION

The fact that semantics must play a crucial role in the sentiment interpretation of text is rather obvious, as even just considering the plain meaning of words can be very indicative ("*I liked it*" vs. "*I hated it*"). However, things are not that simple or straightforward for at least two reasons: (1) *meaning* is not so easy to define, detect, and extract automatically, and (2) sentiment analysis is often not just a matter of distinguishing positive from negative opinions, especially in recent developments.

In the 2015 SemEval campaign, four shared tasks were organized within the *Sentiment Analysis* track: a rather general task on *sentiment analysis in Twitter* (task 10 [1], with four subtasks), a task focused on figurative language, entitled "*Sentiment Analysis of Figurative Language in Twitter*" (task 11 [2]), an *aspect-based sentiment analysis* task (task 12 [3]), where systems had to identify aspects of entities and the sentiment expressed for each aspect, and a rather different task focused on events' polarity, entitled "*CLIPEval Implicit Polarity of Events*" (task 9 [4]). Within the ongoing SemEval 2016, there is also a task on detecting *stance* in tweets (task 6[1])—that is, detecting the position of the author with respect to a given target (against/in favor/neutral)—and one on determining sentiment intensity (task 7[2]). Some of these tasks provide datasets in more than one language. Additionally, a shared task on concept-level sentiment analysis has been organized recently in the context of the European Semantic Web Conference [5]. This fervent, current action on stimulating research, resources, and tools in this field by organizing more numerous and more complex tasks tells us not only that interest in sentiment analysis is growing but also that sentiment analysis is no longer just about detecting whether a given review or tweet is objective or subjective, and in the latter case it is whether positive or negative. Rather, it requires a more complex analysis and interpretation of messages that in turn must rely on deeper processing and understanding. Thus although it is true that semantics and semantic processing play a crucial role in this, we must see *how* this happens, from several points of view.

First, and following intuition, words are sentiment informative at a plain lexical semantics level ("good" is positive, "bad" is negative). This is reflected in the creation of sentiment and emotion lexica

[1] http://alt.qcri.org/semeval2016/task6/
[2] http://alt.qcri.org/semeval2016/task7/

Sentiment Analysis in Social Networks. http://dx.doi.org/10.1016/B978-0-12-804412-4.00003-6

and corpora that can be used in system development, also for languages other than English. Second, deeper linguistic processing is required to perform finer-grained tasks. For instance, in aspect-based sentiment analysis, entities and aspects must be identified, as well as relations among them, and one cannot rely on lexical semantics only; also, in irony detection, systems must incorporate some module that deals with figurative language. Third, even deeper text processing might not suffice for the level of analysis required, and might need to be complemented by *reasoning* over concepts, which could be done by exploitation of semantic resources that the Semantic Web community has to offer in this sense, such as web ontologies and semantic networks. Fourth, sentiment analysis stretches out to, and intersects with, other related areas, such as emotion and personality detection, so the semantics of words and text has to be determined at different levels of affect interpretation.

In this chapter, we review a large collection of semantic resources for sentiment analysis and show how semantics plays various roles in the development of sentiment-aware tools and resources. Specifically, we discuss how state-of-the-art semantic processing is used and adapted to fit the requirements of progressively finer-grained tasks; for example, how semantic information is exploited in statistical models, how advances in semantic similarity models can be ported to sentiment analysis, and how automated reasoning and semantic metadata processing can be used in this field. Through this review we highlight the interaction of sentiment analysis with related affect resources and processing.

2 SEMANTIC RESOURCES FOR SENTIMENT ANALYSIS

Affective information expressed in our texts is multifaceted. Both sentiment and emotion lexicons, and psycholinguistic resources available for English, refer to various affective models and capture different nuances of affect, such as sentiment polarity, emotional categories, and emotional dimensions. Such lexica are usually lists of words with which a positive or negative or emotion-related label (or score) is associated. Besides flat vocabularies of affective words, other resources include and model semantic, conceptual, and affective information associated with multiword natural language expressions, by enabling concept-level analysis of sentiment and emotions conveyed in texts. In our view all such resources represent a rich and varied lexical knowledge of affect, under different perspectives. Therefore we offer here a comprehensive description of such different resources and of their use in the context of sentiment analysis to distinguish between different opinions and sentiment.

2.1 CLASSICAL RESOURCES ON SENTIMENT

One of the first and most widely used resources is the Subjectivity Lexicon [6],[3] which is a list of subjectivity clues compiled from several sources, annotated both manually and automatically. This lexicon is the core of OpinionFinder, one of the first systems for the automatic detection of polarity. Another widely used resource is the Opinion Lexicon[4] compiled by Bing Liu. The list contains approximately 6800 English words that are classified as either positive or negative. Both resources were compiled manually and are thus quite accurate. However, they make two simplifications: first, they encode sentiment information in terms of a sharp division between a positive and negative sentiment

[3] http://mpqa.cs.pitt.edu/lexicons/subj_lexicon/
[4] https://www.cs.uic.edu/~liub/FBS/sentiment-analysis.html

value rather than providing a scale of positivity/negativity; second, they associate sentiment with words rather than with their *senses*. We will discuss the latter issue first, and the former later.

The lack of sense distinctions is quite a limitation from a semantic perspective. Indeed, lexical entries are often polysemous, so the same string might actually have a completely opposite sentiment depending on the context in which it is used. For example, “crazy” can be used in a negative as well as a positive way, which is strictly context dependent. Creating a flat list where words are assigned a binary polarity value will not account for the complexity of such semantic aspects.

One step further in creating sentiment resources is thus assigning polarity values at the *sense* level rather than the *word* level. This is the principle behind the annotation scheme developed in [7], which gave rise to the Subjectivity Sense Annotations,[5] a sense-aware lexicon. This resource actually addresses sentiment at one level up with respect to polarity, as it classifies a given word sense as objective or subjective, without specifying, in the latter case, its polarity value. As a plus, though, this lexicon includes part-of-speech information. Apart from sense distinctions of exactly the same lexical entries, there can also be ambiguities related to parts of speech, as a given word could be, for example, a noun and also an adjective, and could exhibit different polarity features accordingly. A good example is “novel,” which has a neutral polarity when used as a noun but a rather positive one when used as an adjective.

The creation of SentiWordNet [8] addresses ambiguities both within and across parts of speech. SentiWordNet is an extension of WordNet [9], where three sentiment scores are assigned to *each synset*, thus to word senses: positivity, negativity, and objectivity. In addition to the assigning of sentiment to senses, one nice feature of SentiWordNet is that sentiment is a *gradual* concept rather than a categorical one, thus addressing the first limitation of the simplest resources highlighted above. The same synset can exhibit the three values (positivity, negativity, neutrality) with different degrees, summing up to 1. The main drawback of SentiWordNet is that its sentiment scores are for the most part automatically assigned, so the presence of noise is not trivial. It is also known to have a *positive bias* [10], and so is not so good for detecting negative opinions.

The very first sentiment-aware resource we know of—namely, General Inquirer[6] [11], whose compilation was started in 1966—accounted for sense distinctions, although polarity was conceived as a binary value rather than a gradual concept.

Another way to indirectly address word sense disambiguation is to have *domain-dependent lexica*. Often a given word will exhibit only one sense in a specific domain (the classic example of “bank” is indicative here, as in a financial context it is unlikely this term would be used in its “river bank” sense). There are two large lexica that are built around very specific domains, exploiting contributors’ reviews: the Yelp Lexicon, with a focus on restaurant reviews, derived from the very large Yelp challenge dataset,[7] and the Amazon Lexicon, built on laptop reviews posted on *Amazon* from June 1995 to March 2013 [12,13]. Additionally, there are sentiment lists tailored for the financial domain, developed chiefly by Loughran and McDonald [14] after they noted that three quarters of the words with a negative polarity contained in standard, domain-independent resources did not actually have a negative connotation in the financial domain.

[5] http://mpqa.cs.pitt.edu/lexicons/subj_sense_annotations/
[6] http://www.wjh.harvard.edu/~inquirer/
[7] http://www.yelp.com/dataset_challenge

Among the resources that conceive polarity as a gradual rather than categorical value, we mention AFINN, which is a manually compiled list of 2477 English words and phrases rated with an integer between -5 (very negative) and $+5$ (very positive),[8] and the Sentiment140 lexicon, which contains about 60,000 unigrams and 10 times more bigrams. Differently from the AFINN lexicon, Sentiment140 is compiled automatically by exploiting tweets with emoticons. This resource is part of a suite of lexica all automatically compiled by Said Mohammad and coworkers,[9] under the NRC general label [15,16], which also includes the NRC Hashtag Sentiment Lexicon, built by exploiting the sentiment of hashtags, and two more "double" resources built on the previous ones; namely, the NRC Hashtag Affirmative Context Sentiment Lexicon (and NRC Hashtag Negated Context Sentiment Lexicon), and the Sentiment140 Affirmative Context Lexicon (and Sentiment140 Negated Context Lexicon), where the average sentiment of a term occurring in an affirmative context is separate from the average sentiment of the same term occurring in a negated context.

Within the NRC tools there is also a small manually (crowdsourced) produced lexicon of about 1500 entries, the MaxDiff Twitter Sentiment Lexicon [17]. Handmade resources are always smaller than automatically built ones, as they are costly to create, but they are obviously more reliable. The ways of reducing noise in the automatic acquisition of resources have been explored by, for instance, the use of sentiment proxies such as emoticons [18] or, usually, the adoption of a semiautomatic approach, bootstrapping manually annotated sets of seeds [19,20]. This is also the approach behind the original creation of SentiWordNet. Finally, using seeds, Chen and Skien [21] automatically produced sentiment lexica for more than 130 languages, reporting an overlap with existing resources of more than 90%.

For the same economical reasons we have just mentioned, and because tools for basic text processing are more advanced for English than for other languages, the creation of lexica for other languages can benefit from the automatic porting of data and techniques from already existing English resources [22]. An example is the transfer of SentiWordNet to other languages. Because English Wordnet is easily aligned to WordNets in other languages, porting SentiWordNet becomes rather straightforward, and it has been shown that sense transfer is robust across languages in about 90% of cases [23] (see also [24] for a WordNet-based, almost unsupervised approach for generating polarity lexica in multiple languages). Transfer of resources can be done in two ways: to create a stable, reusable resource, or "on the fly." An example of the former is Sentix for Italian (via SentiWordNet), which is currently the most commonly used resource for this language in sentiment analysis systems [25,26]. As for the latter, resources are translated and used in the context of a sentiment analysis system, but are not necessarily stabilized and distributed for further use; for example, Hernández Farías et al. [27] use a large number of English lexica/resources automatically translated into Italian to exploit information as features in their supervised system, but no specific evaluation or refining of such resources is performed.

2.2 BEYOND THE POLARITY VALENCE: EMOTION LEXICA, ONTOLOGIES, AND PSYCHOLINGUISTIC RESOURCES

Recently, a variety of affective lexica have been proposed to offer information about affect expressed in text according to finer levels of granularity (eg, referring not simply to positive or negative

[8] https://github.com/abromberg/sentiment_analysis/blob/master/AFINN/AFINN-111.txt
[9] http://saifmohammad.com/WebPages/lexicons.html

sentiment polarity but to emotional categories such as joy, sadness, and fear). Moreover, a variety of psycholinguistic resources are available that can give some additional measure about the emotional disclosure in social media texts, according to different theoretical perspectives on emotions. All such affect-related resources could be useful with the purpose to increase the coverage of different aspects of affect in textual content. We organize the description of such resources into three groups: the first group is related to information about emotions by referring to a finer-grained categorization model; the second group includes *psycholinguistic* resources and other resources that refer to different perspectives on affect, according to dimensional approaches to emotion modeling; the third group includes *knowledge-based* resources and ontologies, which have been developed with the twofold aim to help sentiment analysis systems to grasp the conceptual and affective information associated with natural language opinions, and to use Semantic Web and linked data technologies to provide structured, reusable, and meaningful sentiment analysis results.

2.2.1 Emotion lexica: finer-grained affective lexica based on emotional categories

Theories in the nature of emotion suggest the existence of basic or fundamental emotions such as anger, fear, joy, sadness, and disgust. Different approaches propose different basic or fundamental sets, each having its own specific eliciting conditions and its own specific physiological, expressive, and behavioral reaction patterns. Accordingly, available resources refer to different models of emotions well grounded in psychology, such as the ones proposed by Plutchik [28] and Ekman [29].

One of the first resources referring to a finer-grained model of affect is WordNet-Affect,[10] which was developed through the selection and labeling of WordNet synsets representing affective concepts [30]. A number of WordNet synsets are assigned to one or more affective labels (called a-labels). In particular, the affective concepts representing emotional states are individuated by synsets marked with the a-label emotion. There are also other a-labels for those concepts representing moods, traits, situations eliciting emotions, and so on. The newer version, WordNet-Affect 1.1, includes more than 900 synsets and proposes also a taxonomy of emotions, where the hierarchical structure is modeled on the WordNet hyperonym relation. Starting from WordNet-Affect, for task 14 at SemEval 2007, a new version of the resource has been provided [31]. This resource includes only a portion of WordNet-Affect, as it was reannotated at a finer-grained level with use of the six emotional category labels from [29]: joy, fear, anger, sadness, disgust, and surprise. The resource EmoLex[11] has been developed as part of the NRC suite of lexica. It is a word-emotion association lexicon [32] built via crowdsourcing; the annotations were manually done through Amazon's Mechanical Turk. It contains 14,182 words labeled according to Plutchik's eight primary emotions [28]—joy, sadness, anger, fear, trust, surprise, disgust, and anticipation—and also annotations for negative and positive sentiments. Emolex was originally annotated at a word-sense level. Then the word-level lexicon was created by the taking of the union of emotions associated with all the senses of a word. Very recently, NRC Emolex has been provided also in more than 20 languages as a result of translation of the English terms with use of Google Translate. The underlying assumption is that, despite some cultural differences, affective norms are stable across languages.

[10] http://wndomains.fbk.eu/wnaffect.html
[11] http://www.saifmohammad.com/WebPages/lexicons.html

Another sense-level affective lexicon is SentiSense,[12] which attaches emotional meanings to concepts from the WordNet lexical database [33]. It is composed by a list of 2190 synsets tagged with emotional labels from a set of 14 emotional categories, which refer to a merger of models by Arnold [34], Plutchik, and Parrot [35]: joy, fear, surprise, anger, disgust, love, anticipation, hope, despair, sadness, calmness, like, hate, and ambiguous. Such emotional categories are also related via an antonym relationship. SentiSense has been developed semiautomatically with use of several semantic relations between synsets in WordNet.

2.2.2 Psycholinguistic resources and other accounts of affect

Psycholinguistic resources can be also helpful for capturing the emotional content of the text, allowing the mapping of words onto psychologically valid dimensions. One of the most well-known psycholinguistic resources is the Linguistic Inquiry and Word Count (LIWC)[13] dictionary [36], which assigns one or more psychological categories, such as positive emotion and negative emotion, to individual words. The 2007 version of LIWC comprises almost 4500 words and word stems distributed in categories for analysis of psycholinguistic features in texts. For example, the word "happy" would be labeled with the categories positive emotion and affect. The most recent evolution, LIWC2015, is composed of almost 6400 words, word stems, and selected emoticons, distributed in a wider set of categories arranged hierarchically. The dictionary has been extended and new categories have been added. The new additions allow the user to better tackle with social media language.

Moreover, some psychological theories propose that the nature of an emotional state is determined by its position in a space of independent dimensions. According to such an approach, emotions can be defined as a coincidence of values on a number of different strategic dimensions. Therefore they are described not by marking a small set of discrete categories but rather by scoring properties such as valence (positive/negative) and arousal (active/passive) in a continuous range of values. Accordingly, some lexical resources have been built by reference to models with a dimensional view on affect-related phenomena.

For instance, the Dictionary of Affect in Language[14] developed by Whissell [37] includes 8742 English words rated on a three-point scale along three dimensions: activation (degree of response that humans have under an emotional state); imagery (how difficult it is to form a mental picture of a given word); pleasantness (degree of pleasure produced by words). Instead, the Affective Norms for English Words[15] [38] is a database developed by the rating of 1034 English words in terms of the Osgood, Suci, and Tannenbaum dimensional theory of emotions along three dimensions of the valence-arousal-dominance model [39]: valence or pleasure (the polarity of the emotional activation ranging from positive to negative); arousal or intensity (the degree of excitement or activation an individual feels toward a given stimulus, ranging from calm to exciting); dominance or control (the degree of control an individual feels over a specific stimulus, ranging from out of control to in control).

[12] http://nlp.uned.es/~jcalbornoz/SentiSense.html
[13] http://www.liwc.net
[14] ftp://perceptmx.com/wdalman.pdf
[15] http://csea.phhp.ufl.edu/media/anewmessage.html

2.2.3 Concept-based resources and ontologies: toward a knowledge-based approach to sentiment analysis

Many researchers are currently devoting efforts to develop ontologies of emotions, in some cases also referring to the Semantic Web initiative [40–43]. The use of structured knowledge via ontologies or semantic networks in the sentiment analysis tasks opens new opportunities for understanding opinions expressed in texts. In particular, such knowledge bases can include semantic information both on the sentiment domain and the emotion domain, or on the concepts and the context related to words used for expressing opinions, enabling interesting possibilities of reasoning on such knowledge (see Section 3.4). Again, these ontologies typically try to mirror models of emotions well established in psychology. However, as a difference from emotion lexica such as Emolex, where emotions are taken as a flat list of concepts, here a more sophisticated knowledge can be encoded. For instance, in ontologies the taxonomic relationships among the emotions can be represented, or other interesting relations such as intensity, similarity, or oppositions. The representation of such relationships between concepts in the ontology enables also interesting possibility to automatically reason on such knowledge about emotions. As we will see, Plutchik's circumplex model [28], which inspired the development of the Emolex lexicon, has also been exploited as a reference of concept-level resources and ontology of emotions with important differences that we will highlight.

Concept-level resources use ontologies or semantic networks to enable semantic text analysis. One of the most used resources in this category is SenticNet[16] [40], which aims to create a collection of commonly used polarity concepts (ie, commonsense concepts with relatively strong positive or negative polarity). Differently from other resources such as SentiWordNet, which also contains null-polarity terms, SenticNet does not contain concepts with neutral or almost neutral polarity. The current version includes 30,000 natural language concepts collected from the Open Mind corpus, and the resource is distributed in Resource Description Framework (RDF) XML format. Each concept is associated with emotion categorization values expressed in terms of the hourglass of emotions model [44], which organizes and blends 24 emotional categories from Plutchik's model into four affective dimensions (pleasantness, attention, sensitivity, and aptitude). Moreover, in SenticNet a polarity value that lies in the interval $[-1, 1]$ is associated with each concept. This value is calculated in terms of the four dimensions of the hourglass of emotions model, and specifies if (and to what extent) the input concept is positive or negative (eg, concepts such as "make a good impression" will have a polarity value close to 1, whereas concepts such as "being fired," "leave behind," or "lose control" will have a polarity value close to -1. In this case the emotion categorization model supports comparison and aggregation among results of an emotional analysis of concepts into polarity values.

EmoSenticNet[17] is another concept-based lexical resource [45] and was automatically built by the merging of WordNet-Affect and SenticNet, with the main aim to have a more complete resource containing not only quantitative polarity scores associated with each SenticNet concept but also qualitative affective labels. In particular, it assigns WordNet-Affect emotion labels related to the six Ekman basic emotions mentioned before (disgust, sadness, anger, joy, fear, and surprise) to SenticNet concepts. The whole list currently includes 13189 annotated entries.

[16]http://sentic.net
[17]http://www.gelbukh.com/emosenticnet/

In [46], an ontology of emotional categories based Plutchik's circumplex model [28] was developed (ArsEmotica Ontology of Emotions), encoded in OWL 2 QL. The ontology has been defined in the context of the ArsEmotica project (see Section 3.4) but it is so generic that it might be used to analyze emotions in running text in any domain. It encodes the emotional categories of Plutchik's model and links them with semantic relationships, which can be represented as a wheel of emotions. In particular, it accounts for various levels of intensity, the similarity and opposite relationships, and compositions of basic emotions (primary dyads). Overall, it distinguishes the 32 emotional concepts of the wheel of emotions. The emotional concepts have been semiautomatically linked to synsets from WordNet-Affect 1.1, and then to Italian lexical entries, by exploitation of MultiWordNet, an Italian WordNet strictly aligned with Princeton WordNet 1.6. The linkage between the language level (lexicon based) and the conceptual level representing the emotional concepts has been formalized by integration of the ontology framework Lexicon Model for Ontologies [47].

Finally, we mention a set of affect-related ontologies in the Semantic Web field that are not lexical resources but were created as a response to a more foundational need to represent all the main features of an emotion, and to standardize the knowledge of emotions, so as to support very broad semantic interoperability among affective computing applications, by allowing the mapping of concepts and properties belonging to different emotion representation models.

The Human Emotion Ontology [41] has been developed in this direction, as has the semantic vocabulary Onyx,[18] which has been proposed to describe emotion analysis processes and results [43].

2.3 SOCIAL MEDIA CORPORA ANNOTATED FOR SENTIMENT AND FINE EMOTION CATEGORIES

The growing interest in the development of automatic systems for sentiment analysis has also prompted the production of annotated corpora that could be exploited for the development of such systems in a learning fashion (see also Section 3).

The design of schemes for the annotation of corpora is always a task in the field of data classification, which leads to theoretical assumptions about the concepts to be annotated. It defines what kind of information must be annotated, the inventory of markers to be used, and the annotation's granularity. As highlighted in [48], in annotated corpora for sentiment analysis this is especially challenging. Research in psychology outlines three main approaches to the modeling of emotions and sentiments: the categorical, the dimensional, and the appraisal-based approach. The most widespread are the categorical and the dimensional ones, which describe emotions by marking a small set of discrete categories and scoring properties in a continuous range of values. Accordingly, the kinds of knowledge usually annotated are the sentiment's polarity (positive vs. negative), category (happiness vs. sadness), the source and target of the sentiment, and its intensity. Annotations can be based on simple broad polarity labels, possibly equipped with intensity ratings, which also helps us to classify texts where mixed sentiments are expressed. They can also be based on labels representing different emotions.

The resources developed within the SemEval evaluation campaigns (regularly since 2013) comprise tweet collections annotated for subjectivity, polarity, emotions, irony, and aspects/properties, and are

[18] https://www.gsi.dit.upm.es/ontologies/onyx/

not restricted to English [25,26,48,49]. In addition to the SemEval related data, other collections have been constructed, especially on news and events or product reviews, which are obviously a very good source of opinionated texts, and a field of interest for any business-related application that works with users' opinions. For example, a widely used dataset is the Movie Review Data,[19] a collection of movie reviews [50]. Sentiment labels (positive vs. negative) are assigned to the global review, while in an extension of this corpus, annotation is done at the sentence level [51]. Another classic resource is the MPQA corpus,[20] based on [52], which contains news articles from a wide variety of sources and which was manually annotated for opinions and other private states. This corpus has been expanded very recently with finer-grained information, and now includes entity and event target annotations [53], in line with the increasing interest in aspect-based sentiment analysis. A pioneer work in this sense is [54]: given a product, Hu and Liu [54] mine and summarize reviews related to specific features of that product, on which customers have expressed their positive or negative opinion, so that sentiment is associated with aspects of rather than the whole product. Extraction of opinions on aspects (specifically service, location, and rooms of hotels) is also the focus in [55], where the authors collected and processed 10,000 TripAdvisor reviews.

Interest in properties and aspect-based sentiment analysis is increasing also in Twitter-based work, as shown by recently organized evaluation campaigns on this [3]. Indeed, even if a lot shorter in characters than a product review, the same tweet can contain both positive and negative information, possibly relating to different entities that are mentioned. Niek Sanders's Twitter dataset[21] is not aspect annotated but a specific topic is assigned to each tweet, and this topic can be possibly used as a proxy for the entity with which the sentiment is associated, as is also done in [25], where specific hashtags are used as proxies.

Still in term of proxies, it is interesting to note that to save annotation effort in the creation of sentiment datasets, there have been experiments with *distant learning*, where class labels are not assigned manually but are rather derived from other information available. Performing sentiment analysis on tweets, Go et al. [56] trained a few classifiers using emoticons as noisy labels, and achieved an accuracy of about 80%. Recently, a new Twitter corpus was released by Mohammad et al. [57], where a multilayer scheme is applied. It contains a set of tweets with annotations concerning different aspects: sentiment (positive or negative), but also finer-grained annotation of emotions (anger, anticipation, disgust, fear, joy, sadness, surprise, and trust), and in addition purpose (to point out a mistake, to support, to ridicule, etc.) and style (simple statement, sarcasm, hyperbole, understatement). The corpus is not generic but focused on the political domain. Mohammad et al. collected tweets labeled with a set of hashtags pertaining to the 2012 US presidential election. The tweets were annotated manually by reliance on crowdsourcing platforms. This is the first dataset including a fine-grained annotation concerning stylistic features related to irony, and also a multilayer annotation concerning affect, since both annotations on sentiment polarity and on emotions are provided. For style, only 23% of the tweets were labeled with a style tag pertinent to the expression of irony, whereas most of them were annotated with the label simple statement, which can be interpreted as a tag for marking nonironic expressions.

[19] http://www.cs.cornell.edu/people/pabo/movie-review-data/
[20] http://mpqa.cs.pitt.edu/corpora/mpqa_corpus/
[21] http://www.sananalytics.com/lab/twitter-sentiment/

Previous corpora annotated for fine-grained emotion categories, such as the one proposed for SemEval 2007 task14 [31], were not focused on social media data. The availability of such a new kind of datasets opens the way to the possibility of building machine learning systems that predict emotion and purpose labels in unseen tweets, as proposed by Mohammad et al. [57], who also presented the first results and a baseline for such new sentiment-related tasks on Twitter.

3 USING SEMANTICS IN SENTIMENT ANALYSIS

Because of the task's nature, semantics is obviously a crucial ingredient of any sentiment analysis system. However, depending on the system's complexity, and depending also on the specific task that is undertaken (see Section 1), semantic information of different types can be accessed and incorporated in various different ways. In Section 3.1 we discuss how lexical information mostly derived by the resources discussed in Section 2 is used, while in Section 3.2 we briefly review how current semantic similarity models are adapted to sentiment processing. In Section 3.3 we explore the deeper semantic processing that is required by finer-grained tasks such as aspect-based sentiment analysis. Finally, in Section 3.4 we review systems that perform sentiment analysis by *reasoning* over semantic resources.

3.1 LEXICAL INFORMATION

Independently of the approach (rule based or statistical), virtually all systems for sentiment analysis rely on information derived from lexical resources. In machine learning such information is used as features, normally in combination with other features, too. The best performing system for subtask B of SemEval 2015's task 10, which is the most standard message-level polarity task, is an ensemble classifier that builds on three well-performing classifiers at SemEval 2013 and one at SemEval 2014 [58]. Although the classifiers were handpicked with attention not only to performance but also to significant differences in feature combinations, all four classifiers employ polarity dictionaries of some sort and in some way. Some systems rely on one resource only; for example, Günther and Furrer [59] only use SentiWordNet, while other systems will try to use and combine information from all available sources [15,60].

As a general strategy, lexical information is collected at the word level (eg, the categorical or gradual polarity of a given token) for each word of the whole tweet, and is then propagated via different combinations at the sentence/message level. Such combinations yield values that are then used as different features. For example, feature values can be the number of tokens in the text with a positive score or the number of those with a negative score, but also derived values such as the global polarity score of a tweet obtained as the average of the tokens' polarities. Also, maximum and minimum scores, when available, are used as features.

Most statistical approaches are support vector machine models that implement such features and similar features, and Taboada et al. [10] highlight that information from basic unigrams appears to be the most useful. Although this shows that words in themselves are very informative for this task, the intrinsic limitations in the use of information from a plain dictionary lookup are quite evident if one thinks that words are used in a specific (syntactic) context, and their polarity can change substantially according to how they relate to other words in the text. Mainly because of the lack of tools that can deal with Twitter, few systems perform word sense disambiguation, so the richer sense-annotated lexica cannot be exploited to their full potential. A way of addressing polysemy without disambiguation was suggested by Basile and Nissim [25], who introduced the concept of *polypathy* (calculated as the standard deviation of the polarity scores of the possible senses of a lemma) as an indicator of variance of polarity scores across a lemma's synsets, which can then be used as a feature or threshold value (tokens that have a too high polypathy could, for example, be ignored, as they provide conflicting information).

Another crucial limitation of not taking context into account is *negation* [61] (recall also the separate lexica for negated and affirmative contexts reviewed earlier). To cope with this aspect, Taboada et al. [10] incorporate information from what they call *contextual valence shifters*, showing an increase in performance. This is a first step toward deeper language processing, which is considered more and more necessary, even for short texts such as tweets.

3.2 DISTRIBUTIONAL SEMANTICS

In the context of sentiment analysis the idea of exploiting the distributional hypothesis—namely, the assumption that words that occur in the same contexts tend to have similar meanings [62,63]—simply boils down to the fact that similarity models that predict, for example, that "amazing" and "wonderful" are similar could be extended to predict that if "amazing" has a positive value, so too will "wonderful," and will be at the opposite end of spectrum to, say, "terrible" and "awful." However, general similarity models usually take into account the lexical and morphosyntactic context of a word but not necessarily the text's polarity. Therefore similarity is potentially accurate at a syntactic and more general semantic level but not necessarily in a sentiment-aware way, so "good" and "bad" end up being very similar. This is true for classic distributional similarity models as well as for the more recent and successful distributed vector representations known as *word embeddings*, at least in their standard formulation [64,65].

As a first attempt at directly incorporating sentiment in learning the distributional context of a word, so that words expressing similar sentiment do indeed end up having similar vector representations, Maas et al. [66] developed a model where they give a sentiment predictor function to each word. On the positive/negative classification of tweets, their system is shown to perform better than models that incorporate embeddings trained in a nonsentiment-aware fashion. An even more powerful and recent model, which outperforms that of Maas et al. [66], was proposed by Tang et al. [67], who train a neural network by associating each *n*-gram with the polarity of a sentence (thus beyond the word level, as in [66]), and show that sentiment-specific word embeddings effectively distinguish words with opposite sentiment polarity. This model performs better than other models that use generally learned embeddings.

3.3 ENTITIES, PROPERTIES, AND RELATIONS

Interest in finer-grained sentiment analysis has also necessarily prompted the need for finer-grained semantic analysis, and therefore deeper language processing. For example, aspect-based sentiment analysis must rely on the identification of specific entities and/or properties of entities in reviews or tweets. To do this, standard techniques for entity detection and classification are employed, such as sequential taggers, possibly retrained for specific domains. Particular attention to (named) entities in sentiment analysis is also shown by the OpeNER EU-funded project,[22] which focuses on named entity recognition within sentiment analysis.

[22] http://www.opener-project.eu/

Further, relations between the entities and events involved must be identified, so as to know what is said of which entity. An obvious way to do this is to exploit dependency relations, although deeper processing of tweets is not so simple because of the idiosyncratic and often ungrammatical language such short texts contain (although recent work based on learning neural knowledge graph embeddings shows an error reduction of more than 26% in semantic parsing of tweets [68]). Similar issues arise when one is developing systems to detect stance, as in order to assess the opinion of someone toward a given target, all relations between the entities involved must be correctly identified and associated with the sentiment expressed. However, a deeper linguistic analysis of text is also beneficial, if not necessary, for standard message- or text-level sentiment analysis, as it helps to treat the issue of *contextual valence shifters* mentioned in Section 3.1 by also accounting for word order and sentence structure. To this end, the Natural Language Processing Group at Stanford University developed a *sentiment treebank*.[23] This treebank has been used to train a recursive neural network built on top of grammatical structures [69], achieving an increase of 5 percentage points on sentence polarity classification. On fine-grained sentiment level they obtained a 9.7% improvement over a bag-of-words baseline, and overall showed the ability to accurately capture the effects of negation and its scope at various levels in the tree structures.

3.4 CONCEPT-LEVEL SENTIMENT ANALYSIS: REASONING WITH SEMANTICS

Approaches enabling concept- and context-based analysis can lead to a better understanding of opinions expressed in textual data, therefore reducing the gap between unstructured information and structured machine-processable data. Concept-level sentiment analysis exploits large semantic knowledge bases (eg, ontologies and semantic networks), together with natural language processing tools and techniques, thus stepping away from blind use of keywords and word co-occurrence counts. Rather, it relies on the implicit features associated with natural language concepts. Unlike purely syntactic techniques, concept-based approaches are able to detect also sentiments that are expressed in a subtle manner (eg, through the analysis of concepts that do not explicitly convey any emotion) but that are implicitly linked to other concepts that do so.

The accent on concept-level sentiment analysis is recent. We can observe different knowledge representations of the concepts related to words and of their affective load, as well as different ways in which systems that embrace this perspective on sentiment analysis exploit and reason about such knowledge, depending also on the specific task they address.

Also, this perspective on sentiment analysis offers a challenge to the Semantic Web community. A periodic shared task on concept-level sentiment analysis has been launched recently [5], and some of the systems described below explicitly rely on the use of Semantic Web and linked data resources and tools to enable automated reasoning (eg, ontological reasoning on the taxonomic structure of an ontology of affect) and semantic metadata processing.

Let us start with Sentilo,[24] an unsupervised, domain-independent system that relies on a knowledge-based approach where existing tools for natural language processing and publicly available lexical/affective resources are combined with Semantic Web techniques for representing and reasoning about knowledge involved in opinion sentences. Sentilo exploits the semantic graph representation

[23] http://nlp.stanford.edu/sentiment/treebank.html
[24] http://wit.istc.cnr.it/stlab-tools/sentilo/

of a sentence enriched with opinion-related information (eg, opinion holder, topics, sentiment scores). Such a semantic representation allows Sentilo to address the finer-grained tasks of detecting holders and topics of an opinion, and additionally to identify and distinguish both main topics and subtopics. Specifically, given an input opinion sentence (natural language text), the system returns the corresponding semantic representation as a FRED graph (RDF/OWL graph) annotated with concepts from a semantic model of opinions. Annotations concerning the opinion holder, opinion features, and opinion topics are produced on the basis of a set of heuristic rules exploiting the FRED semantic representation and by use of lexical and affective resources, such as SentiWordNet and SenticNet. An evaluation of Sentilo as the opinion holder and main/subtopic detection tool has been presented [70]. The system is able to achieve good performance compared with the other tools, highlighting promising results on opinion holder (F1: 95%), subtopic (F1: 78%), and topic (F1: 68%) detection. Sentilo is also available as a REST service that returns RDF as output. This is important to make the output of the system machine processable and reusable.

In other systems the accent is on the use of conceptual and affective knowledge of words, the focus being on sophisticated representations of the affective space that can enable commonsense or ontological reasoning on networks or taxonomies of concepts. This is, for instance, the case for many systems exploiting the resources of the SenticNet suite of resources [71]. Another system that can be placed in this category is ArsEmotica,[25] which is an example of a domain-dependent system focused on cultural heritage. It applies sentiment analysis to resources from online art collections, and exploits tags given by the visitors on social platforms as an additional source of information [46]. The extraction of the emotional semantics from social tags is driven by an ontology-based representation of the emotional categories presented in the previous section. Methods and tools from the Semantic Web and natural language processing provide the building blocks for the creation of a semantic social space to organize artworks according to the ontology of emotions. The output is encoded into World Wide Web Consortium ontology languages. The semantic representation is not limited to affect but concerns also aspects of the specific art domain. This gives the twofold advantage of enabling tractable reasoning on the relationships between emotions and other interesting dimensions of the domain (eg, artists and artworks), and fostering the interoperability and integration of tools developed in the Semantic Web and linked data community. The system has been evaluated against a dataset of tagged multimedia artworks from the ArsMeteo Italian art portal.[26] A SPARQL end point has been implemented to explore the collection and to extract information about the relationships among emotions, artworks and authors by posing queries such as: "*Give me the artworks stirring emotions similar to sadness and belonging to the music genre*." In some systems the focus on conceptual semantics, not limited to affect, is also combined with issues related to the contextual semantics of words. Such approaches do not offer fixed sentiment polarities but assign context-specific sentiment orientation to words. An interesting proposal in this line is SentiCircles, a lexicon-based approach for the detection of sentiment in Twitter posts at both entity level and tweet level [72]. SentiCircle builds a dynamic representation of context to tune the preassigned strength and polarity of words in the lexicon, and incorporates both contextual semantics (ie, semantics inferred from the co-occurrence of words) and conceptual semantics (ie, semantics extracted from background ontologies).

[25] http://di.unito.it/arsemotica
[26] http://www.arsmeteo.org

4 CONCLUSIONS

We started this contribution by pointing out how much interest and research in sentiment analysis have grown in the past few years. We have highlighted that in the recent SemEval campaigns tasks more complex than basic polarity classification have been proposed. And complexity here means finer-grained interpretations of sentiment, be it related to figurative uses such as irony or be it related to singling out which entities or entity properties are actually in focus in an opinionated statement. This is also reflected in resource creation, where we have seen that corpora have been extended in their annotation to include polarity at the sentence level, where before this was done at the global document level, or at the entity level, where before it was done at the message level.

Accordingly, sentiment analysis systems are required to be ever more accurate and more sophisticated at the same time, thus attempting to perform deeper semantic processing that is rather successful on language data cleaner than that from Twitter and for other natural language processing tasks but is still experimental with regard to noisy social media data, with all the subjectivity that sentiment analysis carries. We have provided a survey of the state-of-the-art approaches and resources in this sense, highlighting how semantics is being coded and used in this emerging and growing field. From what we see, sentiment analysis in social media appears to be expanding in a variety of directions, incorporating and adapting processing tools from natural language processing and reasoning tools from the Semantic Web community, as well as including possibly orthogonal aspects such as a whole range of affect-related concepts. What we also see is that semantics necessarily remains the core of this, in the use, coding, and processing of the meaning of words and larger phrases, so we can only hope that this survey serves as a starting point for much more to come.

Our expectations do not concern only the development of sentiment analysis per se, since related tasks, such as personality and irony detection [73–75], or author profiling [76], also have interest in exploiting the lexical knowledge of affect encoded in the resources we have described.

REFERENCES

[1] S. Rosenthal, P. Nakov, S. Kiritchenko, S. Mohammad, A. Ritter, V. Stoyanov, SemEval-2015 task 10: sentiment analysis in Twitter, in: Proceedings of the Ninth International Workshop on Semantic Evaluation (SemEval 2015), Association for Computational Linguistics, Denver, CO, 2015, pp. 451–463.

[2] A. Ghosh, G. Li, T. Veale, P. Rosso, E. Shutova, J. Barnden, A. Reyes, SemEval-2015 task 11: sentiment analysis of figurative language in Twitter, in: Proceedings of the Ninth International Workshop on Semantic Evaluation (SemEval 2015), Association for Computational Linguistics, Denver, CO, 2015, pp. 470–478.

[3] M. Pontiki, D. Galanis, H. Papageorgiou, S. Manandhar, I. Androutsopoulos, SemEval-2015 task 12: aspect based sentiment analysis, in: Proceedings of the Ninth International Workshop on Semantic Evaluation (SemEval 2015), Association for Computational Linguistics, Denver, CO, 2015, pp. 486–495.

[4] I. Russo, T. Caselli, C. Strapparava, SemEval-2015 task 9: CLIPEval implicit polarity of events, in: Proceedings of the Ninth International Workshop on Semantic Evaluation (SemEval 2015), Association for Computational Linguistics, Denver, CO, 2015, pp. 443–450.

[5] D.R. Recupero, E. Cambria, ESWC'14 challenge on concept-level sentiment analysis, in: Semantic Web Evaluation Challenge—SemWebEval 2014 at ESWC 2014, Anissaras, Crete, Greece, May 25–29, 2014, Revised Selected Papers, Communications in Computer and Information Science, vol. 475, Springer, 2014, pp. 3–20.

[6] T. Wilson, J. Wiebe, P. Hoffmann, 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, Vancouver, Canada, 2005, pp. 347–354.
[7] J. Wiebe, R. Mihalcea, Word sense and subjectivity, in: Proceedings of the 21st International Conference on Computational Linguistics and the 44th Annual Meeting of the Association for Computational Linguistics, 2006, pp. 1065–1072.
[8] S. Baccianella, A. Esuli, F. Sebastiani, SentiWordNet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining, in: Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10), European Language Resources Association (ELRA), Valletta, Malta, 2010, pp. 2200–2204.
[9] C. Fellbaum, WordNet: An Electronic Lexical Database, Bradford Books, Cambridge, MA, 1998.
[10] M. Taboada, J. Brooke, M. Tofiloski, K. Voll, M. Stede, Lexicon-based methods for sentiment analysis, Comput. Linguist. 37 (2) (2011) 267–307.
[11] P.J. Stone, D.C. Dunphy, M.S. Smith, D.M. Ogilvie, The general inquirer: a computer approach to content analysis, Am. J. Sociol. 73 (5) (1966) 634–635.
[12] J. J. McAuley, J. Leskovec, Hidden factors and hidden topics: understanding rating dimensions with review text, in: Proceedings of the 7th ACM Conference on Recommender Systems, 2013, pp. 165–172.
[13] S. Kiritchenko, X. Zhu, C. Cherry, S. Mohammad, NRC-Canada-2014: detecting aspects and sentiment in customer reviews, in: Proceedings of the Eighth International Workshop on Semantic Evaluation (SemEval 2014), Association for Computational Linguistics and Dublin City University, Dublin, Ireland, 2014, pp. 437–442.
[14] T. Loughran, B. McDonald, When is a liability not a liability? Textual analysis, dictionaries, and 10-ks, J. Financ. 66 (1) (2011) 35–65.
[15] S. Mohammad, S. Kiritchenko, X. Zhu, 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, GA, 2013, pp. 321–327.
[16] X. Zhu, S. Kiritchenko, S. Mohammad, NRC-Canada-2014: recent improvements in the sentiment analysis of tweets, in: Proceedings of the Eighth International Workshop on Semantic Evaluation (SemEval 2014), Association for Computational Linguistics and Dublin City University, Dublin, Ireland, 2014, pp. 443–447.
[17] S. Kiritchenko, X. Zhu, S.M. Mohammad, Sentiment analysis of short informal texts, J. Artif. Intell. Res. 50 (2014) 723–762.
[18] A. Pak, P. Paroubek, Twitter for sentiment analysis: when language resources are not available, in: Proceedings of the 23rd International Workshop on Database and Expert Systems Applications, 2011, pp. 111–115.
[19] G. Pitel, G. Grefenstette, Semi-automatic building method for a multidimensional affect dictionary for a new language, in: Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08), European Language Resources Association (ELRA), Marrakech, Morocco, 2008.
[20] T.D. Smedt, W. Daelemans, “Vreselijk mooi!” (terribly beautiful): a subjectivity lexicon for Dutch adjectives, in: N. Calzolari, K. Choukri, T. Declerck, M.U. Dogan, B. Maegaard, J. Mariani, J. Odijk, S. Piperidis (Eds.), Proceedings of LREC 2012, European Language Resources Association (ELRA), 2012, pp. 3568–3572.
[21] Y. Chen, S. Skiena, Building sentiment lexicons for all major languages, in: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Association for Computational Linguistics, Baltimore, Maryland, 2014, pp. 383–389.
[22] C. Banea, R. Mihalcea, J. Wiebe, Porting multilingual subjectivity resources across languages, Trans. Affect. Comput. 4 (2) (2013) 211–225.

[23] C. Banea, R. Mihalcea, J. Wiebe, Sense-level subjectivity in a multilingual setting, Comput. Speech Lang. 28 (1) (2014) 7–19.
[24] I. San Vicente, R. Agerri, G. Rigau, Simple, robust and (almost) unsupervised generation of polarity lexicons for multiple languages, in: Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics (EACL2014), Gothenburg, Sweden, 2014, pp. 88–97.
[25] V. Basile, M. Nissim, Sentiment analysis on Italian tweets, in: Proceedings of the Fourth Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, ACL, Atlanta, GA, 2013, pp. 100–107.
[26] V. Basile, A. Bolioli, M. Nissim, V. Patti, P. Rosso, Overview of the Evalita 2014 SENTIment POLarity classification task, in: Proceedings of the Fourth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2014), Pisa, Italy, 2014.
[27] D.I. Hernández Farías, D. Buscaldi, B. Priego-Sánchez, IRADABE: adapting English lexicons to the Italian sentiment polarity classification task, in: First Italian Conference on Computational Linguistics (CLiC-it 2014) and the Fourth International Workshop EVALITA 2014, 2014, pp. 75–81.
[28] R. Plutchik, The nature of emotions, Am. Sci. 89 (4) (2001) 344–350.
[29] P. Ekman, An argument for basic emotions, Cogn. Emot. 6 (3–4) (1992) 169–200.
[30] C. Strapparava, A. Valitutti, WordNet-Affect: an affective extension of WordNet, in: Proceedings of the International Conference on Language Resources and Evaluation (LREC'04), vol. 4, European Language Resources Association (ELRA), 2004, pp. 1083–1086.
[31] C. Strapparava, R. Mihalcea, SemEval-2007 task 14: affective text, in: Proceedings of the Fourth International Workshop on Semantic Evaluations, SemEval '07, Association for Computational Linguistics, Stroudsburg, PA, 2007, pp. 70–74.
[32] S.M. Mohammad, P.D. Turney, Crowdsourcing a word-emotion association lexicon, Comput. Intell. 29 (3) (2013) 436–465.
[33] J. Carrillo de Albornoz, L. Plaza, P. Gervás, SentiSense: an easily scalable concept-based affective lexicon for sentiment analysis, in: Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12), European Language Resources Association (ELRA), Istanbul, Turkey, 2012, pp. 3562–3567.
[34] M.B. Arnold, Emotion and Personality, vol. 1, Columbia University Press, New York, NY, 1960.
[35] W.G. Parrot, Emotions in Social Psychology: Essential Readings, Psychology Press, Philadelphia, 2001.
[36] J.W. Pennebaker, M.E. Francis, R.J. Booth, Linguistic Inquiry and Word Count: LIWC 2001, Lawrence Erlbaum Associates, Mahwah, NJ, 2001.
[37] C. Whissell, Using the revised dictionary of affect in language to quantify the emotional undertones of samples of natural languages, Psychol. Rep. 2 (105) (2009) 509–521.
[38] M.M. Bradley, P.J. Lang, Affective norms for English words (ANEW): instruction manual and affective ratings, tech. rep., Center for Research in Psychophysiology, University of Florida, Gainesville, FL, 1999.
[39] J.A. Russell, A. Mehrabian, Evidence for a three-factor theory of emotions, J. Res. Pers. 11 (3) (1977) 273–294.
[40] E. Cambria, D. Olsher, D. Rajagopal, SenticNet 3: a common and common-sense knowledge base for cognition-driven sentiment analysis, in: Proceedings of AAAI Conference on Artificial Intelligence, 2014, pp. 1515–1521.
[41] M. Grassi, Developing HEO human emotions ontology, in: Proceedings of the 2009 Joint COST 2101 and 2102 International Conference on Biometric ID Management and Multimodal Communication, Springer-Verlag, Berlin, Heidelberg, 2009, pp. 244–251.
[42] V. Francisco, P. Gervas, F. Peinado, Ontological reasoning for improving the treatment of emotions in text, Knowl. Inform. Syst. 25 (2010) 421–443.
[43] J.F. Sánchez-Rada, C.A. Iglesias, Onyx: a linked data approach to emotion representation, Inform. Process. Manage. 52 (1) (2016) 99–114.

[44] E. Cambria, A. Livingstone, A. Hussain, The hourglass of emotions, in: COST 2102 Training School, Revised Selected Papers, Lecture Notes in Computer Science, vol. 7403, Springer, 2011, pp. 144–157.
[45] S. Poria, A. Gelbukh, A. Hussain, N. Howard, D. Das, S. Bandyopadhyay, Enhanced SenticNet with affective labels for concept-based opinion mining, IEEE Intell. Syst. 28 (2) (2013) 31–38.
[46] F. Bertola, V. Patti, Ontology-based affective models to organize artworks in the social semantic web, Inform. Process. Manage. 52 (1) (2016) 139–162.
[47] J. McCrae, D. Spohr, P. Cimiano, Linking lexical resources and ontologies on the semantic web with LEMON, in: The Semantic Web: Research and Applications, Springer, New York, NY, 2011, pp. 245–259.
[48] C. Bosco, V. Patti, A. Bolioli, Developing corpora for sentiment analysis: the case of Irony and Senti-TUT, IEEE Intell. Syst. 28 (2) (2013) 55–63.
[49] J. Steinberger, T. Brychcın, M. Konkol, Aspect-level sentiment analysis in czech, in: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL 2014), 2014, pp. 24.
[50] B. Pang, L. Lee, A sentimental education: sentiment analysis using subjectivity, in: Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL 2004), 2004, pp. 271–278.
[51] B. Pang, L. Lee, Seeing stars: exploiting class relationships for sentiment categorization with respect to rating scales, in: Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL 2005), 2005, pp. 115–124.
[52] J. Wiebe, T. Wilson, C. Cardie, Annotating expressions of opinions and emotions in language, Lang. Resour. Eval. 39 (2–3) (2005) 165–210.
[53] L. Deng, J. Wiebe, MPQA 3.0: an entity/event-level sentiment corpus, in: R. Mihalcea, J.Y. Chai, A. Sarkar (Eds.), Proceedings of HLT-NAACL 2015, The Association for Computational Linguistics, 2015, pp. 1323–1328.
[54] M. Hu, B. Liu, Mining and summarizing customer reviews, in: Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD ’04, ACM, New York, NY, 2004, pp. 168–177.
[55] I. Titov, R.T. McDonald, A joint model of text and aspect ratings for sentiment summarization, in: Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics (ACL 2008), vol. 8, 2008, pp. 308–316.
[56] A. Go, R. Bhayani, L. Huang, Twitter sentiment classification using distant supervision, CS224N Project Report, Stanford 1 (2009) 12.
[57] S.M. Mohammad, X. Zhu, S. Kiritchenko, J. Martin, Sentiment, emotion, purpose, and style in electoral tweets, Inform. Process. Manage. 51 (4) (2015) 480–499.
[58] M. Hagen, M. Potthast, M. Büchner, B. Stein, Webis: an ensemble for Twitter sentiment detection, in: Proceedings of the Ninth International Workshop on Semantic Evaluation (SemEval 2015), Association for Computational Linguistics, Denver, CO, 2015, pp. 582–589.
[59] T. Günther, L. Furrer, GU-MLT-LT: sentiment analysis of short messages using linguistic features and stochastic gradient descent, in: Proceedings of the Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013), Association for Computational Linguistics, Atlanta, GA, 2013, pp. 328–332.
[60] Y. Miura, S. Sakaki, K. Hattori, T. Ohkuma, Teamx: a sentiment analyzer with enhanced lexicon mapping and weighting scheme for unbalanced data, in: Proceedings of the Eighth International Workshop on Semantic Evaluation (SemEval 2014), 2014, pp. 628–632.
[61] R. Saurí, A Factuality Profiler for Eventualities in Text, ProQuest, Ann Arbor, MI, 2008.
[62] Z. Harris, Distributional structure, Word 10 (23) (1954) 146–162.
[63] M. Sahlgren, The distributional hypothesis, Ital. J. Linguist. 20 (1) (2008) 33–54.
[64] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, in: Proceedings of the ICLR Workshop, 2013, pp. 1–12.

[65] T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, in: Advances in Neural Information Processing Systems (NIPS), 2013, pp. 3111–3119.
[66] A.L. Maas, R.E. Daly, P.T. Pham, D. Huang, A.Y. Ng, C. Potts, Learning word vectors for sentiment analysis, in: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies—Volume 1, HLT '11, Association for Computational Linguistics, Stroudsburg, PA, 2011, pp. 142–150.
[67] D. Tang, F. Wei, N. Yang, M. Zhou, T. Liu, B. Qin, Learning sentiment-specific word embedding for Twitter sentiment classification, in: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, vol. 1, 2014, pp. 1555–1565.
[68] L. Heck, H. Huang, Deep learning of knowledge graph embeddings for semantic parsing of Twitter dialogs, in: Proceedings of the Second IEEE Global Conference on Signal and Information Processing (DRAFT), IEEE Institute of Electrical and Electronics Engineers, 2014, pp. 597–601.
[69] R. Socher, A. Perelygin, J. Wu, J. Chuang, C.D. Manning, A. Ng, C. Potts, Recursive deep models for semantic compositionality over a sentiment treebank, in: Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Seattle, WA, 2013, pp. 1631–1642.
[70] A. Gangemi, V. Presutti, D. Reforgiato Recupero, Frame-based detection of opinion holders and topics: a model and a tool, IEEE Comp. Int. Mag. 9 (1) (2014) 20–30.
[71] E. Cambria, B.W. Schuller, B. Liu, H. Wang, C. Havasi, Knowledge-based approaches to concept-level sentiment analysis, IEEE Intell. Syst. 28 (2) (2013) 12–14.
[72] H. Saif, Y. He, M. Fernandez, H. Alani, Contextual semantics for sentiment analysis of Twitter, Inform. Process. Manage. 52 (1) (2016) 5–19.
[73] S. Poria, A.F. Gelbukh, B. Agarwal, E. Cambria, N. Howard, Common sense knowledge based personality recognition from text, in: Advances in Soft Computing and Its Applications, Proceedings of the 12th Mexican International Conference on Artificial Intelligence, MICAI 2013, Part II, Springer, 2013, pp. 484–496.
[74] D.I. Hernández Farías, J. Benedí, P. Rosso, Applying basic features from sentiment analysis for automatic irony detection, in: Pattern Recognition and Image Analysis, Proceedings of IbPRIA 2015, Lecture Notes in Computer Science, vol. 9117, Springer, 2015, pp. 337–344.
[75] E. Sulis, D.I.H. Farías, P. Rosso, V. Patti, G. Ruffo, Figurative messages and affect in Twitter: Differences between #irony, #sarcasm and #not, Knowledge-Based Systems, Elsevier, 2016, In Press. http://dx.doi.org/10.1016/j.knosys.2016.05.035
[76] F.M. Rangel, P. Rosso, On the impact of emotions on author profiling, Inform. Process. Manage. 52 (1) (2016) 73–92.

CHAPTER

LINKED DATA MODELS FOR SENTIMENT AND EMOTION ANALYSIS IN SOCIAL NETWORKS

4

C.A. Iglesias[a], J.F. Sánchez-Rada[a], G. Vulcu[b], P. Buitelaar[b]

Technical University of Madrid, Madrid, Spain[a] National University of Ireland, Galway, Ireland[b]

1 INTRODUCTION

Sentiment analysis is now an established field of research and a growing industry [1]. However, language resources for sentiment analysis are being developed by individual companies or research organizations and are normally not shared, with the exception of a few publicly available resources such as WordNet-Affect [2] and SentiWordNet [3].

Domain-specific resources for multiple languages are potentially valuable but not shared, sometimes because of intellectual property and license considerations, but often because of technical reasons, including interoperability. Several initiatives have addressed interoperability of language resources since the late 1980s, such as Text Encoding Initiative [4], but there is not yet a widely accepted global solution for integrating and combining heterogeneous linguistic resources from different sources [5]. In this respect the data interoperability problem has been addressed by linked data technologies, which have gained wide acceptance. *Linked data* [6] refers to best practices and technologies for publishing, sharing, and connecting structured data on the web. This approach has been followed by the linking open data (LOD) project, a grassroots community effort supported by the World Wide Web Consortium (W3C) whose aim is to bootstrap the Web of Data by identifying existing datasets available under open licenses, converting them to Resource Description Framework (RDF) format following the linked data principles, and publishing them on the web. The data cloud that originated from this initiative is known as the *LOD cloud*. Several communities such as the Open Linguistics Working Group [5] proposed the idea of adopting linked data principles for representing, sharing, and publishing open linguistic resources with the aim of developing a subcloud of the LOD cloud of linguistic resources, known as the *linguistic linked open data (LLOD) cloud* [7].

In addition, the use of linked data for modeling linguistic resources provides a clear path to their semantic annotation and linking with semantic resources of the Web of Data. This is especially important for making sense of social media streams, whose semantic interpretation is particularly challenging because they are strongly interconnected, temporal, noisy, short, and full of slang [8]. Moreover, several authors [9] have shown that the use of semantics in sentiment analysis outperforms

Sentiment Analysis in Social Networks. http://dx.doi.org/10.1016/B978-0-12-804412-4.00004-8