This is the final peer-reviewed accepted manuscript of:

Reforgiato Recupero, D., Presutti, V., Consoli, S., Gangemi, A., & Nuzzolese, A. G. (2015). Sentilo: frame-based sentiment analysis. Cognitive Computation, 7(2), 211-225.

The final published version is available online at:

<http://dx.doi.org/10.1007/s12559-014-9302-z>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Sentilo: Frame-based Sentiment Analysis

Diego Reforgiato Recupero ·
Valentina Presutti ·
Sergio Consoli · Aldo Gangemi ·
Andrea Giovanni Nuzzolese

Received: date / Accepted: date

Abstract Sentilo is an unsupervised, domain-independent system that performs sentiment analysis by hybridizing natural language processing techniques and semantic web technologies. Given a sentence expressing an opinion, Sentilo recognizes its holder, detects the topics and sub-topics that it targets, links them to relevant situations and events referred to by it, and evaluates the sentiment expressed on each topic/subtopic. Sentilo relies on a novel lexical resource which enables a proper propagation of sentiment scores from topics to subtopics, and on a formal model expressing the semantics of opinion sen-

D. Reforgiato Recupero
National Research Council (CNR), Institute of Cognitive Sciences and Technologies, Semantic Technology Laboratory, Via Gaifami 18, 95028 Catania, Italy
E-mail: diego.reforgiato@istc.cnr.it

V. Presutti
National Research Council (CNR), Institute of Cognitive Sciences and Technologies, Semantic Technology Laboratory, Via Nomentana, 56 - 00161 Rome, Italy
E-mail: valentina.presutti@cnr.it

S. Consoli
National Research Council (CNR), Institute of Cognitive Sciences and Technologies, Semantic Technology Laboratory, Via Gaifami 18, 95028 Catania, Italy
E-mail: sergio.consoli@istc.cnr.it

A. Gangemi
LIPN, University Paris 13, Sorbone Cité, UMR CNRS, Paris, France
National Research Council (CNR), Institute of Cognitive Sciences and Technologies, Semantic Technology Laboratory, Via Nomentana, 56 - 00161 Rome, Italy
E-mail: aldo.gangemi@cnr.it

A.G. Nuzzolese
Department of Computer Science and Engineering, University of Bologna, Mura Anteo Zamboni, 7 - 40126 Bologna, Italy
National Research Council (CNR), Institute of Cognitive Sciences and Technologies, Semantic Technology Laboratory, Via Nomentana, 56 - 00161 Rome, Italy
E-mail: andrea.nuzzolese@istc.cnr.it

tences. Sentilo provides its output as a RDF graph, and whenever possible it resolves holders' and topics' identity on Linked Data.

Keywords Opinion mining · Sentic computing · Sentiment analysis · Conceptual frames

1 Introduction

The success of Web social media and review-based crowd-sourcing sites characterizes the Web as a huge focus group containing people's beliefs, judgments, speeches, attitudes that can be of enormous value to gather financial predictions and, for companies and organizations, to market their products and results, identify new opportunities, as well as manage their reputations. Therefore the study of intelligent algorithms capable of automatically mining opinions from natural language content is more and more attracting the interest of academia and industry. This is the goal of *Sentiment Analysis* (SA) [28], which has a substantial overlap with *opinion mining*, a rather recently developing research field whose aim is to detect and extract subjective information, such as opinions or emotions, in source materials.

An opinion can be defined as *an intentional statement by somebody (holder) on some fact (topic) that is expressed with a possible sentiment*. A SA system should be able to extract and characterize opinions by recognizing the *attitude* (positive, negative or objective) of an opinion holder on a certain topic, or by evaluating the overall *tonality* of a document; it can be *document-based* or *sentence-based*.

A number of research efforts and investments are in place in this domain, such as the EU FP7 EuroSentiment project¹, which aims at providing a shared set of language resources for fostering sentiment analysis.

The first and most common approaches to SA come from traditional natural language processing (NLP), examples are [37, 39]. They hardly can cope with some aspects of opinions such as subtle linguistic forms, expression of positive and negative nuances at the same time, implicit judgements that can be derived from explicit ones. These issues call for a cognitive and social perspective of the problem to be connected to the NLP one. In other words, the idea is to shift from a word-level to a concept-level analysis of opinions. This intuition has been the basis of a novel, multi-disciplinary approach to SA, called *sentic computing*, which claims the importance of including semantic features into opinion mining: claim supported by evidence that SA algorithms performance improves if they are augmented with semantic features [29, 31, 36, 21, 13]. [10] provides a nice survey on, and discusses new trends in, opinion mining and SA.

The main difference between a traditional NLP approach and a sentic computing approach to SA is that the former mainly relies on parts of text in which opinions are explicitly expressed, such as polarized terms, affective words, and

¹ EuroSentiment EU FP7 project. <http://eurosentiment.eu/>, 2014

their co-occurrence frequencies. On the contrary, the latter performs a fine-grained analysis of opinion sentences so as to identify and analyze all relevant concepts and their mutual relations that can either explicitly or implicitly convey the expression of emotions. Let us analyze the following example:

“People hope that the President will resign.”

A human would easily understand that the people referred to by this sentence have a rather negative opinion on “the President” because they envision his/her resignation. This simple sentence however lacks of terms explicitly indicating a positive or negative opinion, e.g. about “the President”, making it hard for a NLP-based tool (e.g. [31]) to catch it. However, the term “hope” evokes a positive attitude towards what is referred to by the subordinate proposition “the President will resign”. This means that “people” refers to the holder of a positive opinion about a possible “resign” event (i.e., main topic) whose agent is “the President” (i.e. a subtopic). Intuitively, a subtopic is an entity that is indirectly targeted by an opinion sentence. For example, in this case the opinion holder indirectly expresses an opinion on “the President”, while it directly expresses an opinion on a “resign” event. Being a resignation a generally negative event for its agent, a positive judgement of it implies a negative one on its agent. The above rationale can be performed by sentic computing approaches, making them more powerful than NLP-based approaches in determining the subjective information conveyed by opinion sentences.

In [18] we have introduced *Sentilo*, a sentic computing approach to opinion mining. *Sentilo* produces a formal representation (i.e. a RDF graph) of an opinion sentence that allows to distinguish its holders and topics (whose identities are resolved on Linked Data) with very high accuracy (holder detection 95%, main topic detection 68%, subtopic detection 78%)².

One of the goals of our research is to develop a method for fine-grained SA of sentences, meaning that given an opinion sentence, we want to assess a sentiment score for each identified topic as well as for the overall sentence. In this article, we describe an upgraded version of this approach and of its implementation that addresses this goal, which also includes the development of a novel lexical resource. *Sentilo* is able to perform a deep analysis of opinion sentences like the one exemplified above.

This capability has a significant potential impact on industrial applications that use sentiment analysis, e.g. in TripAdvisor, where users indicate their rating to e.g. hotels, and write some comment about them. Comments may contain details about the rationales behind the expressed rating, focusing on specific aspects of a hotel. Let us consider an example of a user giving an average rating to a hotel with the following comment: *The hotel rooms were good but children entertainment was insufficient..*

Most sentiment analysis tools would be able to recognize that the comment has a neutral polarity overall. However, it would be desirable to automatically characterize and detail what specific aspects of the hotel were judged and

² F1 measures.

how: it would enable the automatic classification of subjects under review, e.g. hotels, based on their specific aspects' ratings, hence providing a much richer and effective service to TripAdvisor users. Sentilo supports this task. Referring to the above example, Sentilo would recognize a positive opinion about the hotel rooms and assess a negative opinion for its "children entertainment" service.

The contribution of this article can be summarized as follows:

- An extended version of *OntoSentilo*, an ontology for opinion sentences introduced in [18];
- *SentiloNet*: a new lexical resource enabling the evaluation of opinions expressed by means of events and situations;
- *Sentiment Propagation*: a novel scoring algorithm for opinion sentences;
- an upgraded version of Sentilo's prototype implementation³;
- an empirical evaluation of Sentilo on a corpus of user-based hotel reviews.

Additionally, Sentilo prototype has been endowed with two different user-oriented graphical interfaces, one showing the RDF graph enriched with opinion-related concepts, and the other hiding such details but providing the SA evaluation for each relevant topic.

The paper is organized as follows: Section 2 discusses related work. Section 3 introduces our sentic computing approach, named *Sentilo*. Section 4 describes an extension to *OntoSentilo*, and the novel resource *SentiloNet*. Section 5 presents a novel algorithm for computing sentiment scores of individual topics as well as the overall tone of a sentence. Details about the system implementation are given in Section 6, where we also show the graphical user interfaces built on top of *Sentilo*. Section 7 presents and discusses evaluation results. Finally, conclusions and discussions about future directions are reported in Section 8.

2 Related Work

SA approaches can be grouped into three main categories: *keyword spotting*, in which text is classified into categories based on the presence of fairly unambiguous words [15, 46]; *lexical affinity*, which assign arbitrary words a probabilistic affinity for a particular concept [40, 47]; and *statistical methods*, which calculate the valence of keywords, punctuation and word co-occurrence frequencies on the base of large training corpora [19, 20]. Most opinion mining and sentiment analysis systems in the literature are centered on the extraction of the most relevant text fragments containing subjective opinions through machine learning approaches [1, 24], fuzzy logic models [23], as well as feature selection [32]. Other works exploit polarity classification of opinions (typically positive, negative, neutral) in a target document [34, 41, 45], while others deal with the extraction of moods from informal text resources such as blogs [30, 42]. For a detailed survey on opinion mining and SA the reader can refer to [33, 28].

³ <http://wit.istc.cnr.it/stlab-tools/sentilo/>

The problem with these approaches is that they mainly rely on parts of text in which opinions are explicitly expressed through positive and negative terms. (e.g. good, nice, excellent, fortunate, correct, superior, best, bad, nasty, poor, unfortunate, wrong, inferior, worst, etc.). In many cases, opinions are expressed implicitly through context and domain dependent concepts, making the performance of NLP-based approaches limited. This has been the main motivation behind the idea of sentic computing.

Sentic computing is a novel multi-disciplinary approach to SA first proposed in [8], which envisions the development of biologically-inspired, psychologically-motivated computational approaches and claims the importance of performing semantic-based analysis and fine grained classification of opinions for tackling the problem [10]. In sentic computing, the analysis of natural language is based on *affective* ontologies [12,7] and *sensitiveness* reasoning tools [9,11]. For additional details on sentic computing the reader can refer to [8].

In line with the sentic computing direction, solving the tasks of detecting opinion topics have proved to positively impact the performance of SA algorithms [27,6,44,21]. However, none of such algorithms provides a proper evaluation of topic detection as a task per se.

Topic detection and sentiment attribution to specific topics can be obtained by looking at the *compositionality* of a text. The importance of deep parsing of a text in order to attribute sentiment to specific topics is shown by both [38] and [18], which show clear improvements on previous non-compositional work. The system described in the first one takes sentence compositionality into account through a deep parser that generates a *sentiment treebank* for a specific domain (e.g. movie reviews), while the second relies on open domain semantic parsing and semantic web machine reading, as described in this paper as well.

To the best of our knowledge only [18] combines *holder* and topic detection by providing a formal representation of opinion sentences that includes a clear identification of opinion topics and subtopics.

2.1 Lexical and semantic resources for SA

SA therefore relies on the use of lexical and semantic resources as background knowledge. We list and briefly describe the ones that are currently used by Sentilo's approach presented as contribution of this paper.

WordNet [16] is a large lexical database of English. Nouns, verbs, adjectives and adverbs are grouped into sets of synonyms (synsets), each expressing a distinct concept. Synsets are interlinked by means of conceptual-semantic and lexical relations. It is available online along with its APIs, which are compatible with different programming languages.

SenticNet [12] is a publicly available semantic and affective resource for concept-level opinion and sentiment analysis. SenticNet has been built by following a sentic computing procedure: it exploits an ensemble of graph-mining and dimensionality-reduction techniques to bridge the conceptual and affective gap between word-level natural language data and concept-level opinions and sentiments conveyed by them. Sentic API provides the semantics and sentics associated with more than 14,000 common sense concepts. We use SenticNet to retrieve the polarity scores of terms expressing sentiment.

SentiWordNet [2] is a lexical resource for opinion mining. SentiWordNet assigns to each synset of WordNet three sentiment scores: positivity, negativity, objectivity. Values range from -1 to +1 depending on whether the underlying synset represents a very negative or very positive opinion. We use SentiWordNet to retrieve the polarity scores of terms expressing sentiment.

VerbNet (VN) [5] is the largest on-line verb lexicon currently available for English. It is a hierarchical domain-independent, broad-coverage verb lexicon with mappings to other lexical resources such as WordNet and FrameNet⁴. VerbNet is organized into (frame-like) verb classes extending Levin's [26] classes. Each verb class includes a set of "frames" each characterized by a syntactic form and a set of semantic predicates each associated with its thematic roles. We use VerbNet for producing a frame-based semantic representation of opinion sentences.

DBpedia [25] is the RDF version of structured information extracted from Wikipedia. We use DBpedia for unambiguously resolving, when possible, the identity of opinion holders and topics.

3 Background: Sentilo, a sentic computing approach

Sentilo is a sentic computing approach to SA. It provides a formal representation of opinion sentences, in the form of RDF graphs, according to an ontology defining the main concepts and relations characterizing opinion sentences. In an earlier version [18] it focused on detecting holders and topics in opinion sentences and proved to address these tasks with high accuracy (holder detection 95%, main topic detection 68%, subtopic detection 78%)⁵. Let us recall that we distinguish main topics, i.e. direct targets of an opinion, from subtopics, i.e. indirect targets of an opinion. Subtopics always have a dependency relation on a main topic.

In this article, we present an upgraded version of Sentilo, which extends its features in a twofold way: (i) an enriched formal representation of opinion sentences, (ii) a SA algorithm able to compute topic-level as well as sentence-level sentiment scores.

⁴ The framenet project. <http://framenet.icsi.berkeley.edu>, 2002

⁵ F1 measures.

Fig. 1 Pipeline of Sentilo.

Figure 1 shows the components of the pipeline of Sentilo. Some of them (depicted as purple boxes) are novel and extend the earlier version (depicted as green boxes) described in [18]. For an accurate and detailed description of existing components the reader can refer to [18]. However, for the sake of completeness we briefly describe here the core ones.

Sentilo approach to SA is based on the neo-Davidsonian assumption that events and situations are to be considered first class entities in the description of the world. As far as SA is concerned, this means that events and situations provide contextual information for evaluating sentiment expressions in opinion sentences, by reinforcing or weakening them, as well as by making them emerge when they are implicit. In order to computationally reproduce such assumption, Sentilo uses *Boxer* [4], a tool that transforms natural language text to a logical form according to Discourse Representation Theory (DRT) [22]. DRT is a formal theory of meaning based on an event-based model for representing natural language. Boxer relies on VerbNet [5] for identifying and formalizing events and their associated thematic roles. However, the core component of Sentilo is *FRED*⁶ [35], which transforms such logical form to RDF by complying to Semantic Web and Linked Data design principles [3], and by extending the representation model with event- and situation- semantics as formally defined by DOLCE+DnS⁷ ontology [17].

Let us consider the following sentence:

John thinks that the summer will become wonderful.

Figure 2 depicts the RDF graph produced by *FRED*⁸ for this sentence. It provides a formal representation of the sentence, which includes the event occurrence `fred:think_1` of type `fred:Think` and the event occurrence `fred:become_1` of type `fred:Become`, (both typed as `dul:Event`). The meaning

⁶ FRED, <http://wit.istc.cnr.it/stlab-tools/fred>, December 2014

⁷ Dolce Ultra Lite Ontology. <http://ontologydesignpatterns.org/ont/dul/DUL.owl>

⁸ prefix `dul:` stands for <http://www.ontologydesignpatterns.org/ont/dul/DUL.owl> and prefix `rdfl:` stands for <http://www.w3.org/1999/02/22-rdf-syntax-ns#>; prefix `fred:` refers to a locally defined namespace that can be customized by users.

Fig. 2 Semantic representation of the sentence *John thinks that the summer will become wonderful*.

of the two concepts is disambiguated through an alignment to VerbNet concepts⁹. The entities having the thematic roles of agents for these events are recognized as “John”, and “summer”, respectively. The quality “wonderful” is correctly detected and associated with “summer”. Notably, FRED also represents modalities, in this case the mode “will” is interpreted as “necessary”. Currently, this information is not exploited by Sentilo, although it could be in future developments.

The *Opinion model annotator* component (depicted in Figure 1) enriches this formal representation with concepts defined in *OntoSentilo*, an ontology for opinion sentences, depicted in Figure 3. For a detailed description of this ontology the reader can refer to [18]. Briefly, *OntoSentilo* defines concepts and relations that characterize the entities composing an opinion sentence. The opinion triggering event, identified by the `hasOpinionTrigger` relation, is typically referred to by a verb, which explicitly indicates the presence of an opinion in the sentence, e.g. to think, to like, to hope, etc. Such verbs can be neutral, e.g., to think, or carry a sentiment themselves e.g., to like. Opinion holders can be either explicit or implicit in a sentence, and they are represented by the relation `hasHolder`. Main topics and subtopics of an opinion are represented by the relations `hasTopic` and `hasSubTopic`, respectively. Finally opinion features are represented as values of the relation `hasOpinionFeature`.

In the upgraded version of Sentilo, presented in this paper, we have extended the ontology for opinion sentences by further elaborating the dependencies between main topics and subtopics (i.e. `dependsOn` relation in Figure 3a). The extension is described in Section 4.

Figure 4 depicts a fragment of the RDF graph representing the sentence “John thinks that the summer will become beautiful” and enriched with concepts from *OntoSentilo*. All concepts and relations added by Sentilo belongs

⁹ prefix `vn.data:` refers to VerbNet [5]

Fig. 3 OntoSentilo: an ontology for opinion sentences.

Fig. 4 Semantic representation of the sentence *John thinks that the summer will become wonderful* enriched with opinion-related concepts.

to its local namespace referred to by the prefix `sentilo:`. Sentilo identifies and formally represents the holder of the opinion, i.e. “John”, the main topic of the sentence, i.e. the event occurrence `fred:become_1`, and its subtopic, i.e. the event occurrence `fred:summer_1`. Opinion features are identified as values of the relation `dul:hasQuality`, in this case `fred:Wonderful` is a quality of the event occurrence `fred:summer_1`.

4 Extending the ontology for opinion sentences

As discussed in sections 1 and 3, our approach distinguishes main topics from subtopics of an opinion sentence. This distinction is of primary importance as far as our SA approach is concerned because our aim is to assess a sentiment score value for each identified topic in an opinion sentence as well as the tone of the overall sentence.

As Sentilo represents sentences with an event/situation-based approach, i.e. frame-based, it can benefit from the semantics of the frame-based relations between topics and subtopics for correctly and more deeply analyzing the sentiment of opinion sentences. In this section we focus on this aspect.

Let us consider the example of Section 1:

People hope that the President will be condemned by the judges.

Fig. 5 RDF graph for the sentence *People hope that the President will be condemned by the judges*.

In this case the triggering event *hope* carries a positive sentiment tone of the expressed opinion. However, this is not enough for a complete SA of this sentence. So far we can only state that the sentence expresses a positive opinion on the event *condemn*, which according to OntoSentilo is a main topic.

If the main topic was a *simple* entity such as a person, a place, an object, etc., typically there would not be subtopics depending on it. This is the case of sentences such as: “*I like Rome*”, “*I support your decision*”, etc. If instead, the main topic is an event or a situation such as in this case, it will have a frame-like structure, and other entities belonging to its structure could be indirect targets of the expressed opinion, i.e. sub-topics. We need to further investigate the frame structure of the main topic in order to first identify all potential subtopics¹⁰ and then select those ones that are actually affected by the expressed sentiments.

Figure 5 shows the RDF graph for the example sentence. It can be noticed that the event *condemn* involves two entities having two different roles. Respectively: *the judges* (represented by **fred:judge_1**) having the role **vn.role:Agent**, and *the President* (represented by **fred:President**) having the role **vn.role:Theme**. A human would easily understand that while the agentive role, i.e. *the judges*, is not indirectly affected by the opinion expressed by means of *hope*, the passive role, i.e. *the President*, actually is. In other words, if people hope his/her condemnation it means they have a negative opinion on him/her. While their opinion on the judges is not expressed, neither implicitly, by this sentence. Notice, that while *hope* carries a positive sentiment polarity, such polarity changes its sign when applied on its passive role.

Hence, there are two issues to be solved in order to correctly interpret this sentence:

- how to distinguish subtopics that are indirectly affected by an opinion?
- how to evaluate the polarity of the sentiment indirectly expressed on them?

In order to address these issues, we have extended Sentilo approach in a twofold way: (i) we have extended OntoSentilo by further elaborating the

¹⁰ Notice that this process can be recursive, and the role of main topic/subtopic in such cases would be contextual to the current iteration.

relation **dependsOn** between main topics and subtopics (see Figure 3a); (ii) we have created a novel resource of annotated verbs, named SentiloNet¹¹ (see Figure 1), which enables the selection of subtopics that are actually indirectly affected by opinions expressed in a sentence, as well as the evaluation of their polarity.

4.1 Topic-subtopic semantic relationships

From a formal representation perspective, we have extended OntoSentilo with the following relations:

- **sentilo:participatesIn**: is a relation connecting all (potential) subtopics that belong to the frame structure of a **dul:Situation** (through the relation **boxing:involves**) or of a **dul:Event** (by playing a thematic role), given that the situation/event is a main topic in the sentence. This relation simply allows to navigate the graph from subtopics to main topics. It makes it explicit the dependence between them by expressing the intensional semantics of “participating in an event or situation”;
- **sentilo:playsSensitiveRole**: is a relation connecting a subtopic to a main topic, indicating that such subtopic would be indirectly affected by a possible opinion directly expressed on the main topic;
- **sentilo:isPositivelyAffectedBy**: is a relation connecting a subtopic to an event/situation (being it a main topic) indicating that a possible opinion on the main topic would indirectly affect the subtopic by keeping the same polarity;
- **sentilo:isNegativelyAffectedBy** is a relation connecting a subtopic to an event/situation (being it a main topic) indicating that a possible opinion on the main topic would indirectly affect the subtopic by changing its polarity;

4.2 SentiloNet, a new resource to identify opinions for events

In order to automatically enrich the formal representation of an opinion sentence with the new defined relations, we need to provide our method with the proper background knowledge. To this aim we introduce the concepts of *Role sensitivity* and *Factual impact*. These concepts have been the basis for the design of a novel resource of annotated verbs, named *SentiloNet*.

Role sensitivity : a role is sensitive with respect to an event (referred to by a verb) if it is indirectly affected by an opinion directly expressed on the event. As far as the annotation of a verb (frame) is concerned, the *sensitivity* is an attribute of its thematic roles. The value of the sensitivity attribute of a role with respect to a verb can be either *true* or *false*, meaning that the role is sensitive or is not, respectively.

¹¹ An excerpt of SentiloNet can be downloaded from <http://www.stlab.istc.cnr.it/documents/sentilo/sentilonet.zip>

Factual impact : indicates that an event (referred to by a verb) has either a positive or negative impact on a specific role. As far as the annotation of a verb is concerned, the *factual impact* is an attribute of its sensitive roles. The value of this attribute for a role can be *positive*, *negative*, meaning that the inherited opinion will keep its polarity or change it, respectively.

The current version of SentiloNet includes 1,100 annotated verbs. Given the high number of different thematic roles of verbs, we have devised a heuristics that allowed us to manually annotate a good amount of verbs in a rather limited amount of time. Our reference is the set of thematic roles used by FRED. They come from VerbNet and from an internal resource of verbs used by Boxer (see Figure 1). We have created two classes of roles, i.e. AGNT and PTNT, each including a set of roles that are considered equivalent with respect to their agentive or passive attitude, respectively. The extension of the two classes is the following:

- AGNT = {boxer:Agent, vn.role:Experiencer, vn.role:Actor, vn.role:Actor1, vn.role:Actor2, vn.role:Cause, vn.role:Agent}
- PTNT = {boxer:Patient, vn.role:Topic, vn.role:Beneficiary, vn.role:Patient, vn.role:Patient1, vn.role:Patient2}

The roles `boxer:Theme` and `vn.role:Theme` are treated in a special way as they can have an agentive or passive semantics, depending on context. Operationally, we assign them to *AGNT* if, in a given sentence, the identified event has no defined agentive role. We assign it to *PTNT* otherwise.

SentiloNet indicates, for 1,100 verbs, the value of *sensitivity* and *factual impact* attributes for each class of roles.

Figure 6 shows (a fragment of) the RDF graph representing the sample sentence “*People hope that the President will be condemned by the judges.*” We omit the part of the graph, which includes the annotations of opinion holder, opinion trigger, topics and subtopics, for the sake of readability. However, for completeness, we highlight with red rounded-rectangles the graph nodes that are originally annotated with such roles. The reader can easily inspect the whole resulting graph by running Sentilo demo with the sample sentence¹².

In Figure 6 the graph showed in Figure 5 is enriched by Sentilo, using SentiloNet and (extended) OntoSentilo as background knowledge. The resulting graph correctly shows that the *judges* participate in the *condemn* event, however they do not play a sensitive role in it. In fact, according to the rationale of the “sensitivity” concept, if an opinion is expressed on a condemn event occurrence, the entity playing an agentive role in it is not affected by this opinion. In the specific example, the opinion holder expresses a positive opinion on this event occurrence, which does not always imply any implicit opinion on its agent, i.e. the judges. Differently, the role played by “the President” is sensitive and negatively affected by such event. In fact, expressing a positive sentiment on the condemnation of someone cognitively implies having a negative opinion

¹² <http://wit.istc.cnr.it/stlab-tools/sentilo/service>

Fig. 6 RDF graph for the sentence *People hope that the President will be condemned by the judges.* enriched with topic-subtopic dependency relations.

on him/her. Sentilo captures very well these relations thanks to the SentiloNet resource, which provides a cognitive-oriented background knowledge as far as the concepts of *role sensitivity* and *factual impact* are concerned.

5 Sentiment Propagation: a sentic computing algorithm for sentiment analysis

The upgraded version of Sentilo includes an algorithm for assessing a sentiment score at topic as well as sentence level, named *Sentiment propagation*. As we will explain later in this section, the computation of these scores depends on the structure of the RDF graph representing the opinion sentence as it takes into account the frame-based semantic relations between main topics and subtopics.

5.1 Gathering individual sentiment score

A first step towards the design of this algorithm is to assign an individual sentiment score (if applicable) to each element in an opinion sentence graph. To this aim we rely on SentiWordNet [2] and SenticNet [12], as shown in Figure 1. We assign a score to adjectives and adverbs that are identified by `dul:hasQuality` relation values, and to instances of `dul:Event` that are recognized as trigger events, i.e., identified by `sentilo:hasOpinionTrigger` relation values.

We have investigated and implemented two alternative approaches for score selection¹³. The first approach assigns a score retrieved by querying the *polarity* attribute of a concept in SenticNet. The second one combines SenticNet and SentiWordNet scores.

We are currently performing a set of experiments and annotation efforts in order to rigorously compare the performances of the two methods. However,

¹³ Users can choose between the two by means of a selection box included in the graphical user interface of Sentilo prototype available at <http://wit.istc.cnr.it/sentilo-release/sentilo>.

from an initial set of empirical observations on using Sentilo with the two approaches we noticed that the method that combines the scores from the two resources shows to be more reliable. This confirmed our expectations based on the following rationale: sometimes, SenticNet misses a score value for a required concept. Moreover, it provides one score per concept without distinguishing its possible different nuances. Hence, SenticNet score approximates an average value for the scores of all possible senses, or possibly indicates the most probable one. For example, let us consider the SenticNet score value (polarity attribute) for the concept *stink*, which is +0.029. Clearly, this value does not capture appropriately the negative nuance of this concept. SentiWordNet instead, being word-based, provides different score values for the different senses of *stink*, including the negative value -0.38. However, disambiguating the sense of a word is a time-consuming task, and only relying on the most-frequent sense may result in retrieving a wrong value. For this reason, combining the SentiWordNet scores of most frequent senses, and the SenticNet score can provide an appropriate balanced value.

We have devised a simple heuristics for computing this combined score. Currently, this is the default method applied by Sentilo for individual score assignment. The algorithm that implements the combined score computation is described in Figure 7. Let us use a running example for explaining the computation of a combined score, by simulating the execution of the algorithm described in Figure 7 for the word “amazing”, which would return the value 0.304.

The input of the algorithm is a term w (adjective, adverb or verb), e.g., “amazing”, and the output is $score \in [-1, +1]$, e.g. 0.304. First, we retrieve the SenticNet polarity value of w and store it in the variable $sNet$. For example, if $w = \text{“amazing”}$, $sNet = 0.357$. This can be seen by using the SenticNet API at <http://sentic.net/api/en/concept/amazing/>. Then, we query WordNet [16], for retrieving the most frequently used senses of w and store them in a set T . The frequency of a sense is given by the value of the attribute *tag_count*. After retrieving the first most frequent sense, we want to include in our set only those senses that have a frequency value high enough to be significative. For this reason, we include in T only those senses s having a tag_count_s value not smaller than $1/10$ of tag_count_{s-1} .

For example, the word “amazing” has the senses and respectively, the frequencies shown in Table 1¹⁴. In this case, we store both senses in T :

$$T = \{amazing_1, amazing_2\}.$$

Then, we query SentiWordNet for retrieving the score values associated with each sense in T (see Table 1), compute their average value, and store it in a variable sWN . In our example, $sWN = 0.25$.

Finally, we compute the average value of $sNet$ and sWN , which will give us the value of the combined score. For the word “amazing”, the combined score

¹⁴ We also include in the table the respective sentiment scores

Word	WordNet Sense ID	tag_count_s	Sentiment score
<i>amazing</i> ₁	302359789	5	0.125
<i>amazing</i> ₂	301282510	5	0.375

Table 1 Senses, frequencies, and sentiment scores of the word “amazing”, according to SentiWordNet.

```

score = CombinedScore(w){
  sNet = SenticNet score for w;
  n = number of w senses;
  T = {};
  for i ← 1 to n do
    si = extract next sense of w from WordNet in decreasing order of
    tag.count;
    if  $tag\_count[s_{i-1}] > 10 \times tag\_count[s_i]$  then
      break;
    T = T ∪ {si};
  end
  T' = SentiWordNet score values for each element in T;
  sWN = AVERAGE(T' {...});
  return AVERAGE(sWN, sNet);
}
```

Fig. 7 Sentilo default algorithm for opinion expressing word scoring.

is:

$$CombinedScore("amazing") = (sNet + sWN)/2 = (0.357 + 0.25)/2 = 0.357$$

This value becomes the object value of a relation `sentilo:hasScore` between the RDF entity referred to by *w* in the opinion sentences, and its score value.

5.2 Sentiment propagation algorithm

Given an entity, identified as a topic of an opinion (either a main or subtopic), we compute its sentiment score by combining the scores of all its associated opinion features (i.e. values of `dul:hasQuality` relations), which are extracted from the RDF graph representing the opinionated sentence. If a topic participates in an event or a situation occurrence, we propagate their sentiment scores to it, according to the semantics expressed by the frame-based thematic role (e.g., `vn.role:Agent`) that it plays, its sensitiveness and factual impact attribute values.

Sentilo sentiment score $sc_{Sentilo}$ of a topic *t* can be defined as a function *f* taking the following arguments:

$$sc_{Sentilo}(t) = f(\sum_{i=0}^n sc(q_i(t)), \sum_{j=0}^m type_j(t), \sum_{k=0}^l ctx_k(t), truth(t), mod(t), sc(trig(sent)))$$

where

- $sc(x)$ is the score of an entity x as provided by the CombinedScore algorithm;
- $q_i(t)$ is an object value of a triple $\mathbf{t}$ `dul:hasQuality` q_i . Such triples represent direct opinion features, i.e. adjectives and adverbs, associated with entities composing the opinion sentence;
- $type_j(t)$ is a type of t expressed in the RDF graph by means of `rdf:type` triples;
- $cxt_k(t)$ is a context of t , if any. It can be either a situation or an event, which t participates in;
- $truth(t)$ is a truth value associated with t , where t is typically an event or situation occurrence, or a quality. If its value is *false* it means that the entity is negated. E.g. in a sentence such as “*John is not a good guy*”, a RDF triple `situation_1 boxing:hasTruthValue boxing:False` would be included in the graph, and its effect would be to change the sign of the sentiment score assigned to the feature *good*;
- $mod(t)$ is a marked modality of a topic t , if any. E.g. in a sentence such as *I would like a dog*, an RDF relationship `fred:like_1 boxing:hasModality boxing:Necessary` would be included. At this time, Sentilo propagation algorithm does not yet use this information, but its abstract model, the f function, includes it;
- $trig(sent)$ is an opinion trigger expression in the sentence containing t .

Function f is an abstract model referring to the *Sentiment Propagation Algorithm* (SP) implemented by Sentilo for computing $sc_{Sentilo}$ of a topic t . Fig. 8 depicts a flowchart describing the main steps of the SP algorithm, which are described in details in the rest of this section. The detailed description refers to a running example based on the sample opinion sentence of Figure 6 and Figure 9.

1. Given a frame-based RDF representation of an opinion sentence, we identify its topics and subtopics. A detailed description of the algorithm for topics and subtopics identification can be found in [18].
Figure 6 shows topics and subtopics identified by Sentilo for the sample sentence, i.e. `fred:condemn_1` (topic), `fred:judge_1` (subtopic), `fred:President` (subtopic);
2. All trigger events are assigned with a sentiment score computed by means of the CombinedScore() algorithm (cf. Figure 7). In our reference sample, we assign a score value of 0.180 to the trigger event `fred:hope_1` as shown in Figure 9;
3. Let *pos* and *neg* be two empty arrays that will be filled with sentence-level sentiment scores;
4. We create a sorted list L of topics and subtopics as follows: subtopics being the subject of any relation in the set: `{playsSensitiveRole,`

Fig. 8 Sentilo Sentiment Propagation algorithm: flowchart.

`isPositivelyAffectedBy`, `isNegativelyAffectedBy`}¹⁵ are placed ahead in the list, all the remaining subtopics are appended after, and main topics are inserted at last. According to Figure 6

$L = \{\text{fred:President}, \text{fred:judge_1}, \text{fred:condemn_1}\};$

5. Topics in L are processed starting from the entry in the first position as follows: if an entry t_i (being either a topic or a subtopic) has a `sentilo:playsSensitiveRole` relation to another element t_j of L not yet processed, then the algorithm passes to process the next element in the list and t_i is moved to the tail of L . For example, at the first iteration $t_i = \text{fred:President}$, which has a `sentilo:playsSensitiveRole` relation to `fred:condemn_1`, which in turn is included in L . Consequently our list is modified as $L = \{\text{fred:judge_1}, \text{fred:condemn_1}, \text{fred:President}\}$. The rationale behind this is that if t_i has a `sentilo:playsSensitiveRole` to t_j it means that the computation of t_i score is a function of the score of t_j , hence the latter has to be computed before the former;
6. Let pos_t and neg_t be two empty arrays that will be filled with individual sentiment scores computed for each topic t_i ;
7. If the current topic t is a subtopic (or a topic participating in a trigger event), we check if it plays a sensitive role in some situation or event, i.e. there is a `sentilo:playsSensitiveRole` relation between t and a situation or an event;

¹⁵ We omit the prefix `sentilo:` for the sake of readability and brevity.

8. If condition in step 7 is true, we identify the frames i.e., situations and events, having an assigned sentiment score sc_f , in which t plays a sensitive role, i.e. t is the subject of a `sentilo:playsSensitiveRole` relation to that situation or event.
 We propagate sc_f to the subtopic, considering possible negations `boxing:hasTruthValue`, according to its factual impact value, i.e. if the factual impact value is negative (i.e., `sentilo:isNegativelyAffectedBy`), we change the score sign. The resulting sentiment score is added to the appropriate pos_t or neg_t array. In our sample, when t is `fred:condemn_1`, condition 6 becomes true as this topic participates with a sensitive role to the triggering event `fred:hope_1`. Hence, we propagate the score of `fred:hope_1` to `fred:condemn_1` as it is shown in Figure 9.
 In turn, when t is `fred:President`, we propagate the score of `fred:condemn_1` to `fred:President` by changing its sign, as the subtopic `fred:President` is negatively affected by the event `fred:condemn_1`. Figure 9 shows that `fred:President` is associated with a score of value -0.180 ;
9. We check if t is associated with one or more opinion features, e.g. if it is the subject of `dul:hasQuality` relations. If this is the case, we add the opinion features' score to the appropriate pos_t or neg_t array of t , considering possible negation, i.e. `boxing:hasTruthValue` relation, the presence of which would cause the change of the score sign. In our example, `fred:judge_1` is not associated with any opinion feature;
10. Once all sentiment scores affecting t have been identified, we compute the positive sentiment score $sc(t)_{pos}$ associated with t as the average value of the sentiment scores in pos_t array and the negative sentiment score $sc(t)_{neg}$ as the average value of the sentiment scores in neg_t array and formally represent them by means of `sentilo:hasScore` relations in the RDF graph. The two scores are then added to the pos and neg arrays for contributing to the evaluation of the overall sentence tone;
11. If t is a trigger event carrying a positive (negative) sentiment, we add its sentiment score to the pos (neg) array, considering possible `boxing:hasTruthValue` relations which would invert the sign of the score;
12. Once all elements in L have been processed, we compute the overall sentence sentiment positive and negative scores as the average values of the sentiment scores in pos array and the average value of the sentiment scores in the neg array, respectively.

Summarizing: referring to the sample sentence “*People hope that the President will be condemned by the judges*”, after running the Sentiment Propagation algorithm, Sentilo returns the RDF graph, a fragment of which is depicted in Figure 9. The graph correctly shows that the event `fred:hope_1` has a positive sentiment score (the value 0.180 has been calculated according to the

Fig. 9 Sentilo Sentiment Propagation algorithm result for the sentence *People hope that the President will be condemned by the judges.*

CombinedScore() algorithm show in Figure 7). This score is propagated to the event `fred:condemn_1`, and in turn it is propagated (c.f. current topic t plays a sensitive role in the event `fred:condemn_1`) to its subtopic `fred:President` by changing its sign to negative (given the `sentilo:isNegativelyAffectedBy` relation from `fred:President` to `fred:condemn_1`). The subtopic `fred:President` inherits a negative score from `fred:condemn_1` because the thematic role `vn.role:Theme` for the verb *condemn* is annotated with a negative factual impact value (cf. Section 4.2). The score -0.18 is therefore assigned to the subtopic *President* (as shown in the graph) and to the sentence-level score (not shown in the graph).

6 Implementation details

Sentilo has been developed on top of FRED [35]. Apache Felix¹⁶ has been used as the reference OSGi framework for deploying Sentilo modules as JAVA OSGi bundles. A dedicated server with C&C, Boxer, FRED and Sentilo has been set up and it is freely available online¹⁷. Sentilo is available as a REST service to client applications as well as a human-oriented demo¹⁸.

The human-oriented demo serves mainly as a demonstrator of its capability, while the REST service has a clear commercial potential: the granularity of its results allows a client application to aggregate opinion analysis coming from different sources, on a specific or a general domain, based on shared holders and/or topics. Potential applications could be of interest of all stakeholders dealing with opinions or reviews that have interest in performing data analytics on such opinions during time. Additionally, political parties as well as companies putting new products on the market are potential users of Sentilo as they can use it to monitor the impact of a new brand/product on target users by analyzing their opinions.

In order to explain what the graphical user interface shows, let us consider the following sample sentence:

¹⁶ Apache Felix: <http://felix.apache.org/>.

¹⁷ Sentilo, <http://wit.istc.cnr.it/sentilo-release/sentilo>

¹⁸ Sentilo Advanced User Interface, <http://wit.istc.cnr.it/stlab-tools/sentilo/ui/sentence.html>

Tim Burton thinks that Johnny Depp is a great actor.

Figure 10 shows Sentilo graphical user interface (targeted at generic web users). On the top-left of the page there is a text area where a user can type an opinion sentence. On the right side of the text area there are two speedometers (which we have named “sentilometers”), which initially are set to zero. The first one is green and measures the positive sentiment of the whole sentence, the second one is red and measures the negative sentiment of the whole sentence. Both can have values between 0 and 1 indicating a low intensity (values closer to 0) or high intensity (values closer to 1).

Sentilo identifies all topics that are in the scope of an opinion and computes a score for each of them. Such topics and associated scores are visualized, under the text area. For each topic, the sentilometers are three: the first (green) that measures (with values from 0 to 1) the positive charge of the sentiment associated with the topic, the second (red) that measures (with values from 0 to 1) the negative charge of the sentiment associated with the topic, and the third (red to green) that measures the average sentiment score of the topic. Furthermore, if a topic identity can be resolved on Linked Data, a descriptive image for the topic is shown, as well as additional information of possible interest to users.

In Figure 10, *Tim Burton* is correctly identified as the opinion holder and its image is displayed because *Tim Burton* identity can be resolved on DBpedia. The sentilometers for *Tim Burton* are set to 0 because there is no opinion expressed on him. On the other hand, *Johnny Depp* is correctly recognized as a topic, he is resolved on DBpedia and his image is shown. His associated positive sentilometer indicates a positive opinion expressed by the term *great* which has a positive score according to our sentiment score computing algorithm.

A technical-oriented graphical interface is also available: it shows the resulting RDF graph.

Fig. 10 Sentilo interface with *sentilometers*, targeted at generic web users: results for the sentence *Tim Burton thinks that Johnny Depp is a great actor*

7 Results

We have evaluated Sentilo performance at computing overall sentence sentiment polarity on a corpus of user-based hotel reviews selected from TripAdvisor (retrieved on January 10th, 2014). We have randomly selected two sets of 50 reviews. The first set filled with the most recent 50 5-stars rated reviews, the other set filled with the most rated 1-star rated reviews. The resulting corpus can be freely downloaded¹⁹.

We excluded, for the moment, intermediate rated reviews as the average sentiment for them would not be clearly comparable with the user-annotation (star-rating). Such reviews, probably containing positive as well as negative opinions will be useful for a further evaluation of Sentilo performance at the topic-level. However, this task requires a time-consuming manual annotation effort that we leave to future work.

With this experiment we wanted to assess Sentilo ability to automatically compute sentence-level sentiment analysis on user-based reviews. We have compared Sentilo results with user-based rating. Each review contains more than one sentence. In order to determine the polarity of the overall tone of a review, we count the number of sentences with overall positive sentiment score, and the number of sentences with overall negative score. We also compute the average positive score (from all positive sentences), and the average negative score (from all negative sentences). The absolute values are then compared, and the higher one determines the overall tone of a review. As an example, here we report the averaged sentiment results for the following review:

“My son and I stayed at Le 123 Sebastopol for 5 nights. The location of the hotel is fantastic. They were so incredibly accommodating to us.

¹⁹ <http://www.stlab.istc.cnr.it/documents/sentilo/reviewsposneg.zip>

	# correlating	# non-correlating	% correlating
Positive	49	1	0.98
Negative	32	18	0.64
Overall	81	19	0.81

Table 2 Results for the correlation between Sentilo computed sentiment scores for user-based reviews and user-assigned open rating.

There were many nights when I stayed up all night and hung out in the lobby working. The staff was so nice. They brought me water, food — all without asking. The staff definitely went above and beyond to be helpful to make sure we had a great experience in Paris. I firmly believe that we would not have had that same experience at other hotels in Paris. Le 123 Sebastopol is a quintessential boutique hotel. I can see why it is rated so high by many at Trip Advisor. Only downside..... the standard rooms can be a bit small. We upgraded to a bigger room and were very content."

Sentence-level sentiment scores:

posReview: 1.9 contPos: 8 avg: 0.237

negReview: -0.882 contNeg: 4 avg: -0.2205

This means that for this review Sentilo computed eight positive sentences (contPos) with an overall positive sentiment of 1.9 (average 0.237), four negative sentences (contNeg) with an overall negative sentiment of -0.882 (average -0.2205). In such a case the review has been considered positive as 0.237 is greater than the absolute value of -0.2205.

Performance results are summarized in Table 2 as correlation values. The reader may notice that Sentilo performs better with positive reviews (0.98 correlation) with respect to negative reviews (0.64 correlation). Overall the average correlation value is 0.81.

8 Conclusions and future work

In this paper we have presented an upgraded version of Sentilo, a novel sentic computing system for sentiment analysis introduced in its first version in [18]. Sentilo combines natural language processing techniques with knowledge representation and makes use of affective knowledge resources such as SenticNet [12], SentiWordNet [2] and the SentiloNet resource of annotated verbs, presented as novel contribution in this article. Additional contribution includes an extension of OntoSentilo, an ontology for opinion sentences, which defines frame-based semantic relations between topics and subtopics of opinion sentences. We have introduced the concepts of *sensitiveness* and *factual impact* as attributes of thematic roles of frames. The former indicates if a subtopic is indirectly affected by opinions that are directly expressed on its context,

typically an event or situation occurrence. The latter indicates if the context of a subtopic (typically referred to by a verb) has a positive or negative impact on it. These concepts provide the conceptual foundation for the design of SentiloNet that, in its current version, includes 1,100 annotated verbs.

Based on SentiloNet and OntoSentilo, we have designed and implemented an algorithm for computing sentiment scores at topic- and sentence- level. This algorithm is able to propagate the scores from context to subtopics so as to enable a fine-grained analysis of opinions. Sentilo approach has been implemented and is available as REST service²⁰.

The use of a semantic-web-aware machine reader like FRED [35] allows Sentilo to resolve the identity of entities involved in an opinion on resources like DBpedia and WordNet. Furthermore, FRED foundation on cognitive frames also supports future development in solving tasks such as resolution of sarcasm or other emotions [14] or development of computational models for emotions [43].

We have tested Sentilo on a corpus of user-based reviews retrieved from TripAdvisor, which has shown encouraging results with an average correlation value 0.81. Further experiments and comparisons with other sentiment analysis tools and focused on topic-level sentiment analysis are under development.

Ongoing work concentrates on extracting opinion features from all graph patterns that are generated by Sentilo, and on designing an algorithm to calculate aspect-based opinion scores. Computational intelligence methods (cf. Fig. 1), including *fuzzy reasoning*, *combinatorial optimization on graph mining*, *learning mechanisms*, *sentic pattern discovery* and *analogical reasoning* are under investigation as possible extensions of Sentilo's approach.

Acknowledgment

The work described in this paper was performed with the support of the PRISMA (PiattafoRme cloud Interoperabili per SMArt-government) project, funded by the MIUR (Ministero dell'Istruzione, dell'Università e della Ricerca).

References

1. C. Ovesdotter Alm, D. Roth, and R. Sproat. Emotions from text: Machine learning for text-based emotion prediction. In *Proceedings of HLT/EMNLP*, pages 347–354, Vancouver, Canada, 2005.
2. A. Baccianella, S. Esuli, and F. Sebastiani. SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. In N. Calzolari, K. Choukri, B. Maegaard, J. Mariani, J. Odiijk, S. Piperidis, M. Rosner, and D. Tapias, editors, *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*, Valletta, Malta, 2010.
3. C. Bizer, T. Heath, and T. Berners-Lee. Linked Data - The Story So Far. *International Journal on Semantic Web and Information Systems*, 5(3):1–22, 2009.

²⁰ Sentilo, <http://wit.istc.cnr.it/sentilo-release/sentilo>

4. J. Bos. Wide-coverage semantic analysis with Boxer. In *Proceedings of the 2008 Conference on Semantics in Text Processing (STEP '08)*, pages 277–286, Stroudsburg, USA, 2008.
5. S. W. Brown, D. Dligach, and M. Palmer. VerbNet class assignment as a WSD task. In *Proceedings of the Ninth International Conference on Computational Semantics (IWCS '11)*, pages 85–94, Stroudsburg, USA, 2011.
6. K. Cai, S. Spangler, Y. Chen, and L. Zhang. Leveraging sentiment analysis for topic detection. *Web Intelligence and Agent Systems*, 8(3):291–302, 2010.
7. E. Cambria, M. Grassi, A. Hussain, and C. Havasi. Sentic computing for social media marketing. *Multimedia Tools and Applications*, 59(2):557–577, 2012.
8. E. Cambria and A. Hussain. *Sentic computing: Techniques, tools, and applications*, volume 2. Springer Verlag, Heidelberg, DE, 2012.
9. E. Cambria, A. Hussain, C. Havasi, and C. Eckl. Common sense computing: From the society of mind to digital intuition and beyond. In J. Fierrez, J. Ortega-Garcia, A. Esposito, A. Drygajlo, and M. Faundez-Zanuy, editors, *Biometric ID Management and Multimodal Communication*, volume 5707 of *Lecture Notes in Computer Science*, pages 252–259. Springer Verlag, Heidelberg, DE, 2009.
10. E. Cambria, B. Schuller, Y. Xia, and C. Havasi. New avenues in opinion mining and sentiment analysis. *IEEE Intelligent Systems*, 28(2):15–21, 2013.
11. E. Cambria, Y. Song, H. Wang, and A. Hussain. Isanette: A common and common sense knowledge base for opinion mining. In M. Spiliopoulou, H. Wang, D. Cook, J. Pei, W. Wang, O. Zaiane, and X. Wu, editors, *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, pages 315–322, 2011.
12. Erik Cambria, Daniel Olsher, and Dheeraj Rajagopal. Senticnet 3: A common and common-sense knowledge base for cognition-driven sentiment analysis. In Carla E. Brodley and Peter Stone, editors, *Twenty-Eight AAAI Conference on Artificial Intelligence*. AAAI Press, Palo Alto, California, July 2014.
13. H. Chen and Z. Wuand P. Cudré-Mauroux. Semantic Web Meets Computational Intelligence: State of the Art and Perspectives. *IEEE Computational Intelligence Magazine*, 7(2):67–74, 2012.
14. D. Das and S. Bandyopadhyay. Sentence-Level Emotion and Valence Tagging. *Cognitive Computation*, 4(4):420–435, 2012.
15. C. D. Elliott. *The Affective Reasoner: A Process Model of Emotions in a Multi-agent System*. PhD thesis, Northwestern University, Evanston, USA, 1992. UMI O. No. GAX92-29901.
16. C. Fellbaum. WordNet: An Electronic Lexical Database. *The MIT Press*, 1998.
17. A. Gangemi and V. Presutti. Towards a pattern science for the semantic web. *Semantic Web*, 1(1,2):61–68, 2010.
18. A. Gangemi, V. Presutti, and D. Reforgiato Recupero. Frame-based detection of opinion holders and topics: A model and a tool. *IEEE Computational Intelligence Magazine*, 9(1), 2014.
19. B. Goertzel, K. Silverman, C. Hartley, S. Bugaj, and M. Ross. The Baby Webmind Project. In *Proceedings of The Annual Conference of The Society for the Study of Artificial Intelligence and the Simulation of Behaviour (AISB)*, pages 147–148, 2000.
20. M. Hu and B. Liu. Mining opinion features in customer reviews. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI'04)*, pages 755–760, 2004.
21. R. Johansson and A. Moschitti. Relational features in fine-grained opinion analysis. *Computational Linguistics*, 39(3):473–509, 2013.
22. H. Kamp. A theory of truth and semantic representation. In J. A. G. Groenendijk, T. M. V. Janssen, and M. B. J. Stokhof, editors, *Formal Methods in the Study of Language*, volume 1, pages 277–322. Mathematisch Centrum, Amsterdam, NE, 1981.
23. A. Kazemzadeh, S. Lee, and S. S. Narayanan. Fuzzy logic models for the meaning of emotion words. *IEEE Computational Intelligence Magazine*, 8(2):34–49, 2013.
24. R. Lau, Y. Xia, and Y. Ye. A probabilistic generative model for mining cybercriminal networks from online social media. *IEEE Computational Intelligence Magazine*, 9(1):31–43, 2014.
25. J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, and C. Bizer. DBpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web Journal*, 2014.

26. B. Levin. *English Verb Classes and Alternations A Preliminary Investigation*. University of Chicago Press, Chicago, USA, 1993.
27. C. Lin, Y. He, R. Everson, and S. Ruger. Weakly supervised joint sentiment-topic detection from text. *IEEE Transactions on Knowledge and Data Engineering*, 24(6):1134–1145, 2012.
28. B. Liu. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2012.
29. Semantic Engines LLC. Opinion crawl. <http://opinioncrawl.com/>, 2010.
30. N. Nicolov, F. Salvetti, J. Martin, and M. Liberman. *Computational Approaches to Analysing Weblogs: Papers from 2006 AAAI Spring Symposium*. AAAI Press, Menlo Park, USA, 2006.
31. University of Stanford. Stanford Sentiment Analysis. <http://nlp.stanford.edu/sentiment/>, 2014.
32. B. Pang and L. Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics (ACL '04)*, Barcelona, Spain, 2004.
33. B. Pang and L. Lee. *Opinion mining and sentiment analysis*, volume 2 of *Foundations and trends in information retrieval*. now Publishers Inc., Delft, Netherlands, 2008.
34. B. Pang, L. Lee, and S. Vaithyanathan. Thumbs Up?: Sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing (EMNLP '02)*, volume 10, pages 79–86, Stroudsburg, USA, 2002.
35. V. Presutti, F. Draicchio, and A. Gangemi. Knowledge extraction based on discourse representation theory and linguistic frames. In *Knowledge Engineering and Knowledge Management*, volume 7603 of *Lecture Notes in Computer Science*, pages 114–129. Springer Verlag, Heidelberg, DE, 2012.
36. H. Saif, Y. He, and H. Alani. Semantic sentiment analysis of twitter. In *Proceedings of the 11th International Conference on The Semantic Web (ISWC'12)*, volume Part I, pages 508–524, Boston, MA, 2012. Springer Verlag.
37. Sentiment 140. <http://www.sentiment140.com/>, 2013.
38. Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*, 2013.
39. Social mention. <http://www.socialmention.com/>, 2011.
40. S. Somasundaran, J. Wiebe, and J. Ruppenhofer. Discourse level opinion interpretation. In *Proceedings of the 22nd International Conference on Computational Linguistics (COLING '08)*, volume 11, pages 801–808, Manchester, UK, 2008.
41. S. O. Sood, S. Owsley, K. J. Hammond, and L. Birnbaum. Reasoning through search: a novel approach to sentiment classification. In *Proceedings of the 16th International World Wide Web (WWW) Conference*, Banff, Canada, 2007.
42. S. O. Sood and L. Vasserman. ESSE: Exploring mood on the web. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM)*, Seattle, USA, 2009.
43. D. N. Tam. Computation in emotional processing: Quantitative confirmation of proportionality hypothesis for angry unhappy emotional intensity to perceived loss. *Cognitive Computation*, 3(2):394–415, 2011.
44. I. Titov and R. McDonald. Modeling online reviews with multi-grain topic models. In *Proceedings of the 17th International World Wide Web (WWW) Conference*, pages 111–120, Beijing, China, 2008.
45. P. D. Turney. Thumbs Up or Thumbs Down?: Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics (ACL '02)*, pages 417–424, Philadelphia, Pennsylvania, 2002.
46. J. Wiebe, T. Wilson, and C. Cardie. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2-3):165–210, 2005.
47. T. Wilson, J. Wiebe, and 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), pages 347–354, Vancouver, Canada, 2005.