Cognitive Computation (2019) 11:469–488
https://doi.org/10.1007/s12559-019-9625-x

ReUS: a Real-time Unsupervised System For Monitoring Opinion Streams

Mauro Dragoni[1] · Marco Federici[2] · Andi Rexha[3]

Received: 9 November 2017 / Accepted: 3 January 2019 / Published online: 7 February 2019

Abstract

An actual challenge within the sentiment analysis research area is the extraction of polarity values associated with specific aspects (or opinion targets) contained in user-generated content. This task, called *aspect-based sentiment analysis*, brings new challenges like the disambiguation of words' role within a text and the inference of correct polarity values based on the domain in which a text occurs. The former requires strategies able to understand how each word is used in a specific context in order to annotate it as aspect or not. The latter need to be addressed with unsupervised solutions in order to make a system efficient for real-time tasks and at the same time flexible in order to adopt it in any domain without requiring the training of sentiment models. Finally, the deployment of such a system into real-world scenarios needs the development of usable solutions for accessing and analyzing data. This paper presents the ReUS platform: a system integrating an unsupervised approach, based on open information extraction strategies, for performing real-time aspect-based sentiment analysis together with facilities supporting decision-makers in the analysis and visualization of collected data. The ReUS platform has been validated from a quantitative and qualitative perspectives. First, the aspect extraction and polarity inference capabilities have been evaluated on three datasets used in likewise editions of SemEval. Second, a user group has been invited to judge the usability of the platform. The developed platform demonstrated to be suitable for being used into real-world scenarios requiring (i) the capability of processing real-time opinion-based documents streams and (ii) the availability of usable facilities for analyzing and visualizing collected data. Examples of possible analysis and visualizations include the presentation of lists ranking aspects by the importance of their polarity values computed within the whole data repository. This kind of analysis enables, for instance, the discovery of product issues.

Keywords Sentiment analysis · Opinion mining · Unsupervised aspect extraction · Real-time monitoring

Introduction

The possibility of expressing personal opinions about products and services attracted people on social media for discussing about their preferences and for debating about many topics. This opportunity opened the era of "trusting" the Web for finding which are the best restaurants or the best products that satisfy specific needs. This fact enabled the creation of communities able to influence people in taking decisions, but also to collect feedback and issues about product and services. Hence, in the last years, the effort dedicated to the extraction of extracting useful information from user-generated content and for understanding their actual meaning has significantly grown.

Sentiment analysis is a branch of affective computing research [1] that aims to classify text—but sometimes also audio and video [2]— into either positive or negative—but sometimes also neutral [3]. Most of the literature is on English language [4, 5] but recently an increasing number of works are tackling the multilinguality issue [6], especially in booming online languages such as Chinese [7].

Sentiment analysis techniques can be broadly categorized into symbolic and sub-symbolic approaches: the former

✉ Mauro Dragoni
dragoni@fbk.eu

Marco Federici
marco.federici@uva.nl

Andi Rexha
arexha@know-center.at

1 Fondazione Bruno Kessler, Trento, Italy

3 University of Amsterdam, Amsterdam, The Netherlands

2 Know-Center, Graz, Austria

include the use of lexicons [8] and ontologies [9] to encode the polarity associated with words and multiword expressions; the latter consist of supervised [10], semi-supervised [11], and unsupervised [12] machine learning techniques that perform sentiment classification based on word co-occurrence frequencies. Among these, the most popular recently are algorithms based on deep neural networks [13] and generative adversarial networks [14].

While most works approach it as a simple categorization problem, sentiment analysis is actually a suitcase research problem [15] that requires tackling many NLP tasks, including word polarity disambiguation [16], subjectivity detection [17], personality recognition [18], microtext normalization [19], concept extraction [20], time tagging [21], and aspect extraction [22].

Sentiment analysis has raised growing interest both within the scientific community, leading to many exciting open challenges, as well as in the business world, due to the remarkable benefits to be had from financial [23] and political [24] forecasting, e-health [25] and e-tourism [26], user profiling [27] and community detection [28], manufacturing and supply chain applications [29], human communication comprehension [30] and dialogue systems [31], etc.

Early applications of sentiment analysis aimed at computing an overall polarity that was then associated with a document. The drawback of this strategy is that it was not possible to distinguish which were the subjects of each opinion and how users judged them. For this reason, the research community focused on the extraction of subjects, called *aspects*, from user-generated content. This task would make systems capable inferring the polarity of each aspect in an independent way [32–34].

As example, we can consider the following sentence:

I went into new fast-food in the city center yesterday. Even if the space is cute, the food is of a very poor quality.

The proposed example contains three aspects, *fast-food*, *space*, and *food*, each of them associated with their opinions:

- "space" → "cute"
- "food" → "very poor"
- "fast-food" → NULL. Here, no opinions are directly associated with the aspect. Thus, its polarity is implicitly inferred by aggregating the polarities of the other opinion words detected in the sentence.

The structure shown above is obtained through the execution of two tasks: (i) the extraction of aspects mentioned in each sentence and (ii) the inference of polarity values associated with them. The second task can be coped by means of sentiment lexicons (Linguistic Resources and Tools), but the first one need to be addressed by exploiting different strategies. State-of-the-art approaches discussed in Related Work leverages on the use of supervised learning solutions for detecting the aspects contained in a text. Unfortunately, the real-world clashed with the use of supervised solutions. For building a sentiment model, it is necessary to annotate a corpus which contains the description of aspects' context for all possible domains. Those datasets are available only for few domains. Secondly, sentences contained within the same document can refer to topics about more than one domain; thus, using one single model is not the best solution.

For all these reasons, it is necessary to provide solutions able (i) to extract aspects from raw text effectively, (ii) to infer the associated polarity values with a high accuracy, and (iii) to allow the visualization of all collected information into intelligent dashboards able to show real-time information about products. These tasks have to be done by analyzing opinion-based data streams; thus, also efficient solutions are needed in order to develop a system that is suitable for being delivered into a real-world scenario.

In this paper, we present the ReUS platform, an open information extraction solution coping the challenges described above aiming to provide a flexible, effective, and efficient support to the analysis of opinion-based data streams. The ReUS platform relies on the three pillars below:

- the design of an effective approach for detecting aspects occurring in user-generated content;
- the development of an efficient architecture able to process a huge number of documents in real-time; and,
- the realization of facilities supporting users in the analysis and visualization of the collected data.

State-of-the-art solutions are focused on summarizing opinion-based documents with information about the overall polarity associated with a given product. ReUS aims to support customers as well as managers with facilities allowing a deeper analysis of products like the polarity associated with a single aspect and how this polarity evolved through time. For instance, this kind of features would enable customers to take more thoughtful decisions about which products can be better for their needs and managers in observing if changes on specific features improved the overall quality of a product or not. To the best of our knowledge, the proposed solution represents the first architecture supporting this kind of in-depth analysis and visualization of aggregated data. These activities are supported by a set of facilities allowing the navigation of the extracted data that are dynamically populated by the aspect extraction process.

The paper is structured as follows: Related Work discusses the state of the art on aspect-based opinion mining with a focus on approaches working on social media. Then, Linguistic Resources and Tools introduces the background knowledge bases integrated into the proposed platform; The ReUS Platform describes the overall architecture of ReUS, while in The Document Analyzer, we discuss the open information extraction technique integrated into the platform. Then, Client User Interface presents the web application enabling the analysis of the collected data. In Validation of the ReUS Platform, we provide the validation of the platform from both a quantitative and qualitative perspective. Finally, Conclusion and Future Work concludes the paper.

Related Work

The ReUS platform falls within the context of adopting opinion mining techniques applied to social media streaming. Here, we provide an overview of the main approaches about both topics together with references of existing real-world system integrating Open Information Extraction (OIE) strategies for processing natural language text.

Opinion Mining and Sentiment Analysis

Opinion mining and sentiment analysis have been widely investigated [35, 36], especially in the last decade where proposed techniques started to be integrated into real-world systems.

From the theoretical perspective, first attempts were conducted by designing supervised solutions for analyzing the impact of different types of features extracted from raw texts. In [37], the authors compared well-known machine learning approaches (i.e., support vector machine, maximum entropy, and Naïve Bayes) and observed how they performed on datasets containing n-grams. An extension of this investigation has been presented in [38], were the authors tried to include also part-of-speech (POS) tags or to consider only certain type of words (e.g., adjectives). Results demonstrated how the more knowledge is included in the model, the more mitigated are the differences between the effectiveness of the compared algorithms.

The explosion of user-generated content on the Web was the humus for applying opinion mining approaches on such content for marketing purposes. Preliminary applications aimed to compute the overall polarity of a text. In [39], an approach for extracting opinions and for computing their polarity from social networks has been presented. This work represented a seminal study concerning the application of opinion mining strategies on social media, while in [40], the authors investigated about the impact of analyzing sentiment-based information provided in a multi-modal way and how to detect and discard noisy information from texts.

Results observed by using approaches inferring the overall polarity of user-generated contents [41, 42] opened the challenge of performing a more fine-grained analysis of the texts [43, 44]. Product reviews usually contain opinions about many features of the same products for which users provide different feedback. This kind of analysis, termed *aspect-based* sentiment analysis, has been addressed by applying different strategies. First approaches implemented clustering techniques exploiting knowledge bases for detecting similar words referring to the same aspect [45]. Other solutions validated the use of probabilistic graphical models [46] or the adoption of conditional random fields [47] for learning where aspects occur in sentences. Further, approaches relying on the use of fuzzy logic [48, 49] for modeling the uncertainty associated with polarity values or that applies multi-value aggregation strategies [50] for computing the sentiment values associated with documents and/or aspects have been proposed.

A second challenge linked to the fine-grained analysis of texts is the inference of sentence subjectivity [51]. This task consists in understanding if the opinions contained in the analyzed texts represent a personal view of a specific user or if they are objective considerations about a given entity. Research in this direction is still in an early stage, but the first results are promising.

Recently, pure computer science solutions have been enriched with a social dimension for obtaining a generation of fine-tunes solutions for the inference of implicit emotions contained in multi-modal resources. This research area is called sentic computing [52]. This kind of approach has been applied to several problems ranging from the parsing and analysis of textual user-generated content [53] to the detection of affective statuses within multimedia resources like images, audio, or video [54].

Finally, in the last years, first real-world systems started to become available. An example is the system proposed in [55, 56] called *Sentilo* that integrated NLP techniques combined with semantic web technologies for the extraction of sentiment information from texts. With respect to this system, ReUS is able to perform both the opinion mining and sentiment analysis tasks in an unsupervised way being at the same time flexible in working on different domains and efficient in processing real-time data streams.

Opinion Mining and Sentiment Analysis in Social Media

We summarized above the main research directions and approaches dedicated to the opinion mining and sentiment analysis topic. Here, we want to focus on how these techniques have been applied within the social media field. Social media brought a new way of providing information due to the use of many *slangs* and other ways of communicating between people [40].

Among all the social media, Twitter represented the most used one thanks to the public availability of many multi-modal information [39]. The possibility of accessing this huge amount of data attracted the attention of researchers and fostered the creation of systems thought ad hoc for working in a social media environment. An example of these systems is the *SentiStrenght* platform [57], the integrated machine learning strategies optimized for working with short, slangy, and error-prone texts. Observed results demonstrated the need of this kind of features for managing social media texts.

Aspect-based techniques have been adopted also in the social context where the challenge of working with short text have been addressed from different perspectives. Here, the attention has been put on the creation of two main categories of systems. On the one hand, social media are a good environment for applying computational advertising solutions [58] where messages containing links to proper product pages are automatically generated based on user profiles and the observed overall sentiments associated with product aspects. Examples of system working on Twitter [59] or Foursquare [60] have been developed and validated.

Repositories of reviews like the Amazon website have treated as a social media platform and given the amount of opinion-based documents, the research community designed a plethora of approaches for the analysis of products' reviews. Such approaches include supervised solution [47, 61–63] based on datasets manually annotated with the associations between aspects and their opinion words and rule-based systems [64] where sequential rules associated with polarity values are integrated into the algorithm. Finally, other approaches more focused on the application of pure NLP strategies like the modeling of discourse topics [65–68], the genetic learning of dependency relations [69], or on the discovery of grammatical relationships between aspects and opinion words [70] have been proposed.

The realization of the ReUS platform has been inspired by approaches relying on the use of pure NLP strategies for detecting pairs of the form $\{Aspect, Opinion\}$. Our aim is to realize a fully unsupervised platform and, by analyzing the state of the art, the adoption of pure NLP strategies brought to more effective results. With respect to the approaches presented above, ReUS has been supported in the discovery of $\{Aspect, Opinion\}$ relationships and in the computation of aspects' polarity by the integration of different knowledge bases. The result is an effective and efficient platform deployable in a real-world context.

Open Information Extraction

The field of Open Information Extraction (OpenIE) was spawned by the idea to address limitations of traditional Information Extraction, namely the dependence on predefined relations and labeled training datasets (cf.[71]). Initial work mainly focused on shallow syntactical features such as POS tags (cf. *TextRunner* [72], WOE^{pos} [73] or *ReVerb* [74]). While such approaches demonstrated efficiency, their effectiveness was still limited. The necessity of improving the quality of the extracted grammatical dependencies led to the realization of more fine-grained system integrating structured grammatical knowledge (cf. *Kraken* [75], *Ollie* [76], *ClausIE* [77] or *CSD-IE* [78]). The systems described above were designed and tuned for working only on the English language. Non-English research is limited to widely used languages such as Spanish [79], Chinese [80], and German [81]. [82] used features derived from dependency trees in their system, *DepOE*, and showed that they are suitable for English as well as Romance languages. The improved multilingual OpenIE system *ArgOE* [83] leverages on a the use of a common rule set for extracting sentences in many languages. In [79], the authors described the *ExtrHech* system that relies on semantic constraints for working on the Spanish language. Results show that the system effectively works for both the Spanish and English languages. Finally, semi-supervised solutions have been designed for the realization of OpenIE platform for the Chinese (e.g., the *SCOERE* system [80]) and German (e.g., the *PropDE* system citeFalke16) languages.

Linguistic Resources and Tools

The platform presented in this paper leverages on linguistic resources that are exploited for both the aspect extraction and polarity inference tasks. In particular, we integrated the following resources:

- a list of stopwords used for cleaning documents from useless words;[1]
- a general purpose linguistic knowledge base (WordNet);
- three sentiment lexicons (Sentiment Lexicons); and,
- a natural language processing (NLP) library used for managing document content (The Natural Language Processing Tool).

We detail below the integrated resources and the role they play into the platform.

[1]The list of stopwords we used can be found at http://lextek.com/manuals/onix/stopwords1.html

WordNet

WordNet[2] [84] is a linguistic knowledge base developed for the English language including more than 142,000 concepts representing nouns, verbs, adverbs, and adjectives. Concepts are grouped into *synsets*. Each synset represents a set of concepts sharing the same cognitive meaning, i.e., the possibility of interchanging the list of labels associated with each synset into a sentence without changing its conceptual meaning. The number of synsets contained in WordNet is around 118,000. Each synset is linked to other synsets through one, or more, of the semantic relations that have been defined. Among them, we can find synonymy (two synsets are synonyms when the share similar meanings), hypernymy (a synset A is an hypernym of a synset B, when A represent a broader concept than B), and hyponymy (a synset A is an hyponym of a synset B, when A represent a finer concept than B). Besides semantic relations, each synset is accompanied by its definition (called *gloss*) and by sentences showing its use. In case a gloss has distinct conceptual meanings, it is associated with different synsets.

At first view, WordNet can be seen as a thesaurus. However, with respect to it, there some differences that are worth of highlighting:

- each node represents a specific conceptual meaning. It means that no overlap of meanings among two words is possible. This differs from thesauri where words located at the same level may have similar meanings;
- the semantic relationships between synsets are always labeled. While in thesaurus, cluster of similar words are put in relationship without explicitly mention the kind of relations.

WordNet has been integrated into our platform for two purposes: (i) working as bridge for linking the sentiment lexicons described below, and (ii) to support the detection of compound names in the text.

Sentiment Lexicons

The role of sentiment lexicons is to support the detection of terms contained in a document for which a polarity value can be associated. Such terms, called *opinion words*, are used for inferring the overall polarity of a sentence. The literature presents a wide range of sentiment lexicons, each of them build by following different perspectives. Among them, we focused on the three resources mostly used by the affective computing community: SenticNet [85], the General Inquirer vocabulary[3] [86], and the MPQA dictionary[4] [87]. A further peculiarity of these resources is that they can be easily linked each other thanks to their relations with WordNet.

SenticNet is semantic affective computing resource publicly available created through the adoption of artificial intelligence techniques. The main goal of SenticNet is to support the inference of polarity values of concepts. All information represented in SenticNet are in a semantic flavor in order to enable the integration of reasoning facilities. With respect to other sentiment resources available, SenticNet differs for three aspects. Firstly, polarities are represented by using continuous values ranging within the interval $[-1, 1]$. Secondly, SenticNet does not contain polarity information of concepts associated with WordNet's synsets only, but also of entities (called *commonsense concepts*) describing scenarios represented by concepts aggregations like *a lot of stress*, *look for information*, or *loss of strength*. Thirdly, in SenticNet concepts that are not associated with polarity information (i.e., neutral concepts) are not reported in order to avoid issues in case of applying reasoners implementing open world assumption strategies [88]. SenticNet provides the sentiment representation of 100,000 concepts ready to use by third-party systems.

The second exploited resource is the *General Inquirer* dictionary. This dictionary contains containing around 12,000 words annotated with sentiment information associated with different contexts, namely *valence category*, *semantic dimensions*, *words of pleasure*, and *emotional expressiveness*. Briefly, *valence categories* are the well-known *positive* and *negative* ones. The *semantic dimensions* are annotations like *strong*, *weak*, or *active* that are described by the authors are language universals semantic differentials. The *words of pleasure* are annotations representing a more fine-grained classification of the *valence categories*. Examples of *words of pleasure* are *feel*, *virtue*, and *pain*. Finally, the *emotional expressiveness* are annotations used for overestimating or underestimating the other annotations associated with a word. Words can be annotated with one or more of the annotations described above. The reader may refer to the resource website for gathering more detailed information.

The third resource integrated into our platform is the *MPQA* lexicon. Such a lexicon defines the polarity of 8,200 terms annotated with an intensity level. The *MPQA* includes the three canonical polarity values *positive*, *neutral*, and *negative*, while the intensity levels are two annotations, *weak* and *strong*, that, combined with the polarity values, allow the expression of more fine-grained sentiments without using continuous representations. The concepts contained in the lexicon have been used for annotating

[2] http://wordnet.princeton.edu

[3] http://wjh.harvard.edu/~inquirer/spreadsheet_guide.htm

[4] http://mpqa.cs.pitt.edu/corpora/mpqa_corpus/

10,000 sentences that are available within the resource and that can be used from machine learning algorithms in order to build sentiment models. The reader may refer to [89] for the details about the schema used for annotating the sentences with *MPQA* concepts.

The three resources are partially disjointed and no one is a subset of another one. Indeed, even if the *SenticNet* defines the polarity of 100,000 concepts, it does not contain all the concepts defined in the *General Inquirer* and in the *MPQA*. Among all the concepts available in the three resources mentioned above, we integrated into our platform the definition of the concepts having a polarity value different than zero (i.e., not *neutral*) in at least one of the three resources; then, the aggregation of the polarities extracted from each resource is performed by averaging the (i) the polarity extracted from the *SenticNet* that is a continuous value ranging in the interval $[-1, 1]$; (ii) the polarity of the *MPQA* that uses a three-values scale $[-1, 0, 1]$ which values are halved if a concept is annotated with the intensity level *weak*; and (iii) the polarity of the *General Inquirer* that are eventually transformed like the *MPQA* ones.

The Natural Language Processing Tool

The ReUS platform leverages on the Stanford Core NLP Library [90] for processing the streams of documents in order to transform raw texts into their equivalent structured representations. Among the ecosystems of tools included into the Stanford Core NLP library, we integrated into ReUS the POS Tagger [91] for tagging each word with the appropriate tag (i.e., noun, verb, etc.); the Co-reference Annotator [92] that is used for detecting implicit references of the same entity; the Syntactic Parser [93] that is in charge of extracting the grammatic tree of a sentence; and the Dependency Parser [94] that is responsible of extracting the grammatical relationships occurring between sentences' words. An example of the Dependency Parser output is shown in Fig. 4.

The ReUS Platform

The architecture of the ReUS platform is built upon a set of components enabling the processing and the transformation of raw documents through the implemented pipeline. The current version of ReUS is tuned for working on the document's stream of the Amazon website where each document is a product review generated by a user. The overall architecture of ReUS is shown in Fig. 1. Products' reviews are extracted from the document stream and sent to the ***Data Manager***. This component is in charge of analyzing the raw text and of annotating it with metadata like the review's timestamp and the user's identifier. Data generated by this component are stored with a semantic format into the built knowledge store that makes them available to our ***Web Service*** automatically. Such a service manages query requests and exposes stored reviews in a structured format.

The pipeline's process is organized as follows. Reviews contained into the document stream are collected by the *Data Manager* that in turn integrates two modules. The first module, called *Data Analyzer*, contains our OpenIE algorithm that is in charge of transforming the raw document extracted from the stream into its annotated representation. Annotations included in this stage are (i) the words marked as aspects and (ii) the polarity values associated with each aspect. This module contains several components. We detail them in The Document Analyzer.

Fig. 1 Architecture of the ReUS platform

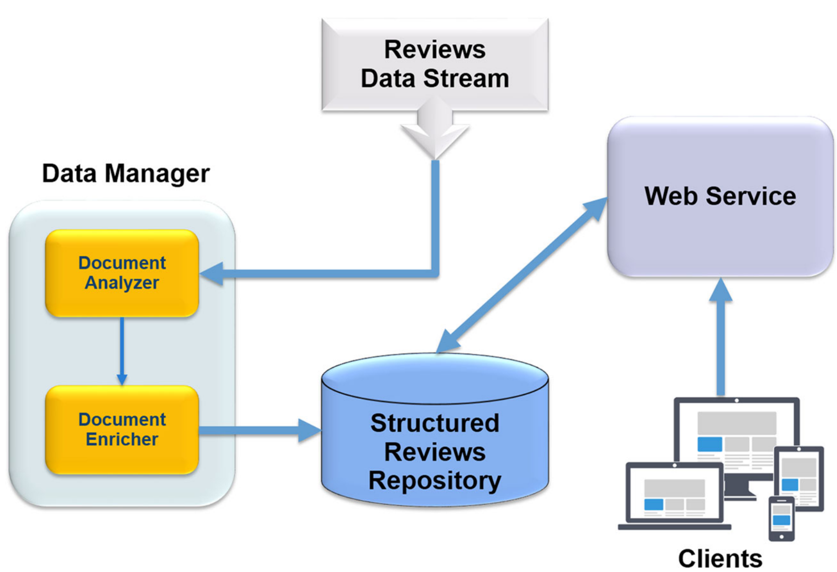

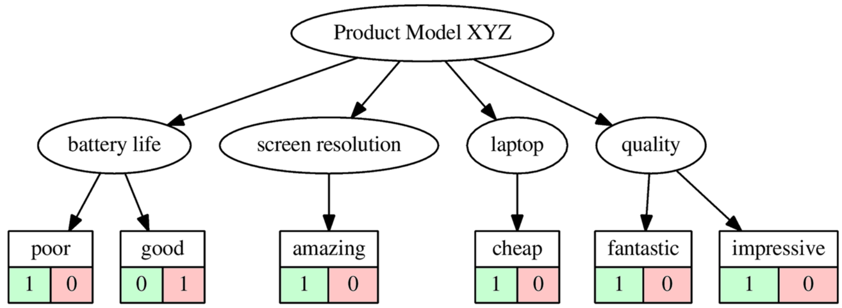

Fig. 2 Tree-based representation of a product review. The root contains the general product information, the first level the list of extracted aspects, and the leaves contain the associated opinion words with the computed polarity value

The second module is the *Document Enricher*. This module is in charge of annotating the structured representation of reviews created by the *Data Manager* with links to Amazon entities. In particular, included are the product's identifier, its name, the product group, and the original score assigned by the user.

Once these tasks have been performed, the *Data Manager* produces the tree structured shown in Fig. 2 and stores it within the reviews' repository.

The root contains overall product information that are collected by means of the Amazon APIs. The first level contains one aspect in each node. Here, synonyms are already resolved; thus, words like *monitor* and *display* have been already clustered under a common term. Finally, for each aspect, the associated opinion words and polarity value are extracted.

The last module of the platform is the *Web Service*. This module exposes a set of APIs enabling the querying of the repository by the integrated client by third-party tools. Through this module, we are able to support the visualization of the product's data almost in real-time. Hence, it is possible to analyze how the reputation of a product, or a particular aspect of it, evolves through time.

The Document Analyzer

We describe in this section the ecosystem of techniques integrated into the ReUS platform for processing the raw text in order to extract candidate aspects. Such an ecosystem is composed by five NLP modules shown in *Parsing Pipeline* package of Fig. 3. Briefly, the five modules provide the following functionalities:

- Extraction of candidate aspects. Raw text is parsed by applying the algorithm described in OpenIE Approach for Aspect Extraction for detecting which are the candidate aspects contained in the text. Once aspects are detected, the extractor looks for the candidate opinion words associated with them.
- Extraction of compound nouns. The extraction of compound names is a challenge tasks aiming to detect text chunks representing the same name. Example of compound names are *hard disk*, *screen resolution*, etc. This step is performed by exploiting the annotation provided by the POS Tagger and WordNet. Briefly, if two, or more, words annotated as nouns occur consecutively, the module looks up in WordNet if the

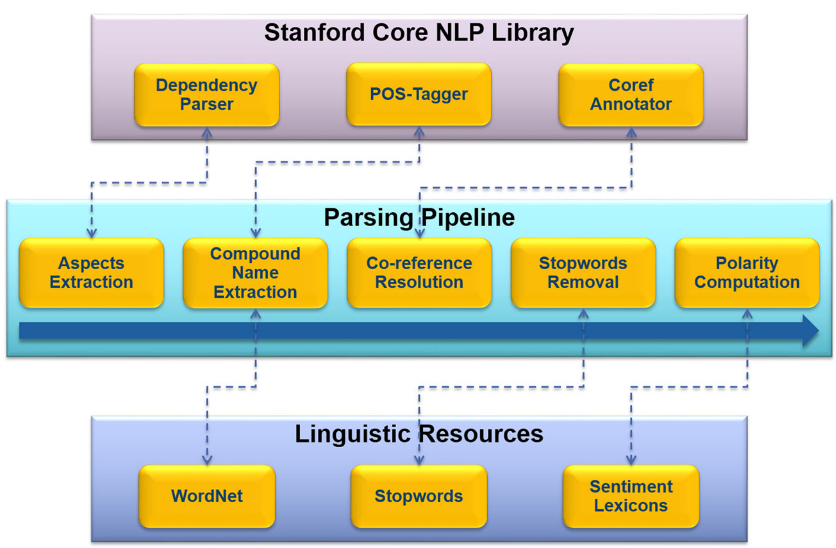

Fig. 3 The natural language processing pipeline implemented within the proposed platform aiming to extract aspects and compute their polarity from the analyzed textual resources

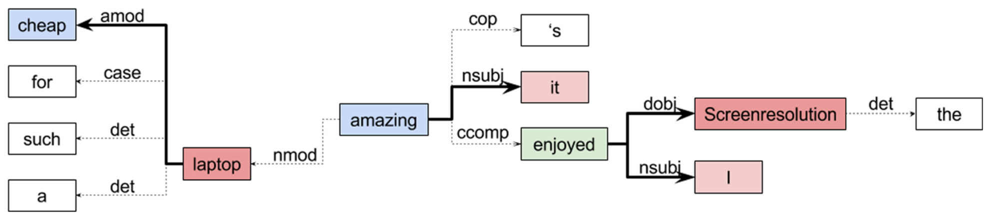

Fig. 4 The figure shows an example of the grammatical dependencies graph created by ReUS

concatenation of the words exists or not. In case it is found, the concatenation of such nouns are considered as a single word.

- Co-reference resolution. Pronouns are often used as replacement of aspect's name. In order to correctly extract the whole set of opinion words associated with an aspect, a co-reference algorithm is necessary. In the current version of ReUS we implemented a preliminary version of this kind of approach. We planned to improve the effectiveness of this module in the next version of ReUS.
- Removal of stopwords. The list of stopwords injected into ReUS is adopted for cleaning the text by removing useless words. This way, only words of interest are preserved in the text in order to ease the polarity computation step.
- Computation of aspect polarity. The last step is in charge of inferring which is the polarity value associated with each extracted aspect. Given a list of n opinion words linked to the aspect A, the overall polarity of A is computed as the aggregation of the n polarity values.

As it is possible to notice, the first three steps of the pipeline are connected with the integrated NLP library containing the algorithms for parsing the raw text. While, the lower layer shows how the adopted linguistic resources are exploited by the pipeline for the processing of document content.

OpenIE Approach for Aspect Extraction

The extraction of aspects leverages on an OpenIE strategy analyzing the morphology of sentences for understanding which are the elements of interest containing it. In Related Work we discussed the state of the art of OpenIE strategies and tools. The algorithm implemented within ReUS extends the facilities available into the *ClausIE* system by adding the management of grammatical rules specifically thought for coping the aspect extraction challenge.

The proposed strategy receives in input the graph of grammatical dependencies built by the NLP tool represented by using the format $\{Relation\_Type, Governor, Dependent\}$.[5] This graph is processed by the OpenIE algorithm that, in tandem, applies to each graph node the three rules shown below: *Rule #1*: when the value of the $Relation\_Type$ element is the adjectival modifier ("amod"), the system generates a link between the $Governor$ and the $Dependent$. The link persists if the $Governor$ has been annotated as aspect and the $Dependent$ is annotated as opinion word, i.e., there is at least one of the sentiment lexicons containing a polarity value different than zero associated to the $Dependent$.

Rule #2: when the value of the $Relation\_Type$ element is the nominal subject ("nsubj"), the system generates a link between the $Governor$ and the $Dependent$. The link persists if the $Dependent$ has been annotated as aspect and the $Governor$ is annotated as opinion word, i.e., there is at least one of the sentiment lexicons containing a polarity value different than zero associated to the $Governor$.

Rule #3: when the value of the $Relation\_Type$ element is the direct object ("dobj"), the system generates a link between the $Governor$ and the $Dependent$. The link persists if the $Dependent$ has been annotated as aspect and the $Governor$ is annotated as opinion word, i.e., there is at least one of the sentiment lexicons containing a polarity value different than zero associated to the $Governor$.

For clarity, we included below a running example showing how the rules are applied by the OpenIE algorithm. Figure 4 shows how each node of the dependency graph is annotated by the NLP library, while in Fig. 5, we show only the nodes that are preserved by the algorithm and that are then stored into the knowledge repository. In both figures, we colored the nodes that were not annotated as *aspect* with light red, the node annotated as *aspect* with red, the nodes annotated as *opinion words*, i.e., words having a polarity value different than zero in at least one of the sentiment lexicon, in green when they are tagged as verbs and in blue when they are tagged as adjectives by the POS Tagger.

For completeness, we want to highlight how the relationships between the nodes *I* and *enjoyed* is treated. The node *I* is annotated as *PP, personal pronoun*; thus, the

[5] Details about triple's elements including their meaning and possible values of each element are available within the official documentation http://nlp.stanford.edu/software/dependencies_manual.pdf

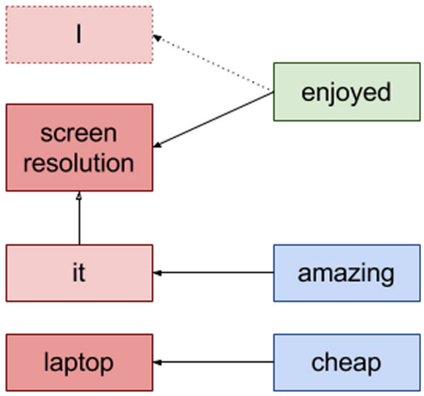

Fig. 5 The figure shows the final set of relationships preserved by ReUS and stored into the knowledge repository.

relationship is kept active until the end of the algorithm in case the co-reference annotator would be able to resolve it. In this specific case, the node is not resolved; hence, the relationship is then discarded by the OpenIE algorithm.

Hence, the relationships judged as valid by the OpenIE algorithm and stored into the knowledge repository are the following:

$$\{\text{cheap}\} \rightarrow \text{laptop}$$
$$\{\text{amazing, enjoyed}\} \rightarrow \text{screenresolution}$$

Client User Interface

The analysis and visualization of collected data is performed through a graphical user interface with which the ReUS platform has been equipped. In the current version of ReUS, we decided to implement a controlled query facility instead of an open one. The rationale behind this choice is that we wanted to support users with intuitive as well as easy functionalities for analyzing the huge amount of data stored into the knowledge repository.

The usage of the interface shown in Fig. 6 is split in the following steps. Users can select from the *Category* drop-down list the product group they want to analyze. Then, if they want, they can insert into the *Aspect* text-box keywords representing the aspect they are interested to analyze. The free-text keyword is matched with the internal dictionary of aspects that ReUS builds on-the-fly during the population of the knowledge repository. The back-end service collects all relevant information and compiled the summary of each product found within the repository.

For each item found in the repository, ReUS provides the following information:

- the full description of the product;
- the number of aspects that has been detected in all reviews related to that specific product;
- the average polarity of the reviews related to that specific product; and,

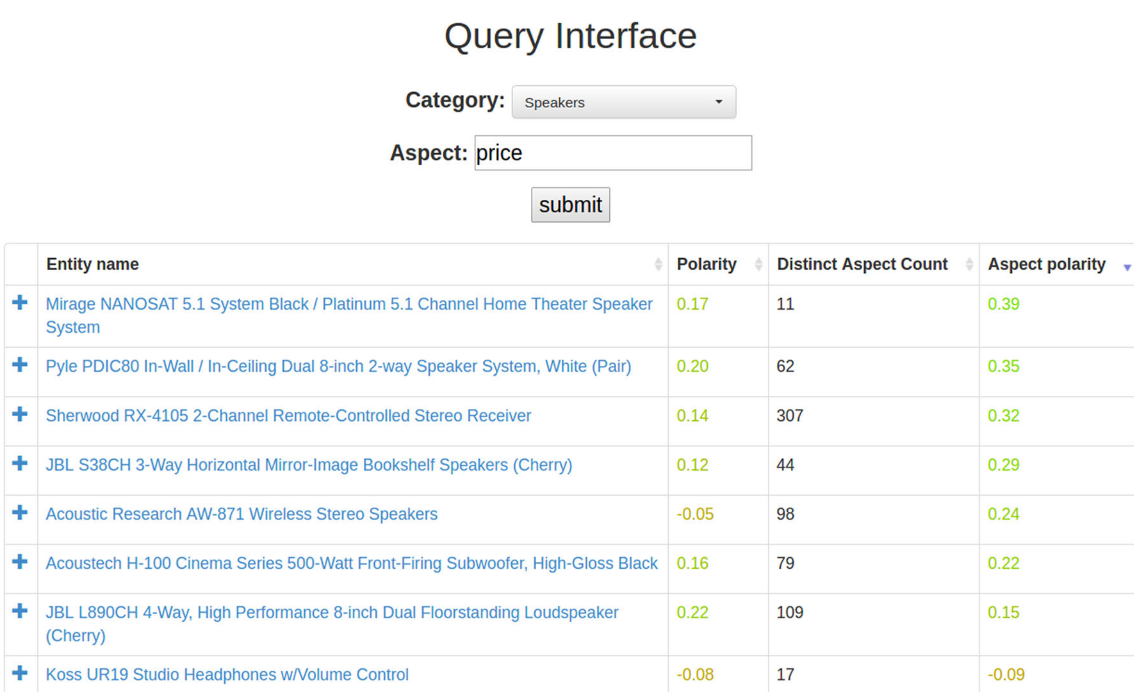

	Entity name	Polarity	Distinct Aspect Count	Aspect polarity
+	Mirage NANOSAT 5.1 System Black / Platinum 5.1 Channel Home Theater Speaker System	0.17	11	0.39
+	Pyle PDIC80 In-Wall / In-Ceiling Dual 8-inch 2-way Speaker System, White (Pair)	0.20	62	0.35
+	Sherwood RX-4105 2-Channel Remote-Controlled Stereo Receiver	0.14	307	0.32
+	JBL S38CH 3-Way Horizontal Mirror-Image Bookshelf Speakers (Cherry)	0.12	44	0.29
+	Acoustic Research AW-871 Wireless Stereo Speakers	-0.05	98	0.24
+	Acoustech H-100 Cinema Series 500-Watt Front-Firing Subwoofer, High-Gloss Black	0.16	79	0.22
+	JBL L890CH 4-Way, High Performance 8-inch Dual Floorstanding Loudspeaker (Cherry)	0.22	109	0.15
+	Koss UR19 Studio Headphones w/Volume Control	-0.08	17	-0.09

Fig. 6 The figure shows the interface developed for allowing users to query the knowledge repository

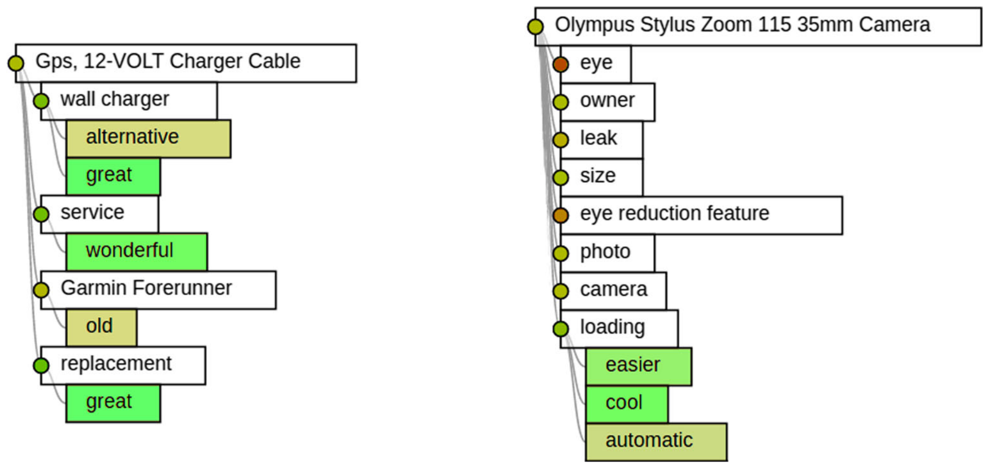

Fig. 7 The figure shows the visualization developed for showing the detailed information associated with each product contained into the knowledge repository

– the polarity associated with the aspect provided by the user within the *Aspect* text-box, if any.

From this visualization, the user can access to a more detailed report of each product. By clicking on the product name, ReUS shows the hierarchical visualization presented in Fig. 7. Here, the user can view all aspects associated with the selected product and expand them in order to see which were the opinions that have been expressed. For simplicity, the polarity value associated with each opinion is provided through a multi-graded colored background of the boxes containing opinion-word text.

Further details about each opinion word can be accessed by clicking on its name. A further table is shown to users (see Fig. 8) where a summary of sentiment information contained in each sentiment lexicon is provided. The usefulness of this view is the possibility of analyzing and comparing how each opinion word is represented within each sentiment lexicon.

The ReUS platform has been developed by adopting state-of-the-art technologies. The AngularJS[6] and Bootstrap[7] libraries have been used in synergy for developing the web interface. The back-end has been developed in Java and services are exposed through REST Servlets producing JSON-based data for showing the results of performed queries. Then, GraphDB[8] has been used for managing the knowledge repository. Finally, the visualization of complex and aggregated product's information was managed by the D3.js[9] library.

Validation of the ReUS Platform

The architecture described in this paper has been validated on three criteria:

C1. the effectiveness in extracting the correct aspects from sentences;
C2. the accuracy in inferring the correct polarity associated with each aspect; and,
C3. the usability of the tool we developed for the experts.

Criterion *C1* is important from the expert's perspective because it allows the discovery of the most important aspects within the reviews' repository and it enables the analysis performed later on. Criterion *C2* allows to observe if the strengths and weaknesses of each aspect have been inferred correctly. Finally, criterion *C3* permitted to collect users' feedback in order to understand if the system has been perceived usable or not.

The evaluation of aspect extraction and polarity detection capabilities has been performed on three benchmarks corresponding to the three editions of SemEval going from 2014 to 2016.[10] Hereafter, we refer to those datasets with the labels *SB2014*, *SB2015*, and *SB2016* for the benchmarks related to 2014, 2015, and 2016 editions of SemEval respectively.

The *SB2014* and *SB2015* datasets allowed to perform the validation of *C1* and *C2* on documents related to the *Laptop* and *Restaurant* domains. While, dataset *SB2016* included the possibility of performing the validation of *C3* also on documents related to the *Hotel* domain.

[6] http://angularjs.org
[7] http://getbootstrap.com
[8] http://graphdb.ontotext.com
[9] http://d3js.org

[10] All datasets are available on SemEval websites:
alt.qcri.org/semeval2014/task4/
alt.qcri.org/semeval2015/task12
alt.qcri.org/semeval2016/task5

−	reception	0.02	4

Opinion name	Polarity	Positive Count	Negative Count	Inferred	General Inquirer	MPQA	SenticNet
great	0.49	1	0	0.07	0.50	0.75	0.86
good	0.20	2	0	0.03	0.50	0.25	0.88
better	-0.03	1	0	-0.01	0.50	0.25	0.13
terrible	-0.61	1	0	-0.10	-0.50	-0.75	-0.90

Fig. 8 The figure shows the summary of sentiment information about each opinion word provided to users

The results observed for *C1* are reported in Effectiveness on Aspect Detection. While, the results observer for *C2* are presented in Accuracy on Computing Aspect Polarity. The precision, recall, and f-measure metrics are reported for the *SB2014* and *SB2015* datasets. While, for the *SB2016* dataset we presented only the obtained f-measure because it was the only metric that were published by the organizers.

Effectiveness on Aspect Detection

Figures 9, 10, 11, 12, and 13 provide the results on the *SB2014*, *SB2015*, and *SB2016* datasets by the systems participated to the SemEval evaluation campaign and by the ReUS platform. For each system, we report a column representing precision, recall, and f1-score. When a column is not present, it means that particular system did not submit results for that dataset or domain.

We mentioned early that the ReUS platform differs from the other participant systems because it implements an unsupervised strategy instead of supervised ones. We can appreciated how ReUS effectiveness is in-line with the performance obtained by other systems. Hence, we may state that the current version of the platform is suitable for being used in real-world contexts and that, by considering the unsupervised nature of the integrated strategy, the addition of new domains would not affect the overall effectiveness.

By analyzing the results obtained on the two domains (i.e., *Laptop* and *Restaurant*), the performance of ReUS significantly changed, while for the *Restaurant* domain, the gap between ReUS and the best system ranges from 2 to 6%; within the *Laptop*, ReUS always obtained the best F1-score on the three datasets.

We run an error analysis of obtained results. The confusion matrices of the results collected for both domains on all datasets are reported in Fig. 14. We can observe how the distribution of the error is different for the two domains. By considering the *Laptop* domain, the errors are equally distributed among false positive and false negative, while for the *Restaurant* domain, 87% of the errors are classified as false positive. This means that ReUS annotates as aspects too many elements of a sentence. This result is in-line with well-known issues of unsupervised algorithms that are inclined to report many false positive errors [95]. Part of the future work will be dedicated to cope with this issue.

Accuracy on Computing Aspect Polarity

The second evaluation concerns the capability of the ReUS platform in inferring the correct polarity with which an

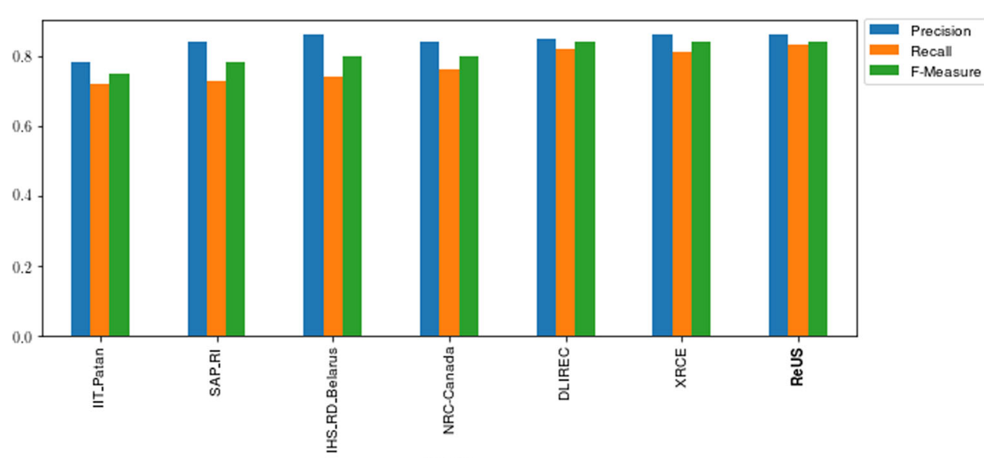

Fig. 9 The graph shows the aspect extraction performance obtained by the systems participated to the SemEval 2014 evaluation campaign and by ReUS on the *Restaurant* domain of the *SB2014* dataset. In order to ease the link between SemEval 2014 systems and the related papers, we report on the *X* axis their acronym

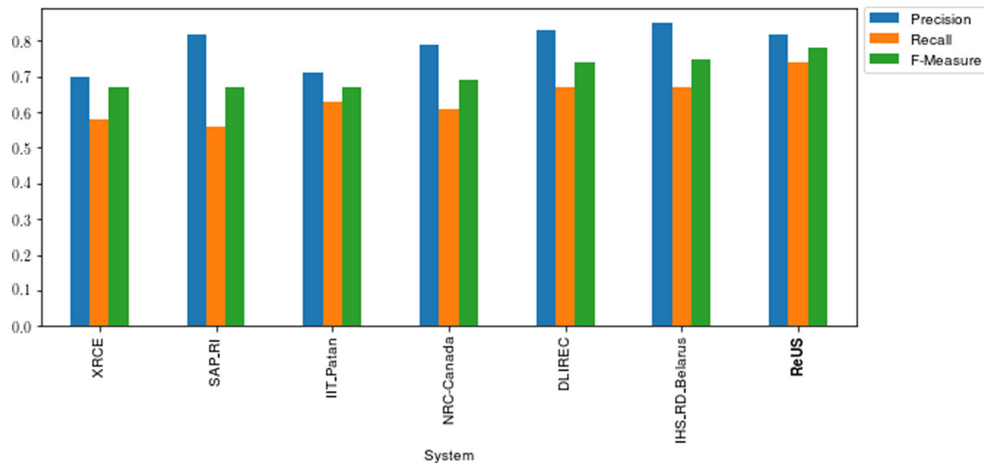

Fig. 10 The graph shows the aspect extraction performance obtained by the systems participated to the SemEval 2014 evaluation campaign and by ReUS on the *Laptop* domain of the *SB2014* dataset. In order to ease the link between SemEval 2014 systems and the related papers, we report on the X axis their acronym

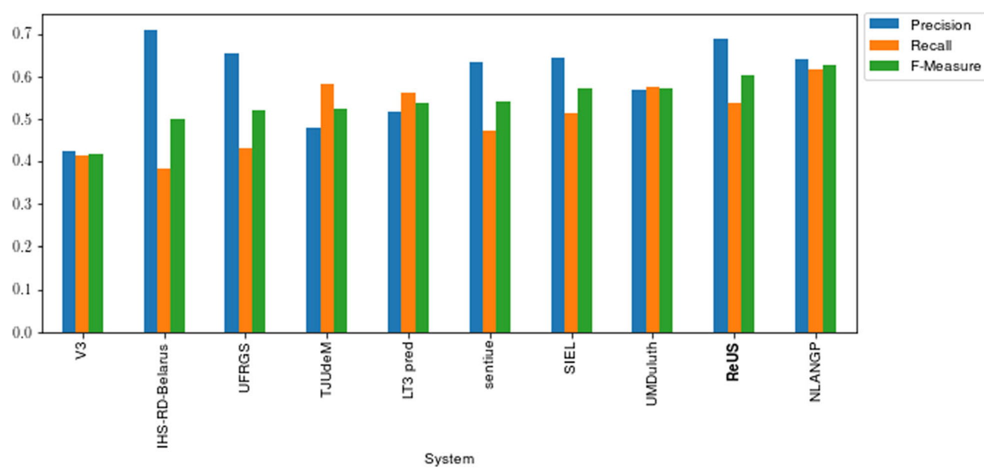

Fig. 11 The graph shows the aspect extraction performance obtained by the systems participated to the SemEval 2015 evaluation campaign and by ReUS on the *Restaurant* domain of the *SB2015* dataset. In order to ease the link between SemEval 2015 systems and the related papers, we report on the *X* axis their acronym

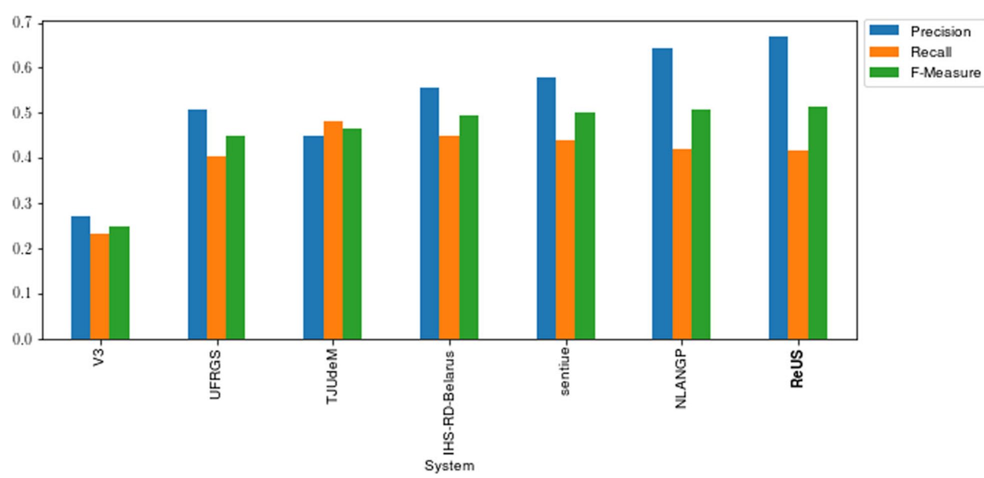

Fig. 12 The graph shows the aspect extraction performance obtained by the systems participated to the SemEval 2015 evaluation campaign and by ReUS on the *Laptop* domain of the *SB2015* dataset. In order to ease the link between SemEval 2015 systems and the related papers, we report on the *X* axis their acronym

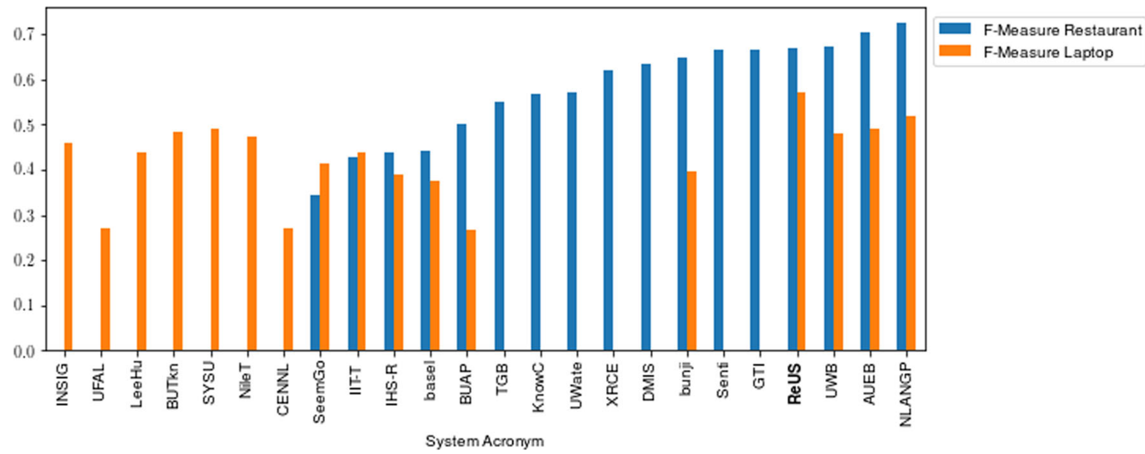

Fig. 13 The graph shows the aspect extraction performance obtained by the systems participated to the SemEval 2016 evaluation campaign and by ReUS on the *SB2016* dataset. In order to ease the link between SemEval 2016 systems and the related papers, we report on the *X* axis their acronym

aspect is mentioned into a text. This evaluation is performed on all the *SB2014*, *SB2015*, and *SB2016* datasets. The polarity of an opinion word is computed by averaging the polarity values extracted from the three resources described in Sentiment Lexicons. Figures 15, 16, and 17 show the accuracies obtained by each system participated to the SemEval evaluation campaigns and by the ReUS platform. Each column shows the accuracy obtained on each domain. When a column is not present, it means that particular system did not submit results for that dataset or domain.

We can notice how also in this case the behavior of the ReUS platform in terms of accuracy is similar to the one observed for the previous task. By considering the *Laptop* domain, the ReUS platform obtained the best performance on the three datasets, while for the other two domains (i.e., *Restaurant* and *Hotel*) there is a gap of 1 to 2% with respect to the best systems. Event if the difference is not statistically significant, in these contexts, the use of supervised approaches still outperform the proposed unsupervised one. Also for the polarity computation task, we performed an error analysis which results are reported in Fig. 18. Here, we can appreciate that there are no significant facts to highlight; indeed, errors are equally distributed among false positive and false negative. Further, we analyzed a subset of test documents for trying to understand why on the *Laptop* domain our algorithm performs better with respect to the other two. We discovered that, generally, reviews belonging to the *Restaurant* and *Hotel* domains are longer and more complex with respect to the *Laptop* ones. This insight will be a starting point for future analysis aiming to still improve the overall effectiveness of the ReUS platform.

User-based Evaluation

Besides the quantitative evaluation presented above, we performed a qualitative one consisted in filling a questionnaire by the group of 57 users invited to use our platform for analyzing the reviews stored in our repository. We reported below a discussion on the two parts of the platform that were analyzed by the users: (i) how the system is efficient in managing data streams and (ii) how the functionalities available within the user interface were understandable.

Platform Efficiency The efficiency of the overall architecture has been measured within a closed scenario where our repository stored around 1,200,000 reviews. By connecting the extractor to the Amazon data stream, we noticed that the scalability of the current version is still limited. In particular, we observed that the main issue was the co-reference analysis of the text that is performed for detecting all instances of an aspect contained within a document and the associated opinion. We tried to solve this problem by parallelizing this activity but the necessity of taking into account the temporal attribute of documents resulted in detrimental effects of the overall effectiveness. For this reason, we planned to develop a scheduler integrating a conflict detector enabling an effective parallelization of the documents to analyze in order to avoid the risk of losing information.

Usability of the Platform We collected feedback concerning the usability of the functionalities that we implemented into the user interface. By analyzing the comments provided within the questionnaire, we noticed that two main issues were highlighted by the users.

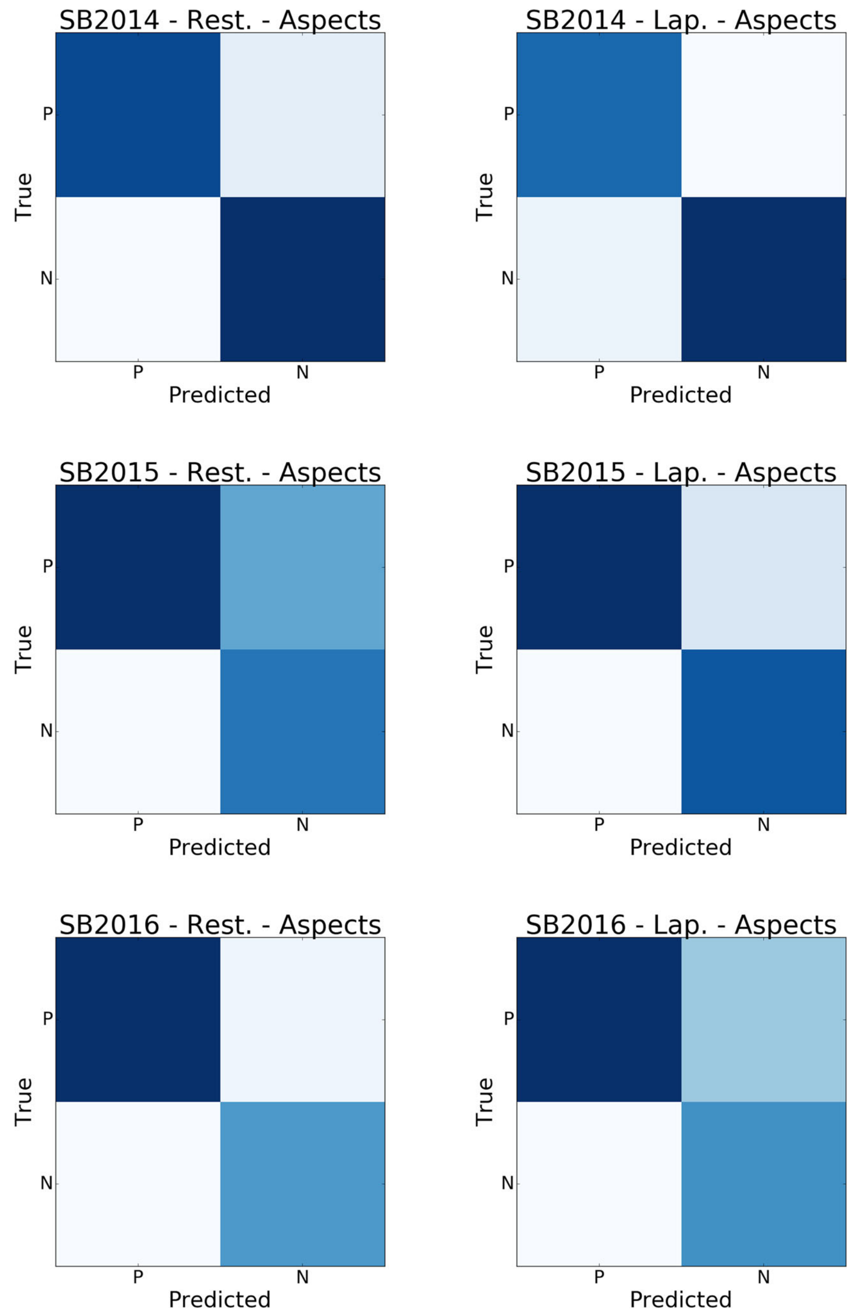

Fig. 14 Confusion matrices showing the error distribution for the aspect extraction task

- Contextual information into the aspect visualization: current version of the platform overlooked the fact that users with different levels of expertise can use the platform. Indeed, based on the their backgrounds, users would require functionalities focusing more on showing, graphically, which are the most important, or interesting, aspects, rather than to visualize further detailed information about them. Examples of such information are the number of users that supported or that rejected an opinion or the polarities associated with each aspect. We planned to include this functionality in future platform releases.
- Visualize how the polarity of each aspect evolves through time: the second consideration provided by users concerns the necessity of having a functionality allowing the observation of aspect polarities *evolve* through time. This feature would be definitely helpful for the decision-making process. For example, it would be used for analyzing how product's changes have been perceived by customers.

The feedback collected from the questionnaires will trigger future actions for evolving the structure of the overall platform. By improving the scalability of document analysis

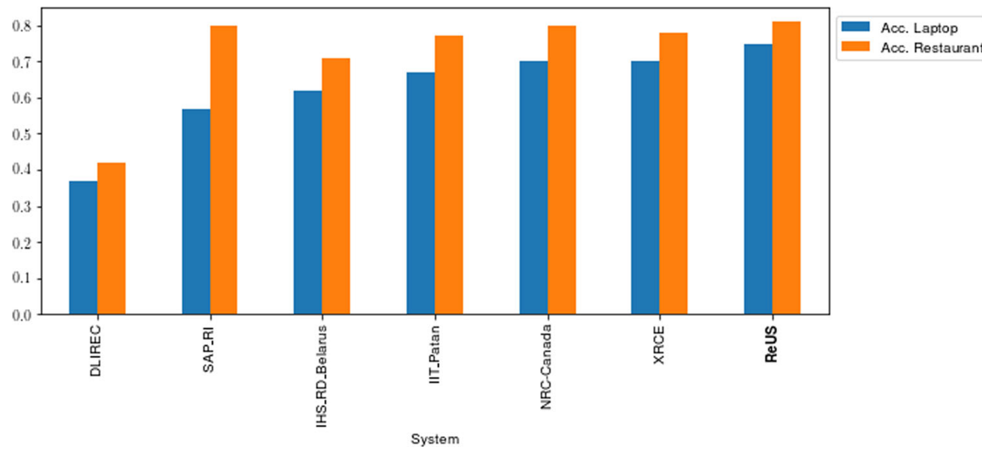

Fig. 15 The graph shows the polarity computation performance obtained by the systems participated to the SemEval 2014 evaluation campaign and by ReUS on the *SB2014* dataset. In order to ease the link between SemEval 2014 systems and the related papers, we report on the *X* axis their acronym

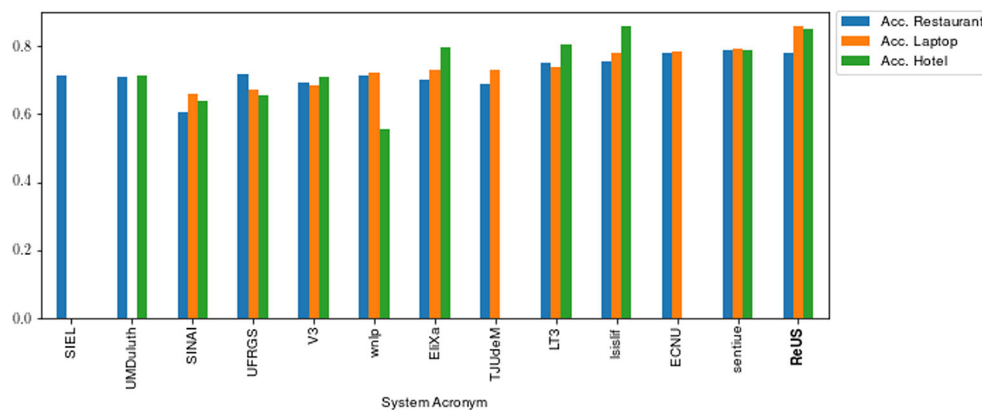

Fig. 16 The graph shows the polarity computation performance obtained by the systems participated to the SemEval 2015 evaluation campaign and by ReUS on the *SB2015* dataset. In order to ease the link between SemEval 2015 systems and the related papers, we report on the *X* axis their acronym

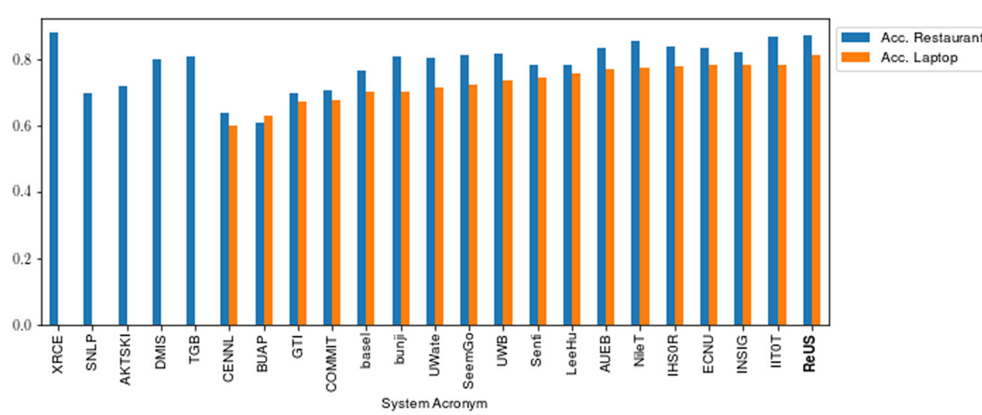

Fig. 17 The graph shows the polarity computation performance obtained by the systems participated to the SemEval 2016 evaluation campaign and by ReUS on the *SB2016* dataset. In order to ease the link between SemEval 2015 systems and the related papers, we report on the *X* axis their acronym

Fig. 18 Confusion matrices showing the error distribution for the polarity computation task

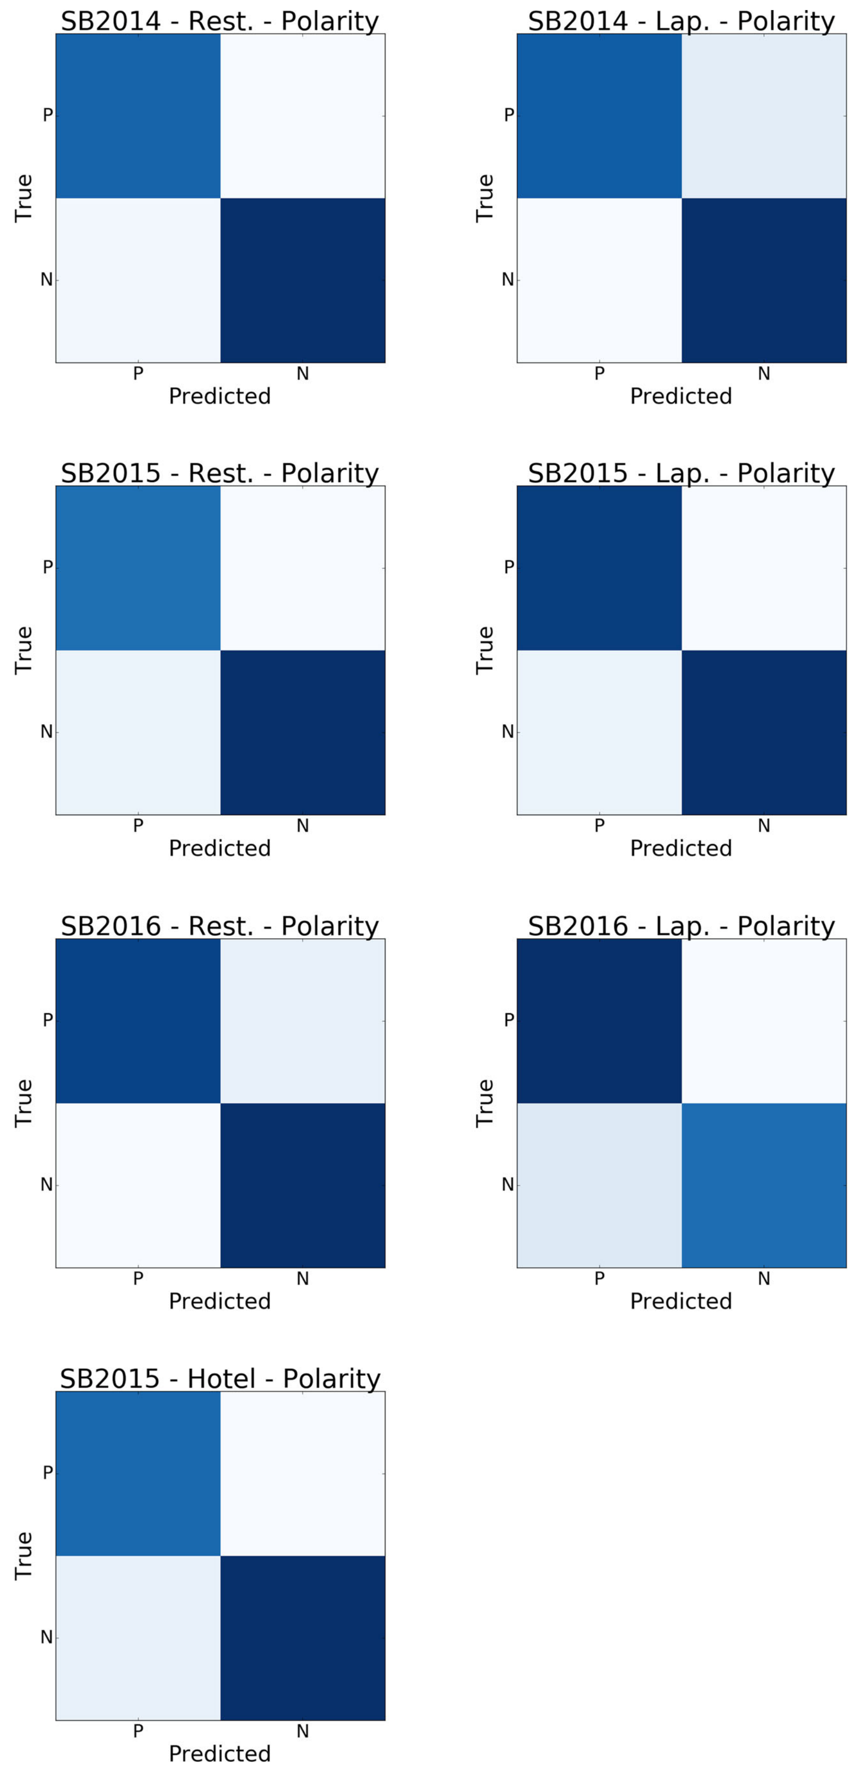

and by adding facilities for mitigating the effort for data analysis will make the system suitable for being delivered into a real-world scenario.

Conclusion and Future Work

In this paper, we presented a real-world system able to monitor streams of opinion-based documents and to extract the mentioned aspects and the polarities associated with them. The proposed architecture implemented an unsupervised strategy for extracting aspects and it integrated multiple knowledge bases for inferring associated polarities. The results confirmed the suitability of both the proposed OpenIE strategy and of the implemented architecture.

We show that in spite of the adoption of unsupervised strategies, the effectiveness of the proposed platform is in-line with the baselines that all implemented supervised solutions. An analysis about possible improvements of the platform focused on the implementation of further semantic-based strategies for improving the natural language parsing component (Fig. 3). An example might be the implementation of a semantic-distance checker able to discover and discard aspects that are not related to the entity referred by the document. Here, the challenge is implementing a solution avoiding the development of learning components. Then, a further analysis could be performed by relaxing the rules we used for tagging a term as aspect or not.

A second evidence we report in the paper is the effectiveness on inferring the polarity values associated with each aspect. The results demonstrated the viability of approaches based on the combination of several sentiment resources. Part of the future work will focus on defining a semantic representation able to easily merge and represent available sentiment resources in order to support the reasoning task.

From the user perspective, the tool we developed for performing the real-time analysis was overall judged interesting and useful thanks to the high number of available functionalities. The feature better perceived by the users was the tree-based view (Fig. 2) due to its capability of easily re-ranked features based on the associated polarity values.

The current version of the platform will be expanded in several directions. One of the techniques we want to implement is an embedded-based clustering strategies for grouping similar words. This way, we should be able to group text chunks having the same conceptual meaning and, at the same time, to preserve the aspect extraction capability.

Finally, the multi-domain research field will be explored for two reasons: First, it will be of interest to carry out refinements on the sentiment resources by taking into account the domain information. This action should improve the overall effectiveness on inferring aspects' polarities; second, to include uncertainty in representing information. Every information we analyze represents only a part of the world and the remaining is totally unknown to a system. Thus, the adoption of a dynamic fuzzy representation of polarities would be able to consider this aspect. The challenge here will be to transform this information in something tangible for the final users.

Compliance with Ethical Standards

Conflict of interests The authors declare that they have no conflict of interest.

Informed Consent Informed consent was not required as no human or animals were involved.

Human and Animal Rights This article does not contain any studies with human or animal subjects performed by any of the authors.

References

1. Poria S, Cambria E, Bajpai R, Hussain A. A review of affective computing from unimodal analysis to multimodal fusion. Information Fusion. 2017;37:98–125.
2. Hazarika D, Poria S, Zadeh A, Cambria E, Morency Louis-Philippe, Zimmermann R. Conversational memory network for emotion recognition in dyadic dialogue videos. In: NAACL; 2018. p. 2122–2132.
3. Chaturvedi I, Cambria E, Welsch R, Herrera F. Distinguishing between facts and opinions for sentiment analysis: survey and challenges. Information Fusion. 2018;44:65–77.
4. Cambria E, Olsher D, Kwok K. Sentic activation: a two-level affective common sense reasoning framework. In: AAAI. Toronto; 2012. p. 186–192.
5. Cambria E. An introduction to concept-level sentiment analysis. Advances in soft computing and its applications, volume 8266 of lecture notes in computer science. In: Castro F, Gelbukh A, and González M, editors. Berlin: Springer; 2013. p. 478–483.
6. Lo SL, Cambria E, Chiong R, Cornforth D. Multilingual sentiment analysis: from formal to informal and scarce resource languages. Artif Intell Rev. 2017;48(4):499–527.
7. Peng H, Ma Y, Li Y, Cambria E. Learning multi-grained aspect target sequence for chinese sentiment analysis. Knowl-Based Syst. 2018;148:167–176.
8. Bandhakavi A, Wiratunga N, Massie S, Deepak P. Lexicon generation for emotion analysis of text. IEEE Intell Syst. 2017;32(1):102–108.
9. Dragoni M, Poria S, Cambria E. OntoSenticNet: a commonsense ontology for sentiment analysis. IEEE Intell Syst. 2018;33(3):77–85.
10. Oneto L, Bisio F, Cambria E, Anguita D. Statistical learning theory and ELM for big social data analysis. IEEE Comput Intell Mag. 2016;11(3):45–55.
11. Hussain A, Cambria E. Semi-supervised learning for big social data analysis. Neurocomputing. 2018;275:1662–1673.
12. Li Y, Pan Q, Yang T, Wang S, Tang JL, Cambria E. Learning word representations for sentiment analysis. Cogn Comput. 2017;9(6):843–851.

13. Young T, Hazarika D, Poria S, Cambria E. Recent trends in deep learning based natural language processing. IEEE Comput Intell Mag. 2018;13(3):55–75.
14. Li Y, Pan Q, Wang S, Yang T, Cambria E. A generative model for category text generation. Inf Sci. 2018;450:301–315.
15. Cambria E, Poria S, Gelbukh A, Thelwall M. Sentiment analysis is a big suitcase. IEEE Intell Syst. 2017;32(6):74–80.
16. Xia Y, Cambria E, Hussain A, Zhao H. Word polarity disambiguation using bayesian model and opinion-level features. Cogn Comput. 2015;7(3):369–380.
17. Chaturvedi I, Ragusa E, Gastaldo P, Zunino R, Cambria E. Bayesian network based extreme learning machine for subjectivity detection. J Franklin Inst. 2018;355(4):1780–1797.
18. Majumder N, Poria S, Gelbukh A, Cambria E. Deep learning-based document modeling for personality detection from text. IEEE Intell Syst. 2017;32(2):74–79.
19. Satapathy R, Guerreiro C, Chaturvedi I, Cambria E. Phonetic-based microtext normalization for twitter sentiment analysis. In: ICDM; 2017. p. 407–413.
20. Rajagopal D, Cambria E, Olsher D, Kwok K. A graph-based approach to commonsense concept extraction and semantic similarity detection. In: WWW; 2013. p. 565–570.
21. Zhong X, Sun A, Cambria E. Time expression analysis and recognition using syntactic token types and general heuristic rules. In: ACL; 2017. p. 420–429.
22. Ma Y, Peng H, Cambria E. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive LSTM. In: AAAI; 2018. p. 5876–5883.
23. Xing F, Cambria E, Welsch R. Natural language based financial forecasting: a survey. Artif Intell Rev. 2018;50(1):49–73.
24. Ebrahimi M, Hossein A, Sheth A. Challenges of sentiment analysis for dynamic events. IEEE Intell Syst. 2017;32(5):70–75.
25. Cambria E, Hussain A, Durrani T, Havasi C, Eckl C, Munro J. Sentic computing for patient centered applications. In: IEEE ICSP; 2010. p. 1279–1282.
26. Valdivia A, Luzon V, Herrera F. Sentiment analysis in tripadvisor. IEEE Intell Syst. 2017;32(4):72–77.
27. Mihalcea R, Garimella A. What men say what women hear: finding gender-specific meaning shades. IEEE Intell Syst. 2016;31(4):62–67.
28. Cavallari S, Zheng V, Cai H, Chang K, Cambria E. Learning community embedding with community detection and node embedding on graphs. In: CIKM; 2017. p. 377–386.
29. Chi Xu, Cambria E, Tan PS. Adaptive two-stage feature selection for sentiment classification. In: IEEE SMC; 2017. p. 1238–1243.
30. Zadeh A, Liang PP, Poria S, Vij P, Cambria E, Morency Louis-Philippe. Multi-attention recurrent network for human communication comprehension. In: AAAI; 2018. p. 5642–5649.
31. Young T, Cambria E, Chaturvedi I, Zhou H, Biswas S, Huang M. Augmenting end-to-end dialogue systems with commonsense knowledge. In: AAAI; 2018. p. 4970–4977.
32. Minqing Hu, Liu B. Mining and summarizing customer reviews. Proceedings of the tenth ACM SIGKDD international conference on knowledge discovery and data mining, Seattle, Washington, USA, August 22-25, 2004. In: Kim W, Kohavi R, Gehrke J, and DuMouchel W, editors. ACM; 2004. p. 168–177.
33. Poria S, Chaturvedi I, Cambria E, Bisio F. Sentic LDA: improving on LDA with semantic similarity for aspect-based sentiment analysis. In: IJCNN; 2016. p. 4465–4473.
34. Poria S, Cambria E, Gelbukh A. Aspect extraction for opinion mining with a deep convolutional neural network. Knowl-Based Syst. 2016;108:42–49.
35. Liu B, Zhang L. A survey of opinion mining and sentiment analysis. Mining text data. In: Aggarwal CC and Zhai CX, editors. Springer; 2012. p. 415-463.
36. Cambria E. Affective computing and sentiment analysis. IEEE Intell Syst. 2016;31(2):102–107.
37. Bo P, Lee L, Vaithyanathan S. Thumbs up? sentiment classification using machine learning techniques. In: Proceedings of EMNLP. Philadelphia; 2002. p. 79–86. Association for Computational Linguistics.
38. Pang B, Lee L. A sentimental education: sentiment analysis using subjectivity summarization based on minimum cuts. In: ACL; 2004. p. 271–278.
39. Go A, Bhayani R, Huang L. Twitter sentiment classification using distant supervision. CS224n Project report, Standford University. 2009.
40. Barbosa L, Feng J. Robust sentiment detection on twitter from biased and noisy data. In: COLING (Posters); 2010. p. 36–44.
41. Dragoni M. Shellfbk: an information retrieval-based system for multi-domain sentiment analysis. In: Proceedings of the 9th international workshop on semantic evaluation, SemEval ’2015, pp 502–509, Denver, Colorado; 2015. Association for Computational Linguistics.
42. Petrucci G, Dragoni M. An information retrieval-based system for multi-domain sentiment analysis. Semantic web evaluation challenges - second semwebeval challenge at ESWC 2015, portorož, Slovenia, May 31 - June 4 (2015), Revised Selected Papers, volume 548 of Communications in Computer and Information Science. In: Gandon F, Cabrio E, Stankovic M, and Zimmermann A, editors. Springer; 2015. p. 234–243.
43. Riloff E, Patwardhan S, Wiebe J. Feature subsumption for opinion analysis. In: EMNLP; 2006. p. 440–448.
44. Wilson T, Wiebe J, Hwa R. Recognizing strong and weak opinion clauses. Comput Intell. 2006;22(2):73–99.
45. Qi Su, Xinying Xu, Guo H, Guo Z, Xian Wu, Zhang X, Swen B, Zhong Su. Hidden sentiment association in chinese web opinion mining. In: WWW; 2008. p. 959–968.
46. Jin W, Ho HH, Srihari RK. Opinionminer: a novel machine learning system for web opinion mining and extraction. In: KDD; 2009. p. 1195–1204.
47. Jakob N, Gurevych I. Extracting opinion targets in a single and cross-domain setting with conditional random fields. In: Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, EMNLP 2010, 9-11 October 2010, MIT Stata Center, Massachusetts, USA, A meeting of SIGDAT, a Special Interest Group of the ACL, pp 1035–1045. ACL; 2010.
48. Dragoni M, Tettamanzi AndreaGB, Pereira CC. Propagating and aggregating fuzzy polarities for concept-level sentiment analysis. Cogn Comput. 2015;7(2):186–197.
49. Dragoni M, Tettamanzi AGB, Pereira CéliadaC. A fuzzy system for concept-level sentiment analysis. In: Valentina presutti, milan stankovic, erik cambria, iván cantador, angelo di iorio, tommaso di noia, christoph lange, diego reforgiato recupero, and anna tordai, editors, semantic web evaluation challenge - semwebeval 2014 at ESWC 2014, anissaras, crete, greece, may 25-29 (2014), revised selected papers, volume 475 of communications in computer and information science, pp 21–27. Springer; 2014.
50. da Pereira CC, Dragoni M, Pasi G. A prioritized and aggregation operator for multidimensional relevance assessment. In: Roberto serra and rita cucchiara, editors, AI*IA 2009: Emergent perspectives in artificial intelligence, XIth international conference of the italian association for artificial intelligence, reggio emilia, italy, december 9-12, 2009, proceedings, volume 5883 of lecture notes in computer science, pp 72–81. Springer; 2009.
51. Aprosio AP, Corcoglioniti F, Dragoni M, Rospocher M. Supervised opinion frames detection with RAID. In: Fabien gandon, elena cabrio, milan stankovic, and antoine zimmermann, editors, semantic web evaluation challenges - second semwebeval

challenge at ESWC 2015, portorož, Slovenia, May 31 - June 4 (2015), Revised Selected Papers, volume 548 of Communications in Computer and Information Science, pp 251–263. Springer; 2015.
52. Cambria E, Hussain A. Sentic computing: a common-sense-based framework for concept-level sentiment analysis. Cham Switzerland: Springer; 2015.
53. Wang QF, Cambria E, Liu CL, Hussain A. Common sense knowledge for handwritten chinese recognition. Cogn Comput. 2013;5(2):234–242.
54. Cambria E, Hussain A. Sentic album: content-, concept-, and context-based online personal photo management system. Cogn Comput. 2012;4(4):477–496.
55. Gangemi A, Presutti V, Recupero DR. Frame-based detection of opinion holders and topics: a model and a tool. IEEE Comp Int Mag. 2014;9(1):20–30.
56. Recupero DR, Presutti V, Consoli S, Gangemi A, Sentilo AndreaGiovanniNuzzolese. Frame-based sentiment analysis. Cogn Comput. 2015;7(2):211–225.
57. Thelwall M, Buckley K, Paltoglou G, Di C, Kappas A. Sentiment in short strength detection informal text. JASIST. 2010;61(12):2544–2558.
58. Fan Teng-Kai, Chang C-H. Sentiment-oriented contextual advertising. Inf Knowl Syst. 2010;23(3):321–344.
59. Dragoni M. A three-phase approach for exploiting opinion mining in computational advertising. IEEE Intell Syst. 2017;32(3):21–27.
60. Sklar M, Concepcion KJ. Timely tip selection for foursquare recommendations. Poster proceedings of the 8th ACM conference on recommender systems, RecSys 2014, Foster City, Silicon Valley, CA, USA, October 6-10 (2014), volume 1247 of CEUR Workshop Proceedings. CEUR-WS.org. In: Chen L and Mahmud J, editors; 2014.
61. Choi Y, Cardie C. Hierarchical sequential learning for extracting opinions and their attributes. I. In: Proceedings of the ACL 2010 conference short papers, ACLShort'10, pp 269–274, Stroudsburg, PA. USA; 2010. Association for Computational Linguistics.
62. Zhang M, Zhang Y, Vo Duy-Tin. Neural networks for open domain targeted sentiment. Proceedings of the 2015 conference on empirical methods in natural language processing, EMNLP 2015, Lisbon, Portugal, September 17-21 (2015), pp 612–621. The Association for Computational Linguistics. In: Màrquez L, Callison-Burch C, Su J, Pighin D, and Marton Y, editors; 2015.
63. Mitchell M, Aguilar J, Wilson T, Durme BV. Open domain targeted sentiment. In: Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL, pp 1643–1654. ACL; 2013.
64. Liu B, Minqing Hu, Cheng J. Opinion observer: analyzing and comparing opinions on the web. In: WWW; 2005. p. 342–351.
65. Mei Q, Ling Xu, Wondra M, Hang Su, Zhai CX. Topic sentiment mixture: modeling facets and opinions in weblogs. In: Proceedings of the 16th international conference on World Wide Web, WWW '07, pp 171–180, New York, NY, USA. ACM; 2007.
66. Titov I, McDonald R. A joint model of text and aspect ratings for sentiment summarization. In: PROC. ACL-08: HLT; 2008. p. 308–316.
67. Li F, Huang M, Zhu X. Sentiment analysis with global topics and local dependency. In: Proceedings of the twenty-fourth AAAI conference on artificial intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010; 2010.
68. Mukherjee A, Liu CX. Aspect extraction through semi-supervised modeling. In: Proceedings of the 50th annual meeting of the association for computational linguistics: long papers - Volume 1, ACL'12, pp 339–348, Stroudsburg, PA, USA. Association for Computational Linguistics; 2012.
69. Dragoni M, Azzini A, Tettamanzi A. A novel similarity-based crossover for artificial neural network evolution. Parallel problem solving from nature - PPSN XI, 11th International Conference, Krakòw, Poland, September 11-15, 2010, Proceedings, Part I, volume 6238 of Lecture Notes in Computer Science, pp 344–353. Springer. In: Schaefer R, Cotta C, Kolodziej J, and Rudolph G, editors; 2010.
70. Yuanbin Wu, Qi Z, Huang X, Wu L. Phrase dependency parsing for opinion mining. In: Proceedings of the 2009 conference on empirical methods in natural language processing: volume 3 - Volume 3, EMNLP '09, pp 1533–1541, Stroudsburg, PA, USA. Association for Computational Linguistics; 2009.
71. Banko M, Cafarella M, Soderland S, Broadhead M, Etzioni O. Open information extraction for the web. In: Proceedings of the international joint conference on artificial intelligence, IJCAI '07; 2007.
72. Yates A, Cafarella M, Banko M, Etzioni O, Broadhead M, Textrunner StephenSoderland. Open information extraction on the web. In: Proceedings of human language technologies: the annual conference of the North American chapter of the association for computational linguistics: demonstrations, NAACL-demonstrations '07, pp 25–26, Stroudsburg, PA, USA. Association for Computational Linguistics; 2007.
73. Wu F, Weld DS. Open information extraction using wikipedia. In: Proceedings of the 48th annual meeting of the association for computational linguistics, ACL'10, pp 118–127, Stroudsburg, PA, USA. Association for Computational Linguistics; 2010.
74. Fader A, Soderland S, Etzioni O. Identifying relations for open information extraction. In: Proceedings of the conference on empirical methods in natural language processing, EMNLP'11, pp 1535–1545, Stroudsburg, PA, USA. Association for Computational Linguistics; 2011.
75. Akbik A, Löser A. Kraken: N-ary facts in open information extraction. In: Proceedings of the joint workshop on automatic knowledge base construction and web-scale knowledge extraction, AKBC-WEKEX'12, pp 52–56, Stroudsburg, PA, USA. Association for Computational Linguistics; 2012.
76. Mausam MS, Bart R, Soderland S, Etzioni O. Open language learning for information extraction. In: Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning, EMNLP-coNLL'12, pp 523–534, Stroudsburg, PA, USA. Association for Computational Linguistics; 2012.
77. Corro LD, Gemulla R. Clausie: clause-based open information extraction. In: Proceedings of the 22Nd international conference on World Wide Web, WWW'13, pp 355–366, New York, NY, USA. ACM; 2013.
78. Bast H, Haussmann E. Open information extraction via contextual sentence decomposition. In: Proceedings of the 2013 IEEE 7th international conference on semantic computing, ICSC'13. IEEE; 2013.
79. Zhila A, Gelbukh A. Open information extraction for spanish language based on syntactic constraints. In: ACL (Student research workshop); 2014. p. 78–85.
80. Wang M, Li L, Huang F. Semi-supervised chinese open entity relation extraction. In: Proceedings of the 3rd IEEE international conference on cloud computing and intelligence systems. IEEE; 2014.
81. Falke T, Stanovsky G, Gurevych I, Dagan I. Porting an open information extraction system from english to german. In: EMNLP; 2016. p. 892–898.
82. Gamallo P, Garcia M, Fern'andez-Lanza S. Dependency-based open information extraction. In: EACL; 2012.

83. Gamallo P, Garcia M. Multilingual open information extraction. Cham: Springer; 2015, pp. 711–722.
84. Fellbaum C. WordNet: an electronic lexical database. Cambridge: MIT Press; 1998.
85. Cambria E, Poria S, Hazarika D, Kwok K. SenticNet 5: discovering conceptual primitives for sentiment analysis by means of context embeddings. AAAI; 2018. p. 1795–1802.
86. Stone P, Dunphy DC, Marshall S. The general inquirer: a computer approach to content analysis. Oxford: MIT Press; 1966.
87. Deng L, Wiebe J. MPQA 3.0: an entity/event-level sentiment corpus. In: Rada mihalcea, joyce yue chai, and anoop sarkar, editors, NAACL HLT 2015, the 2015 conference of the north american chapter of the association for computational linguistics: human language technologies, denver, colorado, USA, May 31 - June 5 (2015), pp 1323–1328. The Association for Computational Linguistics; 2015.
88. Poria S, Gelbukh A, Cambria E, Yang P, Hussain A, Durrani T. Merging senticnet and wordnet-affect emotion lists for sentiment analysis. In: Signal processing (ICSP) (2012) IEEE 11th international conference on, volume 2, pp 1251–1255. IEEE; 2012.
89. Quirk R, Greenbaum S, Leech G, Svartvik J, Crystal D. A comprehensive grammar of the English language, volume 397. Cambridge: Cambridge Univ Press; 1985.
90. Manning CD, Surdeanu M, Bauer J, Finkel J, Bethard SJ, McClosky D. The stanford coreNLP natural language processing toolkit. In: Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations, pp 55–60, Baltimore, Maryland. Association for Computational Linguistics; June 2014.
91. Toutanova CMK, Klein D, Singer Y. Feature-rich part-of-speech tagging with a cyclic dependency network. In: Proceedings of HLT-NAACL 2003, pp 252–259; 2003.
92. Clark K, Manning CD. Entity-centric coreference resolution with model stacking. In: Association for computational linguistics (ACL); 2015.
93. Erhard W. Hinrichs and Dan Roth, editors. Accurate Unlexicalized Parsing. 2003.
94. Chen D, Manning CD. A fast and accurate dependency parser using neural networks. In: Empirical methods in natural language processing (EMNLP); 2014.
95. Liu Q, Gao Z, Liu B, Zhang Y. Automated rule selection for aspect extraction in opinion mining. Proceedings of the twenty-fourth international joint conference on artificial intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31 (2015), pp 1291–1297. AAAI Press. In: Yang Q and Wooldridge M, editors; 2015.