

Extracting Linguistic Features From Opinion Data Streams For Multi-Domain Sentiment Analysis

Mauro Dragoni

Fondazione Bruno Kessler, Trento, Italy
dragoni@fbk.eu

Abstract. The approach described in this paper explores the use of semantic structured representation of sentences extracted from texts for multi-domain sentiment analysis purposes. The presented algorithm is built upon a domain-based supervised approach using index-like structured for representing information extracted from text. The algorithm extracts dependency parse relationships from the sentences containing in a training set. Then, such relationships are aggregated in a semantic structured together with either polarity and domain information. Such information is exploited in order to have a more fine-grained representation of the learned sentiment information. When the polarity of a new text has to be computed, such a text is converted in the same semantic representation that is used (i) for detecting the domain to which the text belongs to, and then (ii), once the domain is assigned to the text, the polarity is extracted from the index-like structure. First experiments performed by using the Blitzer dataset for training the system demonstrated the feasibility of the proposed approach.

1 Introduction

Sentiment analysis is a natural language processing task whose aim is to classify documents according to the opinion (polarity) they express on a given subject [1]. Generally speaking, sentiment analysis aims at determining the attitude of a speaker or a writer with respect to a topic or the overall tonality of a document. This task has created a considerable interest due to its wide applications. In recent years, the exponential increase of the Web for exchanging public opinions about events, facts, products, etc., has led to an extensive usage of sentiment analysis approaches, especially for marketing purposes.

By formalizing the sentiment analysis problem, a “sentiment” or “opinion” has been defined by [2] as a quintuple:

$$\langle o_j, f_{jk}, so_{ijkl}, h_i, t_l \rangle, \quad (1)$$

where o_j is a target object, f_{jk} is a feature of the object o_j , so_{ijkl} is the sentiment value of the opinion of the opinion holder h_i on feature f_{jk} of object o_j at time t_l . The value of so_{ijkl} can be positive (by denoting a state of happiness, bliss, or satisfaction), negative (by denoting a state of sorrow, dejection, or disappointment), or neutral (it is not possible to denote any particular sentiment), or a more granular rating. The term h_i encodes the opinion holder, and t_l is the time when the opinion is expressed.

Such an analysis, may be *document-based*, where the positive, negative, or neutral sentiment is assigned to the entire document content; or *sentence-based* where individual sentences are analyzed separately and classified according to the different polarity values. In the latter case, it is often desirable to find with a high precision the entity attributes towards which the detected sentiment is directed. Based on the scenario in which the opinion is needed, the use of a document-based analysis is preferred with respect to a sentence-based one, and vice versa. In this work, we want to extract the general opinion of an entire document; therefore, our approach relies on a document-based analysis.

A further aspect that it is important to take into account is that, in the classic sentiment analysis problem, the polarity of each document term is considered independently by the domain which the document belongs to. We illustrate the intuition behind domain specific term polarity by considering the following example:

1. The sideboard is **small** and it is not able to contain a lot of stuff.
2. The **small** dimensions of this decoder allow to move it easily.

In these two sentences the adjective “small” is used in two different domains. In the first sentence, we considered the Furnishings domain and, within it, the polarity of the adjective “small” is, for sure, “negative” because it highlights an issue of the described item. On the other hand, in the second sentence, where we considered the Electronics domain, the polarity of such an adjective may be considered “positive”. First attempts exploring how term polarity is conditioned by domain is presented in [3].

Unlike the approaches already discussed in the literature (presented in Section 2), we address the multi-domain sentiment analysis problem from a different perspective. Firstly, we extract semantic and linguistic relationships from document terms, and then, we aggregate them in a structured representation where domain information, and the related polarities, are preserved. Such a structured representation is stored in an index-like repository (from now simply referred as “index”). When the polarity of a new document has to be computed, its structured representation is built and, combined with domain information, it is used for querying the index in order to estimate the polarity of the whole document.

The rest of the work is structured as follows. Section 2 presents a survey on works about sentiment analysis. Section 3 described the proposed approach by explaining how texts are converted in a semantic structured representation, stored during the training phase, and exploited during the test one. Section 4 reports the comparison between the presented approach and three baselines. Finally, Section 5 concludes the paper.

2 Related Work

The topic of sentiment analysis has been studied extensively in the literature [2], where several techniques have been proposed and validated.

Machine learning techniques are the most common approaches used for addressing this problem, given that any existing supervised methods can be applied to sentiment classification. For instance, in [4], the authors compared the performance of Naive-Bayes, Maximum Entropy, and Support Vector Machines in sentiment analysis

on different features like considering only unigrams, bigrams, combination of both, incorporating parts of speech and position information or by taking only adjectives. Moreover, beside the use of standard machine learning method, researchers have also proposed several custom techniques specifically for sentiment classification, like the use of adapted score function based on the evaluation of positive or negative words in product reviews [5], as well as by defining weighting schemata for enhancing classification accuracy [6].

An obstacle to research in this direction is the need of labeled training data, whose preparation is a time-consuming activity. Therefore, in order to reduce the labeling effort, opinion words have been used for training procedures. In [7] and [8], the authors used opinion words to label portions of informative examples for training the classifiers. Opinion words have been exploited also for improving the accuracy of sentiment classification, as presented in [9], where a framework incorporating lexical knowledge in supervised learning to enhance accuracy has been proposed. Opinion words have been used also for unsupervised learning approaches like the one presented in [10].

Another research direction concerns the exploitation of discourse-analysis techniques. [11] discusses some discourse-based supervised and unsupervised approaches for opinion analysis; while in [12], the authors present an approach to identify discourse relations.

The approaches presented above are applied at the document-level[13,14,15,16], i.e., the polarity value is assigned to the entire document content. However, in some case, for improving the accuracy of the sentiment classification, a more fine-grained analysis of a document is needed. Hence, the sentiment classification of the single sentences, has to be performed. In the literature, we may find approaches ranging from the use of fuzzy logic [17,18,19] to the use of aggregation techniques [20] for computing the score aggregation of opinion words. In the case of sentence-level sentiment classification, two different sub-tasks have to be addressed: (i) to determine if the sentence is subjective or objective, and (ii) in the case that the sentence is subjective, to determine if the opinion expressed in the sentence is positive, negative, or neutral. The task of classifying a sentence as subjective or objective, called “subjectivity classification”, has been widely discussed in the literature [21,22,23] and systems implementing the capabilities of identifying opinion’s holder, target, and polarity have been presented [24]. Once subjective sentences are identified, the same methods as for sentiment classification may be applied. For example, in [25] the authors consider gradable adjectives for sentiment spotting; while in [26,27] the authors built models to identify some specific types of opinions.

In the last years, with the growth of product reviews, the use of sentiment analysis techniques was the perfect floor for validating them in marketing activities [28]. However, the issue of improving the ability of detecting the different opinions concerning the same product expressed in the same review became a challenging problem. Such a task has been faced by introducing “aspect” extraction approaches that were able to extract, from each sentence, which is the aspect the opinion refers to. In the literature, many approaches have been proposed: conditional random fields (CRF) [29], hidden Markov models (HMM) [30], sequential rule mining [31], dependency tree kernels [32], clustering [33], and genetic algorithms [34]. In [35], a method was proposed to extract both

opinion words and aspects simultaneously by exploiting some syntactic relations of opinion words and aspects.

A particular attention should be given also to the application of sentiment analysis in social networks [36]. More and more often, people use social networks for expressing their moods concerning their last purchase or, in general, about new products. Such a social network environment opened up new challenges due to the different ways people express their opinions, as described by [37] and [38], who mention “noisy data” as one of the biggest hurdles in analyzing social network texts.

One of the first studies on sentiment analysis on micro-blogging websites has been discussed in [39], where the authors present a distant supervision-based approach for sentiment classification.

At the same time, the social dimension of the Web opens up the opportunity to combine computer science and social sciences to better recognize, interpret, and process opinions and sentiments expressed over it. Such multi-disciplinary approach has been called *sentic computing* [40]. Application domains where sentic computing has already shown its potential are the cognitive-inspired classification of images [41], of texts in natural language, and of handwritten text [42].

Finally, an interesting recent research direction is domain adaptation, as it has been shown that sentiment classification is highly sensitive to the domain from which the training data is extracted. A classifier trained using opinionated documents from one domain often performs poorly when it is applied or tested on opinionated documents from another domain, as we demonstrated through the example presented in Section 1. The reason is that words and even language constructs used in different domains for expressing opinions can be quite different. To make matters worse, the same word in one domain may have positive connotations, but in another domain may have negative ones; therefore, domain adaptation is needed. In the literature, different approaches related to the Multi-Domain sentiment analysis have been proposed. Briefly, two main categories may be identified: (i) the transfer of learned classifiers across different domains [3,43,44], and (ii) the use of propagation of labels through graph structures [45,46,17,47].

All approaches presented above are based on the use of statistical techniques for building sentiment models. The exploitation of semantic information is not taken into account. In this work, we proposed a first version of a semantic-based approach preserving the semantic relationships between the terms of each sentence in order to exploit them either for building the model and for estimating document polarity. The proposed approach, falling into the multi-domain sentiment analysis category, instead of using pre-determined polarity information associated with terms, it learns them directly from domain-specific documents. Such documents are used for training the models used by the system.

3 The Approach

As introduced in Section 1, the proposed system is based on the implementation of an index-like approach, based on the use of structured representations of documents. Such representation is used for either preserving domain information associated with each document and for estimating the polarity of unclassified ones. Document polarity is

estimated through the computation of a Score Status Value [48] (SSV) representing the aggregation of the polarities estimated for each feature extracted from the document. In this section, the steps carried out for implementing our approach are presented.

3.1 Feature Extraction

The first task consists in the detection of the features that are exploited for building the sentiment model. The proposed approach has been designed upon two main desiderata:

1. The need of preserving and exploiting semantic relationships between document terms, requires to find a structured representation of information able to address this issue. In particular, we want to store linguistic information of each term together with its semantic relationships with the other ones;
2. The described approach addresses the problem of sentiment analysis in a multi-domain environment; therefore, each extracted feature has to enclose domain-specific information in order to exploit them during the estimation of document polarity.

Addressing the two pillars described above, requires to parse raw texts in order to extract significant linguistic and semantic information. The proposed solution for extracting the set of features is based on the use of a native natural language processing library, namely the Stanford NLP Core Toolkit [49].

For each document of the training set, we applied the Stanford parser for extracting the terms dependencies. Such dependencies are taken into account for preserving the semantic between terms in the structured representation used for representing document content.

As an example, let's consider the following sentence:

“I came here to reflect my happiness by fishing.”

By applying the Stanford parser, we obtain the following list of dependencies between terms:

```
nsubj(came-2, I-1)
nsubj(reflect-5, I-1)
root(ROOT-0, came-2)
advmod(came-2, here-3)
aux(reflect-5, to-4)
xcomp(came-2, reflect-5)
poss(happiness-7, my-6)
dobj(reflect-5, happiness-7)
prep_by(reflect-5, fishing-9)
```

Each dependency is composed by three elements: the name of the “relation” (R), the “governor” (G) that is the first term of the dependency, and the “dependent” (D) that is the second one. First of all, we removed from the dependencies list, ones containing a stop word¹ as governor or dependent element. Exceptions are made when one of

¹ The list of stop words used in this work is the one provided by Apache with the Lucene and Solr packages

the two terms contained in a dependency is an adjective. From the dependencies list presented above, the pruned list is the following:

```
poss(happiness-7, my-6)
dobj(reflect-5, happiness-7)
prep_by(reflect-5, fishing-9)
```

Then, for each dependency contained in the pruned list, we compile a set of pairs “field - value”. Each pair is a “feature” associated with the dependency extracted from the document. Table 1 show, by using as example the dependency “dobj(reflect-5, happiness-7)”, the list of extracted features.

Field Name	Content
RGD	“dobj-reflect-happiness”
RDG	“dobj-happiness-reflect”
GD	“reflect-happiness”
DG	“happiness-reflect”
G	“reflect”
D	“happiness”

Table 1: Field structure and corresponding content stored in the index.

There are three considerations explaining the rationale of using the presented set of six features.

- The choice of considering the governor and the dependent in both orders is to meet the possibility that the parser may produce different output based on how the text is written within the sentence. Such an order is affected also by the parser used. In our approach we decided to adopt the Stanford parser, but, obviously, any parser producing a list of dependencies like the one presented above can be used.
- For the same reason, we decided to extract features pruned by the relation element, because different parsers may use different kind of dependencies. The meaning of these features (the third and fourth ones) is to track the co-occurrence of terms independently by the relationship between them.
- Finally, the “G” and “D” features are used as backup purpose. Indeed, if, for training a particular model, a small number of samples is available, the use of single terms allows to apply a bag-of-words approach as a backup for computing document polarity. For these two features only nouns, verbs, adverbs, or adjectives are considered.

The set of features extracted from each dependencies is given as input to the component that will combine such features with either the polarity and domain information in order to construct the final representation of each document.

3.2 Structured Representation Construction

Once all features have been extracted, they are passed to the component in charge of structuring and storing them in the model repository that, for simplicity, we call “index”. As mentioned early, to each feature, the domain and polarity information are associated for building its equivalent structured representation. Where, the polarity associated with each feature contained in the model is the average of the polarities of the document in which each feature occurs. This shrewdness is necessary for distinguishing the polarities that each feature may assume in different domains. Indeed, classic approaches based on the use of polarized vocabularies do not consider the possibility that a particular feature may assume different polarities depending on the context in which they occur. An example has been presented in Section 1.

On the light of this, the construction of the structured representation of each feature has to consider two aspects: (i) each feature may appear in different domains, and (ii) for each feature an estimation of the polarity for each domain has to be computed.

Therefore, each feature is translated into the correspondent structured representation shown below. By considering as example the feature “RGD - dobj-reflect-happiness”, we have the following structure:

```
feature-type: RGD
feature-value: dobj-reflect-happiness
domain_1: polarity_1
domain_2: polarity_2
...
domain_n: polarity_n
```

The estimation of $polarity_i$ values associated with each domain is done by analyzing only the explicit information extracted from the training set. Values are computed as:

$$polarity_i(F) = \frac{k_F^i}{T_F^i} \in [-1, 1] \quad \forall i = 1, \dots, n, \quad (2)$$

where F is the feature taken into account, index i refers to domain D_i which the feature belongs to, n is the number of domains available in the training set, k_C^i is the arithmetic sum of the polarities observed for the feature F in the training set restricted to domain D_i , and T_C^i is the number of instances of the training set, restricted to domain D_i , in which feature F occurs.

Once all structured representations are built, they are stored in the repository. Such repository represents a multi-domain model for sentiment analysis purpose.

3.3 Polarity Computation

When an unclassified document needs to be evaluated, a procedure similar to the one adopted for building the model is used for computing its polarity.

A document is given as input to the Stanford parser and the list of dependencies is extracted and pruned by the ones containing stop words. Then, for each valid dependency, we build the related structured representation and we use it for estimating the polarity by analyzing information contained in the model. The final document polarity will be the average of the polarities estimated for each extracted dependency.

Let’s consider the following sentence:

“*I feel good and I feel healthy.*”

After the execution of the Stanford parser and the pruning of exceeding dependencies by using the same strategy described early, we obtain the following set of dependencies:

```

acomp(feel-2, good-3)
acomp(feel-6, healthy-7)

```

From these two dependencies, we generate the following two structures:

```

FEATURE ID: F1
feature-type: RGD; feature-value: acomp-feel-good
feature-type: RDG; feature-value: acomp-good-feel
feature-type: GD; feature-value: feel-good
feature-type: DG; feature-value: good-feel
feature-type: G; feature-value: feel
feature-type: D; feature-value: good

FEATURE ID: F2
feature-type: RGD; feature-value: acomp-feel-healthyd
feature-type: RDG; feature-value: acomp-healthy-feel
feature-type: GD; feature-value: feel-healthy
feature-type: DG; feature-value: healthy-feel
feature-type: G; feature-value: feel
feature-type: D; feature-value: healthy

```

For each structure I presented above, for which the domain D is given, we computed the SSV representing the polarity of the structure I in the domain which the structure belongs to. The Equation below, show how the SSV is computed.

$$\begin{aligned}
SSV(I) = & AVG(DP(RGD_{F1}) + DP(RDG_{F1}) + \\
& DP(GD_{F1}) + DP(DG_{F1}) + \\
& DP(G_{F1}) + DP(D_{F1}) + \\
& DP(RGD_{F2}) + DP(RDG_{F2}) + \\
& DP(GD_{F2}) + DP(DG_{F2}) + \\
& DP(G_{F2}) + DP(D_{F2}))
\end{aligned} \tag{3}$$

where DP is the function extracting the polarity of the feature I for the domain D , and AVG refers to the averaging operation of all detected polarities.

4 Experimental Evaluation

In this Section, we present the results obtained from our experimental campaign where we compared our representation in different settings.

Dataset construction And Baselines The training and testing of the system has been done on two different dataset. For creating the training model, we built structured document representation by using reviews contained in the Blitzer dataset and by applying the DRANZIERA protocol [50]. In particular, we used the balanced version of the dataset in order to same number of positive and negative samples. Concerning the test operation, we created a test set of 32.000 reviews compiled by using the same strategy used for building the Blitzer dataset ². Test set is even balanced with respect to the number of positive and negative opinions. The same philosophy has been used for the domains, where, for each of the 16 domains used in the test set, we had 1.000 positive, and as many negative, reviews.

Our approach (Structured Domain Dependent, SDD) has been compared with three baselines:

- Most Frequent Polarity: the accuracy obtained by the system if it guesses the same polarity for all samples contained in the test set.
- Structured Domain Independent: the accuracy obtained by using the proposed structured representation without considering domain information.
- Bag-Of-Word Domain Dependent: the accuracy obtained by using the classic statistical bag-of-words approach by considering also domain information.

Results and Discussion Table 2 shows the results obtained by the three baselines and by the proposed approach. First column contains the name of the approach, while the second one the accuracy obtained on the test set.

Approach	Accuracy
Most Frequent Polarity (MFP)	0.5000
Structured Domain Independent (SDI)	0.5407
Bag-Of-Word Domain Dependent (BDD)	0.6350
Structured Domain Dependent (SDD)	0.6834

Table 2: Accuracy obtained by our approach with respect to the three chosen baselines.

Results show that the proposed approach leads to better results with respect to all the baselines. Beside this, there is also a significant difference between the accuracies obtained by using domain-dependent features (BDD and SDD approaches) and the one obtained without considering domain information.

By focusing on the two approaches exploiting domain information, in Table 3, we reported the detailed accuracy obtained on each domain by the two approaches exploiting such information. First column contain the name of the domain, second column the number of features for each domain and the last two columns the accuracies obtained by the BDD and SDD approaches respectively.

By observing the results reported in Table 3, no particular correlations between the number of features and the accuracy of the approach can be noticed. Unexpectedly,

² The test set is available at <https://goo.gl/siOJbZ>

Domain	Features	BDD Accuracy	SDD Accuracy
automotive	259,239	0.6230	0.6935
baby	924,365	0.5980	0.5830
beauty	601,163	0.6390	0.6470
cell_phones_service	484,796	0.6115	0.6570
computer_video_games	1,247,408	0.5165	0.5725
electronics	944,796	0.6155	0.7180
gourmet_food	417,309	0.6310	0.6275
health_personal_care	768,616	0.6590	0.7180
jewelry_watches	358,677	0.6375	0.6540
kitchen_housewares	793,167	0.6460	0.7290
musical_instruments	130,005	0.6540	0.7225
office_products	180,172	0.6535	0.7105
software	1,146,081	0.6680	0.7070
sports_outdoors	869,576	0.6540	0.6810
tools_hardware	40,962	0.6830	0.7250
toys_games	833,887	0.6700	0.7885

Table 3: Accuracy obtained in each domain by the BDD and SDD approaches.

the worst result is obtained for the domain having the higher number of features, and one of the best results, obtained on the “tools_hardware” domain, is reported with a very low number of features compared to the others. One of the possible reasons may be the significant presence, in the set of documents used for building the model, of features having uncertain polarity. Indeed, if many features are used in either positive and negative contexts, it is difficult for the system to exploiting such information during the test phase for estimating document polarity. Further investigation in this direction may clarify this aspect.

Finally, we may notice that for the two domains, “gourmet_food” and “baby”, the performance of the bag of words approach, outperform the semantic one.

Approach Limits As we mentioned at the end of Section 2, the approach presented in this paper is a first attempt of exploring the use of structured representation of documents for addressing the sentiment analysis problem. For this reason, we performed a critical analysis of our work in order to highlight which are its limits and to outline a roadmap for future implementations. In particular, we detected three directions for extending the proposed approach:

- Improve dependencies pruning: in the feature extraction process, we pruned part of the dependencies extracted by the Stanford parser. In the light of the results reported in Table 3, we inferred that having a huge number of features is not preparatory for obtaining higher results. Therefore, a more restrictive policy should be implemented in pruning dependencies by trying to detect the most significant features despite the ones causing information overlapping between domains.
- Language coverage: a typical problem affecting the construction of language models is the language coverage of such models. Indeed, without having a large corpus for training the system, a significant number of terms information might be

excluded. This issue is strictly connected with the next one and it may share the possible solution.

- Improve the semantic aspect: one of the possibility for addressing the problem of language coverage, is the adoption of external semantic resources, for instance WordNet, for extending the meaning of each feature. This way, we will be able to reduce the total number of features, due to the use of a concept-based representation of each feature instead of a term-based one, and, at the same time, to increase the language coverage. Working in this direction will mean that the current structured representation will have to be revised accordingly.

5 Conclusion

In this paper, we described a system exploiting a structured representation of document for the problem of multi-domain sentiment analysis. Even if the representation used for structuring documents and the metric adopted for estimating document polarity is quite simple, the system obtained reasonable performances in the provided evaluation. Future work will address the possibility to exploit more sophisticated metrics considering the belonging of a document to a certain domain not in a binary but in a *fuzzy* fashion, measuring some sort of semantic relatedness of the sentence under test with each domain and using such measures as weights for the polarity detection phase. Moreover, we intend to explore the integration of knowledge bases in order to move toward a more cognitive technique able to improve the language coverage of the approach.

References

1. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up? sentiment classification using machine learning techniques. In: Proceedings of EMNLP, Philadelphia, Association for Computational Linguistics (July 2002) 79–86
2. Liu, B., Zhang, L.: A survey of opinion mining and sentiment analysis. In Aggarwal, C.C., Zhai, C.X., eds.: Mining Text Data. Springer (2012) 415–463
3. Blitzer, J., Dredze, M., Pereira, F.: Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In: ACL. (2007) 187–205
4. Pang, B., Lee, L.: A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In: ACL. (2004) 271–278
5. Dave, K., Lawrence, S., Pennock, D.M.: Mining the peanut gallery: opinion extraction and semantic classification of product reviews. In: WWW. (2003) 519–528
6. Paltoglou, G., Thelwall, M.: A study of information retrieval weighting schemes for sentiment analysis. In: ACL. (2010) 1386–1395
7. Tan, S., Wang, Y., Cheng, X.: Combining learn-based and lexicon-based techniques for sentiment detection without using labeled examples. In: SIGIR. (2008) 743–744
8. Qiu, L., Zhang, W., Hu, C., Zhao, K.: Selc: a self-supervised model for sentiment classification. In: CIKM. (2009) 929–936
9. Melville, P., Gryc, W., Lawrence, R.D.: Sentiment analysis of blogs by combining lexical knowledge with text classification. In: KDD. (2009) 1275–1284
10. Taboada, M., Brooke, J., Tofiloski, M., Voll, K.D., Stede, M.: Lexicon-based methods for sentiment analysis. Computational Linguistics **37**(2) (2011) 267–307

11. Somasundaran, S.: Discourse-level relations for Opinion Analysis. PhD thesis, University of Pittsburgh (2010)
12. Wang, H., Zhou, G.: Topic-driven multi-document summarization. In: IALP. (2010) 195–198
13. 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, Denver, Colorado, Association for Computational Linguistics (June 2015) 502–509
14. Petrucci, G., Dragoni, M.: An information retrieval-based system for multi-domain sentiment analysis. In Gandon, F., Cabrio, E., Stankovic, M., Zimmermann, A., eds.: 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., Springer (2015) 234–243
15. Rexha, A., Kröll, M., Dragoni, M., Kern, R.: Exploiting propositions for opinion mining. In Sack, H., Dietze, S., Tordai, A., Lange, C., eds.: Semantic Web Challenges - Third SemWebEval Challenge at ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Volume 641 of Communications in Computer and Information Science., Springer (2016) 121–125
16. Federici, M., Dragoni, M.: A knowledge-based approach for aspect-based opinion mining. In Sack, H., Dietze, S., Tordai, A., Lange, C., eds.: Semantic Web Challenges - Third SemWebEval Challenge at ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Volume 641 of Communications in Computer and Information Science., Springer (2016) 141–152
17. Dragoni, M., Tettamanzi, A.G., da Costa Pereira, C.: Propagating and aggregating fuzzy polarities for concept-level sentiment analysis. *Cognitive Computation* 7(2) (2015) 186–197
18. Dragoni, M., Tettamanzi, A.G.B., da Costa Pereira, C.: A fuzzy system for concept-level sentiment analysis. In Presutti, V., Stankovic, M., Cambria, E., Cantador, I., Iorio, A.D., Noia, T.D., Lange, C., Recupero, D.R., Tordai, A., eds.: 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., Springer (2014) 21–27
19. Petrucci, G., Dragoni, M.: The IRMUDOSA system at ESWC-2016 challenge on semantic sentiment analysis. In Sack, H., Dietze, S., Tordai, A., Lange, C., eds.: Semantic Web Challenges - Third SemWebEval Challenge at ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Volume 641 of Communications in Computer and Information Science., Springer (2016) 126–140
20. da Costa Pereira, C., Dragoni, M., Pasi, G.: A prioritized "and" aggregation operator for multidimensional relevance assessment. In Serra, R., Cucchiara, R., eds.: 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., Springer (2009) 72–81
21. Federici, M., Dragoni, M.: Towards unsupervised approaches for aspects extraction. In Dragoni, M., Recupero, D.R., Denecke, K., Deng, Y., Declerck, T., eds.: Joint Proceedings of the 2th Workshop on Emotions, Modality, Sentiment Analysis and the Semantic Web and the 1st International Workshop on Extraction and Processing of Rich Semantics from Medical Texts co-located with ESWC 2016, Heraklion, Greece, May 29, 2016. Volume 1613 of CEUR Workshop Proceedings., CEUR-WS.org (2016)
22. Riloff, E., Patwardhan, S., Wiebe, J.: Feature subsumption for opinion analysis. In: EMNLP. (2006) 440–448
23. Wilson, T., Wiebe, J., Hwa, R.: Recognizing strong and weak opinion clauses. *Computational Intelligence* 22(2) (2006) 73–99

24. Aprosio, A.P., Corcoglioniti, F., Dragoni, M., Rospocher, M.: Supervised opinion frames detection with RAID. In Gandon, F., Cabrio, E., Stankovic, M., Zimmermann, A., eds.: *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.*, Springer (2015) 251–263
25. Hatzivassiloglou, V., Wiebe, J.: Effects of adjective orientation and gradability on sentence subjectivity. In: *COLING*. (2000) 299–305
26. Kim, S.M., Hovy, E.H.: Crystal: Analyzing predictive opinions on the web. In: *EMNLP-CoNLL*. (2007) 1056–1064
27. Rexha, A., Kröll, M., Dragoni, M., Kern, R.: Polarity classification for target phrases in tweets: A word2vec approach. In Sack, H., Rizzo, G., Steinmetz, N., Mladenic, D., Auer, S., Lange, C., eds.: *The Semantic Web - ESWC 2016 Satellite Events*, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Volume 9989 of *Lecture Notes in Computer Science*. (2016) 217–223
28. Dragoni, M., Recupero, D.R.: Challenge on fine-grained sentiment analysis within ESWC2016. In Sack, H., Dietze, S., Tordai, A., Lange, C., eds.: *Semantic Web Challenges - Third SemWebEval Challenge at ESWC 2016*, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Volume 641 of *Communications in Computer and Information Science.*, Springer (2016) 79–94
29. Jakob, N., Gurevych, I.: Extracting opinion targets in a single and cross-domain setting with conditional random fields. In: *EMNLP*. (2010) 1035–1045
30. Jin, W., Ho, H.H., Srihari, R.K.: Opinionminer: a novel machine learning system for web opinion mining and extraction. In: *KDD*. (2009) 1195–1204
31. Liu, B., Hu, M., Cheng, J.: Opinion observer: analyzing and comparing opinions on the web. In: *WWW*. (2005) 342–351
32. Wu, Y., Zhang, Q., Huang, X., Wu, L.: Phrase dependency parsing for opinion mining. In: *EMNLP*. (2009) 1533–1541
33. Su, Q., Xu, X., Guo, H., Guo, Z., Wu, X., Zhang, X., Swen, B., Su, Z.: Hidden sentiment association in chinese web opinion mining. In: *WWW*. (2008) 959–968
34. Dragoni, M., Azzini, A., Tettamanzi, A.: A novel similarity-based crossover for artificial neural network evolution. In Schaefer, R., Cotta, C., Kolodziej, J., Rudolph, G., eds.: *Parallel Problem Solving from Nature - PPSN XI, 11th International Conference, Kraków, Poland, September 11-15, 2010, Proceedings, Part I*. Volume 6238 of *Lecture Notes in Computer Science.*, Springer (2010) 344–353
35. Qiu, G., Liu, B., Bu, J., Chen, C.: Opinion word expansion and target extraction through double propagation. *Computational Linguistics* **37**(1) (2011) 9–27
36. Dragoni, M.: A three-phase approach for exploiting opinion mining in computational advertising. *IEEE Intelligent Systems* **32**(3) (2017) 21–27
37. Barbosa, L., Feng, J.: Robust sentiment detection on twitter from biased and noisy data. In: *COLING (Posters)*. (2010) 36–44
38. Birmingham, A., Smeaton, A.F.: Classifying sentiment in microblogs: is brevity an advantage? In: *CIKM*. (2010) 1833–1836
39. Go, A., Bhayani, R., Huang, L.: Twitter sentiment classification using distant supervision. *CS224N Project Report*, Stanford University (2009)
40. Cambria, E., Hussain, A.: *Sentic Computing: Techniques, Tools, and Applications*. Volume 2 of *SpringerBriefs in Cognitive Computation*. Springer, Dordrecht, Netherlands (2012)
41. Cambria, E., Hussain, A.: Sentic album: Content-, concept-, and context-based online personal photo management system. *Cognitive Computation* **4**(4) (2012) 477–496
42. Wang, Q.F., Cambria, E., Liu, C.L., Hussain, A.: Common sense knowledge for handwritten chinese recognition. *Cognitive Computation* **5**(2) (2013) 234–242

43. Pan, S.J., Ni, X., Sun, J.T., Yang, Q., Chen, Z.: Cross-domain sentiment classification via spectral feature alignment. In: WWW. (2010) 751–760
44. Yoshida, Y., Hirao, T., Iwata, T., Nagata, M., Matsumoto, Y.: Transfer learning for multiple-domain sentiment analysis—identifying domain dependent/independent word polarity. In: AACL (2011) 1286–1291
45. Ponomareva, N., Thelwall, M.: Semi-supervised vs. cross-domain graphs for sentiment analysis. In: RANLP. (2013) 571–578
46. Huang, S., Niu, Z., Shi, C.: Automatic construction of domain-specific sentiment lexicon based on constrained label propagation. *Knowl.-Based Syst.* **56** (2014) 191–200
47. Dragoni, M., da Costa Pereira, C., Tettamanzi, A.G.B., Villata, S.: Smack: An argumentation framework for opinion mining. In Kambhampati, S., ed.: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016, IJCAI/AAAI Press (2016) 4242–4243
48. da Costa Pereira, C., Dragoni, M., Pasi, G.: Multidimensional relevance: Prioritized aggregation in a personalized information retrieval setting. *Inf. Process. Manage.* **48**(2) (2012) 340–357
49. Manning, C.D., Surdeanu, M., Bauer, J., Finkel, J., Bethard, S.J., McClosky, D.: The Stanford CoreNLP natural language processing toolkit. In: Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations, Baltimore, Maryland, Association for Computational Linguistics (June 2014) 55–60
50. Dragoni, M., Tettamanzi, A., da Costa Pereira, C.: DRANZIERA: an evaluation protocol for multi-domain opinion mining. In Calzolari, N., Choukri, K., Declerck, T., Goggi, S., Grobelnik, M., Maegaard, B., Mariani, J., Mazo, H., Moreno, A., Odijk, J., Piperidis, S., eds.: Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016., European Language Resources Association (ELRA) (2016)