# A Branching Strategy For Unsupervised Aspect-based Sentiment Analysis

Marco Federici[1,2] and Mauro Dragoni[2]

[1] Universitá di Trento, Italy
[2] Fondazione Bruno Kessler, Trento, Italy
federici|dragoni@fbk.eu

**Abstract.** One of the most recent opinion mining research directions falls in the extraction of polarities referring to specific entities (called "aspects") contained in the analyzed texts. The detection of such aspects may be very critical especially when the domain which documents belong to is unknown. Indeed, while in some contexts it is possible to train domain-specific models for improving the effectiveness of aspects extraction algorithms, in others the most suitable solution is to apply unsupervised techniques by making the used algorithm independent from the domain. In this work, we implemented different unsupervised solutions into an aspect-based opinion mining system. Such solutions are based on the use of semantic resources for performing the extraction of aspects from texts. The algorithms have been tested on benchmarks provided by the SemEval campaign and have been compared with the results obtained by domain-adapted techniques.

## 1 Introduction

Opinion Mining is a natural language processing (NLP) task that aims to classify documents according to their opinion (polarity) on a given subject [36]. This task has created a considerable interest due to its wide applications in different domains like marketing, politics, and social sciences. Generally, the polarity of a document is computed by analyzing the expressions contained in the full text by leading to the issue of not distinguishing which are the subjects of each opinion. Therefore, the natural evolution of the opinion mining research field has been focused on the extraction of all subjects ("aspects") from texts in order to make systems able to compute the polarity associated to each aspect in an independent way [25].

Let us consider the following example:

*Yesterday, I bought a new smartphone.*
*The quality of the display is very good, but the buttery lasts too little.*

In the sentence above, we may identify three aspects: "smartphone", "display", and "battery". Each aspect has a different opinion associated with it, in particular:

- "display" → "very good"
- "battery" → "too little"
- "smarthphone" → no explicit opinions, therefore its polarity can be inferred by averaging the opinions associated with all other aspects.

Another important consideration related to this example is that it is easy to detect which is the domain of the analyzed text. In this case, by assuming to have a training set, it should be possible to build domain-specific models for supporting the extraction of the aspects. However, this strategy is in contrast with two considerations coming from real-world scenarios: (i) it is difficult to find annotated dataset related to all possible domains, and (ii) in the same document, it is possible to have sentences belonging to many domains by making the adoption of a domain-specific models not feasible.

To overcome these issues, we propose a set of unsupervised approaches based on natural language processing approaches that do not rely to any domain-specific information. The goal of this study is to provide techniques that are able to reach an effectiveness comparable with supervised systems.

The paper is structured as follows. In Section 2, we provide an overview of the opinion mining field with a focus on aspects extraction approaches. Section 3 presents the natural language processing layer built for supporting the approaches described in Sections 4 and 5. Section 6 discusses the performance of each algorithm; while, Section 7 concludes the paper.

## 2 Related Work

The topic of sentiment analysis has been studied extensively in the literature [31], 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 [35], 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 [10], as well as by defining weighting schemata for enhancing classification accuracy [33].

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 [49] and [42], 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 [32], 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 [48].

Another research direction concerns the exploitation of discourse-analysis techniques. [46] discusses some discourse-based supervised and unsupervised approaches for opinion analysis; while in [50], the authors present an approach to identify discourse relations.

The approaches presented above are applied at the document-level[12,37,43,20], 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 [19,18,38] to the use of aggregation techniques [8,9] 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 [21,45,52] and systems implementing the capabilities of identifying opinion's holder, target, and polarity have been presented [1]. Once subjective sentences are identified, the same methods as for sentiment classification may be applied. For example, in [24] the authors consider gradable adjectives for sentiment spotting; while in [29,44] 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 [16]. 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) [27], hidden Markov models (HMM) [28], sequential rule mining [30], dependency tree kernels [53], clustering [47], and genetic algorithms [14]. In [41], 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 [13]. 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 [2] and [3], 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 [23], 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* [6]. Application domains where sentic computing has already shown its potential are the cognitive-inspired classification of images [5], of texts in natural language, and of handwritten text [51].

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 [4,34,54], and (ii) the use of propagation of labels through graph structures [40,26,19,15].

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 Underlying NLP Layer

A number of different approaches has been tested in order to accomplish aspect extraction task. Each one uses different functionalities offered by the Stanford NLP Library but every technique is characterized by a common preliminary phase.

First of all, *WordNet*[3] [22] resource is used together with Stanford's part of speech annotation to detect compound nouns. Lists of consecutive nouns and word sequences contained in Wordnet compound nouns vocabulary are merged into a single word in order to force Stanford library to consider them as a single unit during the following phases.

The entire text is then fed to the co-reference resolution module to compute pronoun references which are stored in an index-reference map.

The next operation consists in detecting which word expresses polarity within each sentence. To achieve this task *SenticNet*[4] [7], *General Inquirer dictionary* [5] [39] and *MPQA*[6] [11] sentiment lexicons have been used.

While SenticNet expresses polarity values in the continuous range from -1 to 1, the other two resources been normalized: the General Inquirer words have positive values of polarity if they belong to the "Positiv" class while negative if they belong to "Negativ" one, zero otherwise, similarly, MPQA "polarity" labels are used to infer a numerical values. Only words with a non-zero polarity value in at least one resource are considered as opinion words (e.g. word "third" is not present in MPQA and SenticNet and has a 0 value according to General Inquirer, consequently, it is not a valid opinion word; on the other hand, word "huge" has a positive 0.069 value according to SenticNet, a negative value in MPQA and 0 value according to General Inquirer, therefore, it is a possible opinion word even if lexicons express contrasting values). Every noun (single or complex) is considered an aspect as long as it's connected to at least one opinion

[3] https://wordnet.princeton.edu/
[4] http://sentic.net/
[5] http://www.wjh.harvard.edu/ inquirer/spreadsheet_guide.htm
[6] http://mpqa.cs.pitt.edu/corpora/mpqa_corpus/

and it's not in the stopword list. This list has been created starting from the "Onix" text retrieval engine stopwords list[7] and it contains words without a specific meaning (such as "thing") and special characters.

Opinions associated with pronouns are connected to the aspect they are referring to; instead, if pronouns reference can't be resolved, they are both discarded.

The main task of the system is, then, represented by connecting opinions with possible aspects. Two different approaches have been tested with a few variants. The first one relies on the syntactic tree while the second one is based on grammar dependencies.

The sentence "I enjoyed the screen resolution, it's amazing for such a cheap laptop." has been used to underline differences in connection techniques.

The preliminary phase merges words "screen" and "resolution" into a single word "Screenresolution" because they are consecutive nouns. Co-reference resolution module extracts a relation between "it" and "Screenresolution". This relation is stored so that every possible opinion that would be connected to "it" will be connected to "Screenresolution" instead. Figure 1 shows the syntax tree while Figure 2 represents the grammar relation graph generated starting from the example sentence. Both structures have been computed using Stanford NLP modules ("parse", "depparse").

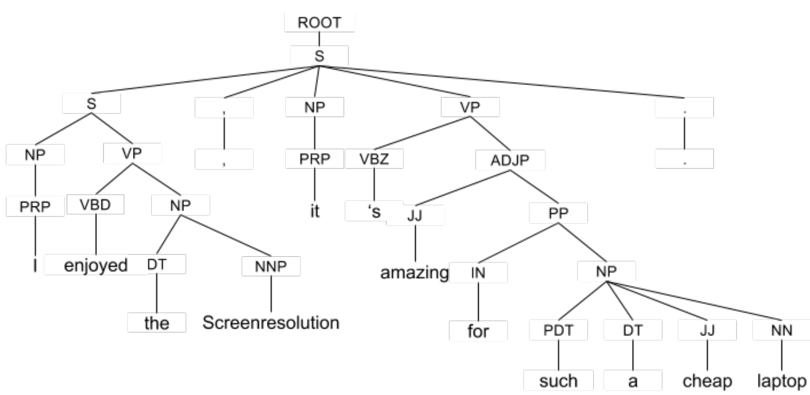


Fig. 1: Example of syntax tree.

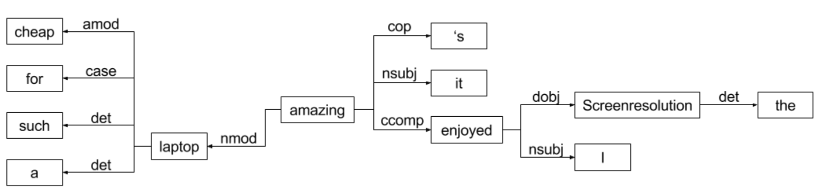


Fig. 2: Example of the grammar relations graph.

[7] The used stopwords list is available at http://www.lextek.com/manuals/onix/stopwords1.html

## 4 Unsupervised Approaches - Syntax-Tree-Based Approach

These typologies of approaches are based on syntax tree structures created by Stanford NLP library. In order to explain how the algorithms connect opinion with aspects a few definition are needed:

- "Intermediate node": tree node which is not a leaf;
- "Sentence node": intermediate node labeled with one of the following:
  - ROOT - Root of the tree
  - S - Sentence
  - SBAR - Clause introduced by a (possibly empty) subordinating conjunction
  - SBARQ - Direct question introduced by a wh-word or a wh phrase
  - SQ - Inverted yes/no question or main clause of a wh-question
  - SINV - Inverted declarative sentence
  - PRN - Parenthetical
  - FRAG - Fragment
- "Noun Phrase node": intermediate node labeled with NP tag

Approaches differ in rules adopted for associating intermediate nodes that define how aspects are extracted by starting from their child nodes.

*Approach 1.1* Each polarized adjective is connected with each possible aspect in the same sentence.

Figure 3 shows she propagation of aspects and opinion in the tree with red lines representing propagation of aspects, blue lines for opinions and purple ones when both are propagated to the upper level.

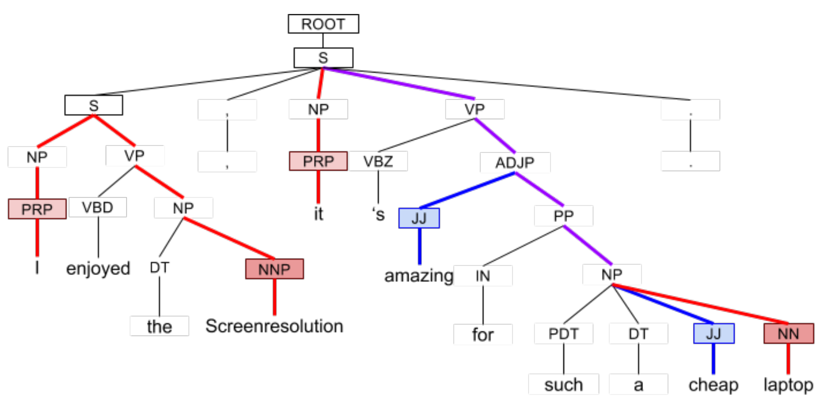


Fig. 3: Parser tree generated by the approach 1.1.

Within the sub-sentence "I enjoyed the Screenresolution" only aspects are detected, consequently, once the Sentence Level node is reached, no connection is done. On the other hand, both polarized adjectives "cheap" and "amazing" are propagated until they reach the top sentence node together with "it" and "laptop" aspects, then, they are connected with each other.

The results are shown in Figure 4.

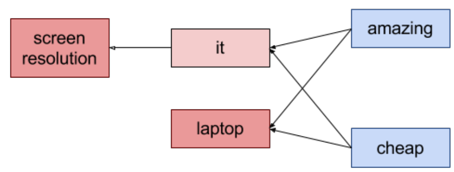


Fig. 4: Relationships generated by the approach 1.1.

*Approach 1.2* Each polarized adjective is connected to each possible aspect within the same sentence or noun phrase.

Influences of this variant are underlined in Figure 5 with the same notation.

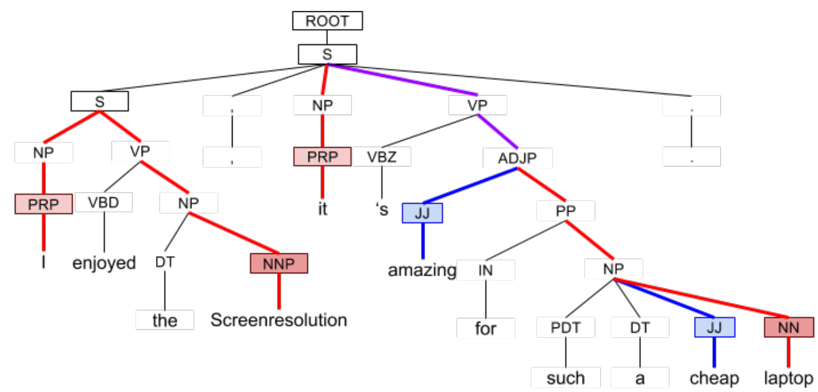


Fig. 5: Parser tree generated by the approach 1.2.

Even if extracted aspects are the same, the opinion "cheap" is associated only with the name "laptop" as shown in Figure 6.

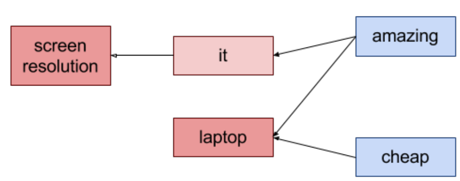


Fig. 6: Relationships generated by the approach 1.2.

*Approach 1.3* When both aspects set and opinion words set related to a node are not empty, each opinion word is connected to the related aspect and removed from the opinion words set. Opinion words and possible aspects are removed anyway in sentence nodes.

Figure 7 shows the effects of the association rules mentioned above.

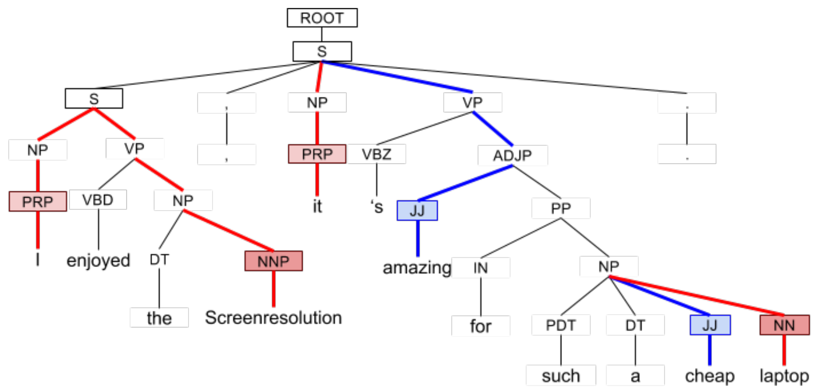


Fig. 7: Parser tree generated by the approach 1.3.

Once again, even if aspects extracted are the same, the connections are different (Figure 8).

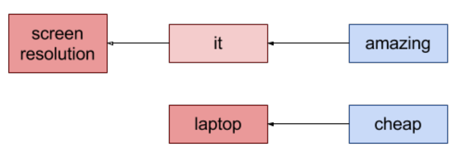


Fig. 8: Relationships generated by the approach 1.3.

## 5 Unsupervised Approaches - Grammar-Dependencies-based Approach

The other set of approaches proposed in this paper exploits grammar dependencies instead of syntax tree to detect aspect-opinion associations. Grammar dependencies computed by Stanford NLP modules (which are represented by the labeled graph in picture [1.2]) can be expressed by triples: $\{Relationtype, Governor, Dependant\}$. One of the most important difference with the previous methodology is represented by the possibility of detecting opinion expressed by word that are not adjectives (such as verbs that are considered by approaches 2.2 and 2.3). Different approaches have been tested in order to detect which kind of triple can be interpreted as a connection between an opinion word and a possible aspect.

*Approach 2.1* The following two rules are implemented:

*Rule 1*: Each adjectival modifier (amod) relation expresses a connection between an aspect and an opinion word if and only if the governor is a possible aspect and the dependant is a polarized adjective.

*Rule 2*: Each nominal subject (nsubj) relation expresses a connection between an aspect and an opinion word if and only if the governor is a polarized opinion and the dependant is a possible aspect.

Figure 9 underlines aspect-opinion connections mined through the process.

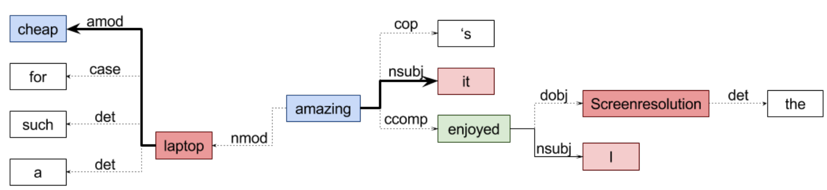


Fig. 9: Parser tree generated by the approach 2.1.

Resulting aspects are shown in Figure 10.

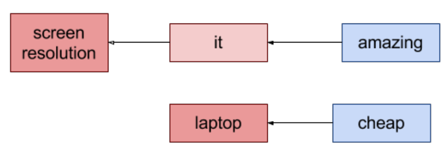


Fig. 10: Relationships generated by the approach 2.1.

*Approach 2.2* The Rules "1" and "2" are both used, in addition a third rule is introduced:

*Rule 3*: Each direct object (dobj) relation expresses a connection between an aspect and an opinion word if and only if the governor is a polarized word and the dependant is a possible aspect.

Figure 11 and 12 shows the results of the aspect detection process with the addition of the direct object relation.

*Approach 2.3* The Rules "1" and "3" are both used, while Rule "2" is changed as follows:

*Rule 2.1*: Each nominal subject (nsubj) relation expresses a connection between an aspect and an opinion word if and only if the governor is a polarized word and the dependant is a possible aspect.

Figure 13 shows results of the modification of the rules. Even if the relation between "enjoyed" and "I" is detected, "I" is not considered as a valid aspect since it's has an unresolved reference in the current context.

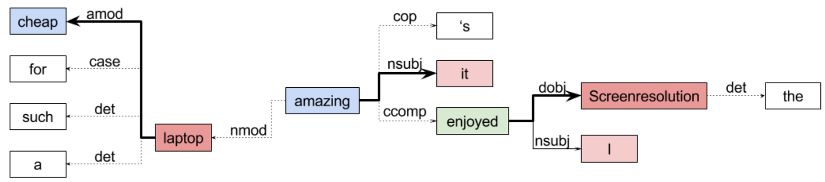


Fig. 11: Parser tree generated by the approach 2.2.

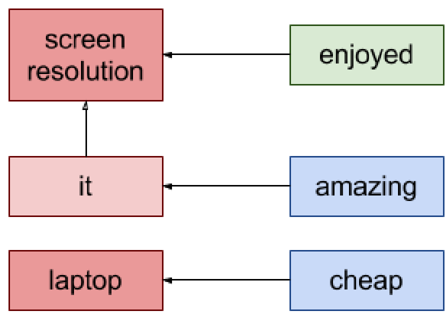


Fig. 12: Relationships generated by the approach 2.2.

Results are the same as the previous example (Figure 14).

## 6 Evaluation

In this Section, we present the evaluation of the proposed system performed by following the DRANZIERA protocol [17]. Each approach has been tested on two datasets provided by the Task 12 of SemEval 2015 evaluation campaign, namely "Laptop" and "Restaurant". To evaluate results a notion of correctness has to be introduced: if the extracted aspects is equal, contained or contains the correct one, it's considered to be correct (for example if the extracted aspect is "screen", while the annotated one is "screen of the computer" or vice versa, the result of the system is considered to correct). Here, we focus our evaluation on two perspectives:

– *Aspect extraction*. The main task in charge to the system is the extraction of aspects from text. Such a task is important for defining, later in the analysis process, which aspects are the most significant ones. This evaluation task focused on measuring the effectiveness of the aspect-extraction approach.
– *Polarity detection*. The computation of the aspect's polarity enables the detection of which product features are strong or weak. The sentiment component is in charge of inferring the polarity of each aspect given the context in which such an aspect is included. Here, we measured the capability of the system of inferring the correct polarity.

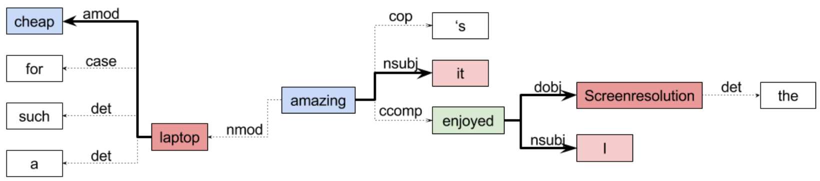


Fig. 13: Parser tree generated by the approach 2.3.

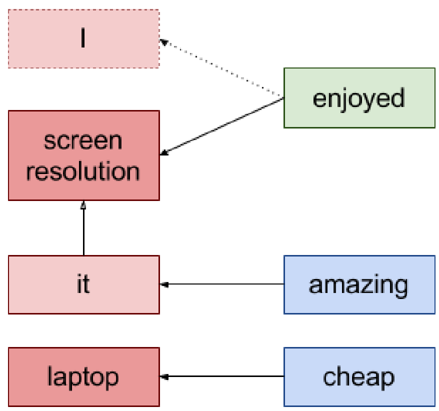


Fig. 14: Relationships generated by the approach 2.3.

### 6.1 Evaluation on Aspect Extraction

Table 1 reports the results obtained by our approach on the aspect extraction benchmark used in SemEval 2015 Task 12. The algorithm has been tested on the "Restaurant" and "Laptop" datasets respectively. The overall performance are in line with the best systems participating in the evaluation campaign and, on the "Laptop" dataset, our aspect extraction approach recorded the best precision and F-measure. It is also important to highlight that all the systems we compared to, apply supervised approaches for extracting aspects, while our approach implements an unsupervised technique. This way, it is possible to implement the system in any environment without the requirement of training a new model.

Concerning the "Restaurant" domain, the gap between our approach and the best ones is given by the conservative strategy implemented for extracting aspects. One of the most common issue in unsupervised aspect-based approach is the extraction of false positive aspects [**?**]. The major consequence of such issue is the poor effectiveness of modules exploiting the outcome of the aspect extraction component. Unfortunately, the adoption of a conservative strategy leads to lower recall values. However, the latter is a preferable solution by considering the massive use of the aspects in the other components of the platform.

Table 1: Results obtained on the aspect extraction task, for the "Restaurant" and "Laptop" datasets, on the SemEval 2015 benchmark. For each dataset, we reported Precision, Recall, and F-Measure. Acronyms refer to the systems participated in the SemEval 2015 competition.

| | **Restaurant** | | | **Laptop** | | |
|---|---|---|---|---|---|---|
| **System Acronym** | **Precision** | **Recall** | **F-Measure** | **Precision** | **Recall** | **F-Measure** |
| IHS-RD-Belarus | **0.7095** | 0.3845 | 0.4987 | 0.5548 | 0.4483 | 0.4959 |
| LT3 pred | 0.5154 | 0.5600 | 0.5367 | - | - | - |
| NLANGP | 0.6385 | **0.6154** | **0.6268** | 0.6425 | 0.4208 | 0.5086 |
| sentiue | 0.6332 | 0.4722 | 0.5410 | 0.5773 | 0.4409 | 0.5000 |
| SIEL | 0.6440 | 0.5135 | 0.5714 | - | - | - |
| TJUdeM | 0.4782 | 0.5806 | 0.5244 | 0.4489 | **0.4820** | 0.4649 |
| UFRGS | 0.6555 | 0.4322 | 0.5209 | 0.5066 | 0.4040 | 0.4495 |
| UMDuluth | 0.5697 | 0.5741 | 0.5719 | - | - | - |
| V3 | 0.4244 | 0.4129 | 0.4185 | 0.2710 | 0.2310 | 0.2494 |
| **SYSTEM** | 0.6895 | 0.5368 | 0.6036 | **0.6702** | 0.4157 | **0.5131** |

### 6.2 Evaluation on Polarity Computation

Table 2 reports the results of the polarity computation technique. The approach has been evaluated on the two datasets mentioned above. Here, we measured the accuracy of the polarity detection algorithm: given the set of opinion words associated with an aspect, such a polarity is computed by aggregating the fuzzy polarities of each opinion words. Results demonstrated the effectiveness of the polarity detection techniques implemented within the system by obtaining the best performance on the "Laptop" dataset, and the third best one on the "Restaurant" dataset. After a detailed analysis of the results, we noticed that the reason for which our approach performs better on the "Laptop" dataset is due to the simple language used for describing product features. Indeed, in the "Restaurant" dataset opinions are expressed in a more articulated way and sometimes the approach fails to detect the right polarity. Improvements in this direction will be part of the future work.

## 7 Conclusions

In this paper, we presented a set of unsupervised approaches for aspect-based sentiment analysis. Such approaches have been tested on two SemEval benchmarks: the "Laptop" and "Restaurant" datasets used in the Task 12 of SemEval 2015 evaluation campaign. Results demonstrated how without using learning techniques the results can be comparable with the ones obtained by trained systems. Future work includes refinement of the proposed approaches in order to make them suitable for real-world implementation.

Table 2: Results obtained concerning the computation of polarities associated with single aspects on the SemEval 2015 benchmark. For each dataset, we reported the accuracy obtained in computing polarities ("positive", or "negative"). Acronyms refer to the systems participated in the SemEval 2015 competition.

| **System Acronym** | **Acc. "Restaurant"** | **Acc. "Laptop"** |
|---|---|---|
| ECNU | 0.7810 | 0.7829 |
| EliXa | 0.7005 | 0.7291 |
| lsislif | 0.7550 | 0.7787 |
| LT3 | 0.7502 | 0.7376 |
| sentiue | **0.7869** | 0.7934 |
| SIEL | 0.7124 | - |
| SINAI | 0.6071 | 0.6585 |
| TJUdeM | 0.6887 | 0.7323 |
| UFRGS | 0.7171 | 0.6733 |
| UMDuluth | 0.7112 | - |
| V3 | 0.6946 | 0.6838 |
| wnlp | 0.7136 | 0.7207 |
| **SYSTEM** | 0.7794 | **0.8589** |

## References


1. Alessio Palmero Aprosio, Francesco Corcoglioniti, Mauro Dragoni, and Marco Rospocher. 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*, pages 251–263. Springer, 2015.
2. Luciano Barbosa and Junlan Feng. Robust sentiment detection on twitter from biased and noisy data. In *COLING (Posters)*, pages 36–44, 2010.
3. Adam Bermingham and Alan F. Smeaton. Classifying sentiment in microblogs: is brevity an advantage? In *CIKM*, pages 1833–1836, 2010.
4. John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL*, pages 187–205, 2007.
5. E. Cambria and A. Hussain. Sentic album: Content-, concept-, and context-based online personal photo management system. *Cognitive Computation*, 4(4):477–496, 2012.
6. E. Cambria and A. Hussain. *Sentic Computing: Techniques, Tools, and Applications*, volume 2 of *SpringerBriefs in Cognitive Computation*. Springer, Dordrecht, Netherlands, 2012.
7. Erik Cambria, Robert Speer, Catherine Havasi, and Amir Hussain. Senticnet: A publicly available semantic resource for opinion mining. In *AAAI Fall Symposium: Commonsense Knowledge*, 2010.
8. Célia da Costa Pereira, Mauro Dragoni, and Gabriella Pasi. 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*, pages 72–81. Springer, 2009.
9. Célia da Costa Pereira, Mauro Dragoni, and Gabriella Pasi. Multidimensional relevance: Prioritized aggregation in a personalized information retrieval setting. *Inf. Process. Manage.*, 48(2):340–357, 2012.

10. Kushal Dave, Steve Lawrence, and David M. Pennock. Mining the peanut gallery: opinion extraction and semantic classification of product reviews. In *WWW*, pages 519–528, 2003.
11. Lingjia Deng and Janyce Wiebe. 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*, pages 1323–1328. The Association for Computational Linguistics, 2015.
12. M. Dragoni. Shellfbk: An information retrieval-based system for multi-domain sentiment analysis. In *Proceedings of the 9th International Workshop on Semantic Evaluation*, SemEval '2015, pages 502–509, Denver, Colorado, June 2015. Association for Computational Linguistics.
13. Mauro Dragoni. A three-phase approach for exploiting opinion mining in computational advertising. *IEEE Intelligent Systems*, 32(3):21–27, 2017.
14. Mauro Dragoni, Antonia Azzini, and Andrea Tettamanzi. A novel similarity-based crossover for artificial neural network evolution. In Robert Schaefer, Carlos Cotta, Joanna Kolodziej, and Günter Rudolph, editors, *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*, pages 344–353. Springer, 2010.
15. Mauro Dragoni, Célia da Costa Pereira, Andrea G. B. Tettamanzi, and Serena Villata. Smack: An argumentation framework for opinion mining. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 4242–4243. IJCAI/AAAI Press, 2016.
16. Mauro Dragoni and Diego Reforgiato Recupero. Challenge on fine-grained sentiment analysis within ESWC2016. In Harald Sack, Stefan Dietze, Anna Tordai, and Christoph Lange, editors, *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*, pages 79–94. Springer, 2016.
17. Mauro Dragoni, Andrea Tettamanzi, and Célia da Costa Pereira. DRANZIERA: an evaluation protocol for multi-domain opinion mining. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Sara Goggi, Marko Grobelnik, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asunción Moreno, Jan Odijk, and Stelios Piperidis, editors, *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.
18. Mauro Dragoni, Andrea G. B. Tettamanzi, and Célia da Costa Pereira. 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*, pages 21–27. Springer, 2014.
19. Mauro Dragoni, Andrea G.B. Tettamanzi, and Célia da Costa Pereira. Propagating and aggregating fuzzy polarities for concept-level sentiment analysis. *Cognitive Computation*, 7(2):186–197, 2015.
20. Marco Federici and Mauro Dragoni. A knowledge-based approach for aspect-based opinion mining. In Harald Sack, Stefan Dietze, Anna Tordai, and Christoph Lange, editors, *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*, pages 141–152. Springer, 2016.
21. Marco Federici and Mauro Dragoni. Towards unsupervised approaches for aspects extraction. In Mauro Dragoni, Diego Reforgiato Recupero, Kerstin Denecke, Yihan Deng, and Thierry Declerck, editors, *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. Christiane Fellbaum. *WordNet: An Electronic Lexical Database.* MIT Press, Cambridge, MA, 1998.
23. Alec Go, Richa Bhayani, and Lei Huang. Twitter sentiment classification using distant supervision. CS224N Project Report, Standford University, 2009.
24. Vasileios Hatzivassiloglou and Janyce Wiebe. Effects of adjective orientation and gradability on sentence subjectivity. In *COLING*, pages 299–305, 2000.
25. Minqing Hu and Bing Liu. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177. ACM, 2004.
26. Sheng Huang, Zhendong Niu, and Chongyang Shi. Automatic construction of domain-specific sentiment lexicon based on constrained label propagation. *Knowl.-Based Syst.*, 56:191–200, 2014.
27. Niklas Jakob and Iryna Gurevych. Extracting opinion targets in a single and cross-domain setting with conditional random fields. In *EMNLP*, pages 1035–1045, 2010.
28. Wei Jin, Hung Hay Ho, and Rohini K. Srihari. Opinionminer: a novel machine learning system for web opinion mining and extraction. In *KDD*, pages 1195–1204, 2009.
29. Soo-Min Kim and Eduard H. Hovy. Crystal: Analyzing predictive opinions on the web. In *EMNLP-CoNLL*, pages 1056–1064, 2007.
30. Bing Liu, Minqing Hu, and Junsheng Cheng. Opinion observer: analyzing and comparing opinions on the web. In *WWW*, pages 342–351, 2005.
31. Bing Liu and Lei Zhang. A survey of opinion mining and sentiment analysis. In C. C. Aggarwal and C. X. Zhai, editors, *Mining Text Data*, pages 415–463. Springer, 2012.
32. Prem Melville, Wojciech Gryc, and Richard D. Lawrence. Sentiment analysis of blogs by combining lexical knowledge with text classification. In *KDD*, pages 1275–1284, 2009.
33. Georgios Paltoglou and Mike Thelwall. A study of information retrieval weighting schemes for sentiment analysis. In *ACL*, pages 1386–1395, 2010.
34. Sinno Jialin Pan, Xiaochuan Ni, Jian-Tao Sun, Qiang Yang, and Zheng Chen. Cross-domain sentiment classification via spectral feature alignment. In *WWW*, pages 751–760, 2010.
35. Bo Pang and Lillian Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *ACL*, pages 271–278, 2004.
36. Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. Thumbs up? sentiment classification using machine learning techniques. In *Proceedings of EMNLP*, pages 79–86, Philadelphia, July 2002. Association for Computational Linguistics.
37. Giulio Petrucci and Mauro Dragoni. An information retrieval-based system for multi-domain sentiment analysis. 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*, pages 234–243. Springer, 2015.
38. Giulio Petrucci and Mauro Dragoni. The IRMUDOSA system at ESWC-2016 challenge on semantic sentiment analysis. In Harald Sack, Stefan Dietze, Anna Tordai, and Christoph Lange, editors, *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*, pages 126–140. Springer, 2016.
39. Stone P.J, D.C. Dunphy, and S. Marshall. *The General Inquirer: A Computer Approach to Content Analysis*. Oxford, England: M.I.T. Press, 1966.
40. Natalia Ponomareva and Mike Thelwall. Semi-supervised vs. cross-domain graphs for sentiment analysis. In *RANLP*, pages 571–578, 2013.

41. Guang Qiu, Bing Liu, Jiajun Bu, and Chun Chen. Opinion word expansion and target extraction through double propagation. *Computational Linguistics*, 37(1):9–27, 2011.
42. Likun Qiu, Weishi Zhang, Changjian Hu, and Kai Zhao. Selc: a self-supervised model for sentiment classification. In *CIKM*, pages 929–936, 2009.
43. Andi Rexha, Mark Kröll, Mauro Dragoni, and Roman Kern. Exploiting propositions for opinion mining. In Harald Sack, Stefan Dietze, Anna Tordai, and Christoph Lange, editors, *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*, pages 121–125. Springer, 2016.
44. Andi Rexha, Mark Kröll, Mauro Dragoni, and Roman Kern. Polarity classification for target phrases in tweets: A word2vec approach. In Harald Sack, Giuseppe Rizzo, Nadine Steinmetz, Dunja Mladenic, Sören Auer, and Christoph Lange, editors, *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*, pages 217–223, 2016.
45. Ellen Riloff, Siddharth Patwardhan, and Janyce Wiebe. Feature subsumption for opinion analysis. In *EMNLP*, pages 440–448, 2006.
46. Swapna Somasundaran. *Discourse-level relations for Opinion Analysis*. PhD thesis, University of Pittsburgh, 2010.
47. Qi Su, Xinying Xu, Honglei Guo, Zhili Guo, Xian Wu, Xiaoxun Zhang, Bin Swen, and Zhong Su. Hidden sentiment association in chinese web opinion mining. In *WWW*, pages 959–968, 2008.
48. Maite Taboada, Julian Brooke, Milan Tofiloski, Kimberly D. Voll, and Manfred Stede. Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2):267–307, 2011.
49. Songbo Tan, Yuefen Wang, and Xueqi Cheng. Combining learn-based and lexicon-based techniques for sentiment detection without using labeled examples. In *SIGIR*, pages 743–744, 2008.
50. Hongling Wang and Guodong Zhou. Topic-driven multi-document summarization. In *IALP*, pages 195–198, 2010.
51. Q. F. Wang, E. Cambria, C. L. Liu, and A. Hussain. Common sense knowledge for handwritten chinese recognition. *Cognitive Computation*, 5(2):234–242, 2013.
52. Theresa Wilson, Janyce Wiebe, and Rebecca Hwa. Recognizing strong and weak opinion clauses. *Computational Intelligence*, 22(2):73–99, 2006.
53. Yuanbin Wu, Qi Zhang, Xuanjing Huang, and Lide Wu. Phrase dependency parsing for opinion mining. In *EMNLP*, pages 1533–1541, 2009.
54. Yasuhisa Yoshida, Tsutomu Hirao, Tomoharu Iwata, Masaaki Nagata, and Yuji Matsumoto. Transfer learning for multiple-domain sentiment analysis—identifying domain dependent/independent word polarity. In *AAAI*, pages 1286–1291, 2011.