

Explaining Sentiment from Lexicon

Sergio Consoli^a, Luca Barbaglia^a and Sebastiano Manzan^a

^aEuropean Commission, Joint Research Centre (JRC), Via E. Fermi 2749, I-21027 Ispra (VA), Italy

Abstract

Lexicon-based Sentiment Analysis relies on sentiment dictionaries which are used to assign a sentiment polarity to the words of an input text. The overall sentiment of the text is then computed by means of a combining function, such as the word count, sum or average. In this short contribution we describe a detailed set of linguistic rules that allow to understand the text fragments which are semantically linked to a given concept of interest in a text. These heuristics have been designed in the spirit of the recent *Interpretable AI* trend, since they allow to understand the origin of sentiment for a specific term, providing more transparency and interpretation of the resulting analysis, and enabling the development of advanced and novel lexicon-based Sentiment Analysis approaches, which is the object of our currently on-going work.

Keywords

Lexicon-based Sentiment Analysis, Natural Language Processing, Interpretability, Rule-based models

1. Introduction

The rapid advances in information and communications technology experienced in the last two decades have produced an explosive growth in the amount of information collected, leading to the new era of *Big Data* [1]. This has brought to the exponential increase in the information available in various domains, allowing for Natural Language Processing (NLP) and novel knowledge generation methods to emerge in different sectors. In particular, utilizing the sentiment extracted from social media has long been the tradition of several studies [2]. As the Web rapidly grows and evolves, people are becoming increasingly enthusiastic about interacting, sharing, and collaborating through social networks, online communities, blogs, and wikis [3]. Therefore, it is critical to correctly interpret sentiments and opinions expressed or reported about social events, political movements, company strategies, marketing campaigns, and any other form or online interaction.

Sentiment Analysis (SA) [4, 5], also known as Opinion Mining, is a Semantic Web technology, directly related to Natural Language Processing, that aims at understanding whether a certain textual message conveys a positive or negative sentiment with respect to a certain topic, or the overall contextual polarity or emotional reaction to a document, interaction, or event [4, 5]. Its outcome might be a quantitative/qualitative polarity (e.g., $[-1 : 1]$, *extr neg*, *neg*, *neut*, *pos*, *extr pos*, etc.) or an emotional state (e.g., *joy*, *anger*, etc.).

X-SENTIMENT: 6th International Workshop on eXplainable SENTiment Mining and EmotioN deTectioN, co-located at ESWC, June 07, 2021, Hersonissos, Greece

✉ sergio.consoli@ec.europa.eu (S. Consoli)

🆔 0000-0001-7357-5858 (S. Consoli); 0000-0001-5930-5392 (L. Barbaglia)

© 2021 Copyright for this paper by its authors.
Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

SA can be performed using both machine learning methods (see e.g., [6, 7, 8]) and lexicon-based methods (see e.g., [9, 10, 11, 12, 13]). Models driven by machine learning algorithms and vector representations have achieved top performance for various SA tasks. Although these models may get very accurate results, they provide however a limited understanding of patterns and features used to correctly classify the input text into sentiment categories. Therefore, these models lack transparency, traceability, and explainability on how the decisions are taken. In addition, another main disadvantage of machine learning models for sentiment analysis consists in their dependence on labelled data used for model training: it is not always easy to ensure that sufficient and correctly labelled data can be obtained for specific domains.

Conversely, lexicon-based approaches to SA are completely unsupervised and do not require any a-priori training corpus. They rely instead on dictionaries of words with assigned positive or negative sentiment polarity scores, also referred to as sentiment dictionaries or lexicons, like, for instance, SentiWordNet [14]¹, SenticNet [15, 16]², and Harvard IV-4³, just to name few popular ones. Most of these sentiment dictionaries are freely available online and have been often reused by the interested scientific and professional communities in several applications. Given a sentence, a lexicon-based SA approach works by assigning positive and negative sentiment polarity values from the dictionary to all the words in the sentence, and a combining function, such as the word count, sum or average [4], is used to aggregate the scores into the overall sentiment of the text. In this way, developers are relieved from collecting and labeling a large, relevant training corpus, at the minor cost of re-using an already existing sentiment dictionary, or constructing, if needed, a customized sentiment lexicon for a specific application. In addition, a lexicon-based SA method can be more easily understood and eventually modified by a human in comparison to a machine learning approach to SA, providing a significant advantage towards interpretability of the model results.

Most lexicon-based SA methods focus on a coarse-grained analysis of the sentiment expressed in the text [4], that is, they assess the entire sentiment of a sentence by considering all expressions of positive and negative sentiment contained in that text. However, coarse-grained methods might not be precise enough in evaluating the sentiment polarity of a specific concept of interest contained in a sentence, given that the sentiment of the entire text is often not expressed towards that specific topic [17].

In currently on-going research, we are investigating a fine-grained perspective to lexicon-based SA. In particular, we are interested in understanding the parts of the text which convey a sentiment connotation with respect to a specific concept of interest, and properly propagating these sentiments towards an overall computed sentiment score for the topic. While the entire approach is currently under development, we report here the set of linguistic polarity rules used to identify the text fragments semantically connected to a specific concept of interest within a sentence, possibly expressing a sentiment connotation towards it. We believe that explicit semantics can be leveraged to explain why a resource has been scored in a specific sentiment category, inducing trustworthiness and avoiding biases, and accompanying the current *model interpretability* trend in AI aiming at opening up the black-box by providing a narrative of the

¹SentiWordNet, version 3.0, available at: <https://github.com/aesuli/SentiWordNet>.

²SenticNet, available at: <https://sentic.net/>.

³The details of the latest version of the Harvard IV-4 dictionary are available at: <http://www.wjh.harvard.edu/~inquirer/homecat.htm>.

Table 1

Used spaCy part-of-speech tags.

TAG	POS	DESCRIPTION
CC	CONJ	conjunction, coordinating
IN	ADP	conjunction, subordinating or preposition
JJ	ADJ	adjective
JJR	ADJ	adjective, comparative
JJS	ADJ	adjective, superlative
MD	VERB	verb, modal auxiliary
NN	NOUN	noun, singular or mass
NNP	PROPN	noun, proper singular
NNPS	PROPN	noun, proper plural
NNS	NOUN	noun, plural
RBR	ADV	adverb, comparative
RBS	ADV	adverb, superlative
VB	VERB	verb

Table 2

Used spaCy dependency parsing classes.

DEP	DESCRIPTION
acl	clausal modifier of noun (adjectival clause)
advcl	adverbial clause modifier
advmod	adverbial modifier
amod	adjectival modifier
attr	attribute
dobj	direct object
neg	negation modifier
opr	object predicate
pcomp	complement of preposition
pobj	object of preposition
prep	prepositional modifier
xcomp	open clausal complement

underlying model [18, 19]. Our methodology under development goes into this direction.

2. Set of linguistic rules for lexicon-based Sentiment Analysis

We provide here details on the semantic rules used to detect the text fragments semantically connected to a specific concept of interest in an input text. These linguistic rules have been derived experimentally after in-depth natural language analysis [20], and are based on both syntax and semantics of the text. Each rule can be seen as a single building block, and the concatenation of these rules enable to explain how a particular sentiment polarity score is evaluated by an underlying lexicon-based SA algorithm, providing more transparency and in-




Figure 1: An illustration on the focused fragment of text for the sentence: *Despite Ronaldo's age, his physical shape looks perfect so far.*

interpretation of the resulting analysis.

The overall process is based on the linguistic features of the *spaCy*⁴ Python library. Tables 1-2 present the main labels assigned by *spaCy* with respect to part-of-speech tagging and dependency parsing, respectively, that we use in our rule scheme. Based on the **POS** (i.e., the detected part-of-speech), **DEP** (i.e., the parsed dependency), and **TAG** (i.e., the tag of the part-of-speech) labels defined in these tables, the algorithm selects a chunk in a sentence only if it contains a certain concept of interest (**Concept**) specified as input, and falls into one of the adopted semantic rules, detailed in the following.

Concept connected to a verb followed by an adjectival complement.

The **Concept** of interest is associated to a verb (**POS = VERB**) which is followed by an adjectival complement relation (**DEP = *acomp***), which means that it connects the verb to an adjectival term which functions as the complement (like an object of the verb) and offers more information about it. The adjective (**POS = ADJ**) can be in the form of:

- a. standard adjective (**TAG = JJ**);

Example (Figure 1): ...despite Ronaldo's age, his physical shape looks perfect so far...

(**Concept**=*shape* → **VERB**=*looks* ^{*acomp*} → **JJ**=*perfect*).

Concept connected to a verb associated to a noun.

In this case **Concept** is connected to a verb (**POS = VERB**) which is linked to a noun (**POS = NOUN**) by means of one of the following relations:

- a. direct object (**DEP = *dobj***), i.e. a clause which connects a transitive verb to a nominal representing the recipient of the action of such predicate;

Example: ...last year Michael received an award for his work...

(**Concept**=*Michael* → **VERB**=*received* ^{*dobj*} → **NOUN**=*award*).

⁴*spaCy*: Industrial-Strength Natural Language Processing in Python. Available at: <https://spacy.io/>.

- b. attribute (**DEP = attr**), i.e. a clause which connects a copula verb to a noun being the non-verb phrase predicate of such verb;

Example: ...in Catalonia taxation has been an heavy deterrent on the development of SMEs...

(**Concept**=*taxation*→**VERB**=*been*^{attr}→**NOUN**=*deterrent*).

Concept connected to a verb followed by an adverbial modifier.

The **Concept** of interest is associated to a verb (**POS = VERB**) which is followed by an adverbial modifier relation (**DEP = advmod**), which means that it connects the verb to a non-clausal adverb or adverbial phrase that serves to modify the predicate. The adverb (**POS = ADV**) can be in the form of:

- a. comparative adverb (**TAG = RBR**);

Example: ...his power will decline further...

(**Concept**=*power*→**VERB**=*decline*^{advmod}→**RBR**=*further*).

- b. superlative adverb (**TAG = RBS**);

Example: ...Tim's attention is best focused on one thing: football!

(**Concept**=*attention*→**VERB**=*focused*^{advmod}→**RBS**=*best*).

Concept connected to a verb followed by an object predicate.

The **Concept** of interest is associated to a verb (**POS = VERB**) which is followed by an object predicate relation (**DEP = oprd**), that is a non-verb phrase predicate in a small clause that functions like the predicate of an object. It means, in other words, that the linked verb is connected to an adjective that qualifies, describes, or renames the object that appears before it. The adjective (**POS = ADJ**) can be in the form of:

- a. standard adjective (**TAG = JJ**);

Example: ...the law has been declared unconstitutional...

(**Concept**=*law*→**VERB**=*declared*^{oprd}→**JJ**=*unconstitutional*).

- b. comparative adjective (**TAG = JJR**);

Example: ...the ECB kept the rates lower than expected...

(**Concept**=*ECB*→**VERB**=*kept*^{oprd}→**JJR**=*lower*).

- c. superlative adjective (**TAG = JJS**);

Example: ...the FED is keeping asset prices the lowest...

(**Concept**=*FED*→**VERB**=*keeping*^{oprd}→**JJS**=*lowest*).

Concept connected to a verb followed by a prepositional modifier.

The **Concept** of interest is connected to a verb (**POS = VERB**) which is followed by a prepositional modifier relation (**DEP = prep**), that is a prepositional phrase that modifies the heading verb. The prepositional modifier is linked to an adposition (**POS = ADP**), which basically establishes a grammatical relationship that links its complement to another word or phrase in the context. An adposition typically establishes a semantic relationship which may be spatial (in, on, under, ...), temporal (after, during, ...), or of some other type (of, for, via, ...). The adposition is then connected to one of the following terms:

- a. a noun (**POS = NOUN**), by means of an object of preposition relation (**DEP = pobj**), i.e. a noun phrase that follows a preposition and completes its meaning;

Example: ...currently Ronaldo is in great shape...

(**Concept**=Ronaldo→**VERB**=is^{prep}→**ADP**=in^{pobj}→**NOUN**=shape).

- b. a verb (**POS = VERB**), by means of a complement of preposition relation (**DEP = pcomp**), i.e. a clause which is not a **pobj** and directly connects the preposition with any dependent completing its meaning;

Example: ...The firm is lowering its profits after paying 1 million euros to the tax office...

(**Concept**=firm→**VERB**=lowering^{prep}→**ADP**=after^{pcomp}→**NOUN**=paying).

Concept connected to a verb followed by an open clausal complement or an adverbial clause modifier.

The **Concept** of interest is in this case associated to a verb (**POS = VERB**) which is connected to another term by means of one of the following relations:

- a. open clausal complement (**DEP = xcomp**), i.e. a predicative or clausal complement without its own subject;

Example: ...news reported that the the FTSE index could keep losing...

(**Concept**=FTSE index→**VERB**=keep^{xcomp}→**VERB**=losing).

- b. adverbial clause modifier (**DEP = advcl**), i.e. a clause which modifies a verb or another predicate (adjective, etc.) as a modifier, not as a core complement;

Example: ...industrial production will reach the bottom as it struggles with the current crisis given by the pandemic...

(**Concept**=industrial production→**VERB**=reach^{advcl}→**VERB**=struggle).

Concept associated to an adjectival clause.

The **Concept** of interest is connected by an adjectival clause (**DEP = acl**), i.e. a finite or non-finite clause that modifies **Concept**, to a term being:

- a. a verb (**POS = VERB**);
Example: ...there exist many tools providing similar benefits...
 (**Concept**=*tools*^{acl}→**VERB**=*providing*).

Concept associated to an adjectival modifier.

In this case **Concept** is connected by an adjectival modifier relation (i.e., **DEP = amod**), i.e. an adjective phrase that modifies the meaning of the **Concept** of interest, to a term being:

- a. an adjective (**POS = ADJ**), in the form of:
- i. standard adjective (**TAG = JJ**);
Example: ...his top performance is encouraging the rest of the team...
 (**Concept**=*performance*^{amod}→**JJ**=*top*).
 - ii. comparative adjective (**TAG = JJR**);
Example: ...there is a larger consumption than in the past years...
 (**Concept**=*consumption*^{amod}→**JJR**=*larger*).
 - iii. superlative adjective (**TAG = JJS**);
Example: ...the manufacturing sector is experiencing the worst decline since World War II...
 (**Concept**=*decline*^{amod}→**JJS**=*worst*).
- b. a verb (**POS = VERB**);
Example: ...overall it appears to be an encouraging agreement...
 (**Concept**=*agreement*^{amod}→**VERB**=*encouraging*).

3. Conclusion and Outlook

This is on-going research. We are aware that the presented linguistic rules are heuristics derived experimentally after in-depth natural language analysis. They further require a rigorous testing, which we aim in our currently on-going work.

In the near future we are planning to implement an advanced lexicon-based SA method leveraging on the described linguistic rules, aiming also at an in-depth performance comparison against other popular SA approaches. This strategy under current development will focus in particular on the economic and financial domains [21], with the goal of providing useful signals to improve forecasting and nowcasting of economic and financial indicators.

4. Acknowledgments

The authors would like to thank the colleagues of the Centre for Advanced Studies at the Joint Research Centre of the European Commission for helpful guidance and support during the development of this research work. The views expressed are purely those of the authors

and may not in any circumstance be regarded as stating an official position of the European Commission. The author Dr. Sergio Consoli is also particularly grateful to the Wagner team in Catania, Italy, for inspiring discussions, precious advices, and support.

References

- [1] V. Marx, The big challenges of big data, *Nature* 498 (2013) 255–260.
- [2] K. Ravi, V. Ravi, A survey on opinion mining and sentiment analysis: Tasks, approaches and applications, *Knowledge-Based Systems* 89 (2015) 14–46.
- [3] F. Neri, C. Aliprandi, F. Capeci, M. Cuadros, T. By, Sentiment analysis on social media, in: *Proceedings of the 2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2012*, 2012, pp. 919–926.
- [4] B. Liu, *Sentiment analysis: Mining opinions, sentiments, and emotions*, Cambridge University Press, 2015.
- [5] E. Cambria, S. Poria, A. Gelbukh, M. Thelwall, Sentiment analysis is a big suitcase, *IEEE Intelligent Systems* 32 (2017) 74–80.
- [6] M. Neethu, R. Rajasree, Sentiment analysis in twitter using machine learning techniques, in: *2013 4th International Conference on Computing, Communications and Networking Technologies, ICCCNT 2013*, 2013, pp. 1–5.
- [7] L. Zhang, S. Wang, B. Liu, Deep learning for sentiment analysis: A survey, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 8 (2018).
- [8] A. Tripathy, A. Agrawal, S. Rath, Classification of sentiment reviews using n-gram machine learning approach, *Expert Systems with Applications* 57 (2016) 117–126.
- [9] C. Khoo, S. Johnkhan, Lexicon-based sentiment analysis: Comparative evaluation of six sentiment lexicons, *Journal of Information Science* 44 (2018) 491–511.
- [10] M. Mostafa, More than words: Social networks’ text mining for consumer brand sentiments, *Expert Systems with Applications* 40 (2013) 4241–4251.
- [11] D. Reforgiato Recupero, V. Presutti, S. Consoli, A. Gangemi, A. G. Nuzzolese, Sentilo: Frame-based sentiment analysis, *Cognitive Computation* 7 (2015) 211–225.
- [12] M. Taboada, J. Brooke, M. Tofiloski, K. Voll, M. Stede, Lexicon-based methods for sentiment analysis, *Computational Linguistics* 37 (2011) 267–307.
- [13] D. Reforgiato Recupero, S. Consoli, A. Gangemi, A. Nuzzolese, D. Spampinato, A semantic web based core engine to efficiently perform sentiment analysis, in: *Lecture Notes in Computer Science*, volume 8798, 2014, pp. 245–248.
- [14] S. Baccianella, A. Esuli, F. Sebastiani, SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining, in: *LREC*, volume 10, 2010, pp. 2200–2204.
- [15] E. Cambria, R. Speer, C. Havasi, A. Hussain, Senticnet: A publicly available semantic resource for opinion mining, in: *2010 AAAI Fall Symposium Series*, volume FS-10-02, 2010, pp. 1–5.
- [16] E. Cambria, Y. Li, F. Z. Xing, S. Poria, K. Kwok, Ensemble application of symbolic and subsymbolic AI for sentiment analysis, in: *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM ’20)*, 2020, pp. 105–114.
- [17] M. Van De Kauter, D. Breesch, V. Hoste, Fine-grained analysis of explicit and implicit

- sentiment in financial news articles, *Expert Systems with Applications* 42 (2015) 4999–5010.
- [18] T. Kim, B. Routledge, Informational Privacy, A Right to Explanation, and Interpretable AI, in: *Proceedings - 2018 2nd IEEE Symposium on Privacy-Aware Computing, PAC 2018*, 2018, pp. 64–74.
- [19] L. Gilpin, D. Bau, B. Yuan, A. Bajwa, M. Specter, L. Kagal, Explaining explanations: An overview of interpretability of machine learning, in: *Proceedings - 2018 IEEE 5th International Conference on Data Science and Advanced Analytics, DSAA 2018*, 2019, pp. 80–89.
- [20] D. Reforgiato Recupero, A. Nuzzolese, S. Consoli, V. Presutti, S. Peroni, M. Mongiovì, Extracting knowledge from text using sheldon, a semantic holistic framework for linked ontology data, in: *WWW 2015 Companion - Proceedings of the 24th International Conference on World Wide Web*, 2015, pp. 235–238.
- [21] L. Barbaglia, S. Consoli, S. Manzan, Monitoring the business cycle with fine-grained, aspect-based sentiment extraction from news, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 11985 LNAI (2020) 101–106.