

Sentiment Analysis is a Big Suitcase

Erik Cambria and Soujanya Poria, Nanyang Technological University
Alexander Gelbukh, Instituto Politécnico Nacional
Mike Thelwall, University of Wolverhampton

With the recent development of deep learning, research in artificial intelligence (AI) has gained new vigor and prominence. Machine learning, however, suffers from three big issues, namely:

1. Dependency: it requires (a lot of) training data and is domain-dependent;
2. Consistency: different training or tweaking leads to different results;
3. Transparency: the reasoning process is unintelligible (black-box algorithms).

In the context of natural language processing (NLP), these issues are particularly crucial because, unlike in other fields, they prevent AI from achieving human-like performance. To this end, AI researchers need to bridge the gap between statistical NLP and many other disciplines that are necessary for understanding human language, such as linguistics, commonsense reasoning, and affective computing. They will have to develop an approach to NLP that is both top-down and bottom-up: top-down for the fact that they should leverage symbolic models such as semantic networks and conceptual dependency representations to encode meaning; bottom-up because they should use sub-symbolic methods such as deep neural networks and multiple kernel learning to infer syntactic patterns from data.

Coupling symbolic and sub-symbolic AI is key for stepping forward to the path from NLP to natural language understanding. Relying solely on machine learning, in fact, is simply useful to make a ‘good guess’ based on past experience, because sub-symbolic methods only encode correlation and their decision-making process is merely probabilistic. Natural language understanding, however, requires much more than that. To use Noam Chomsky’s words, “you do not get discoveries in the sciences by taking huge amounts of data, throwing them into a computer and doing statistical analysis of them: that’s not the way you understand things, you have to have theoretical insights”.

It is necessary to take a holistic approach to sentiment analysis by handling the many sub-problems involved in extracting meaning and polarity from text. While most works approach it as a simple categorization problem, in fact, sentiment analysis is actually a suitcase research problem that requires tackling many NLP tasks. As Marvin Minsky would say, the expression ‘sentiment analysis’ itself is a big suitcase (like many others related to affective computing, e.g., emotion recognition or opinion mining) that all of us use to encapsulate our jumbled idea about how our minds convey emotions and opinions through natural language.

We propose to address the composite nature of the problem via a three-layer structure inspired by the ‘Jumping NLP Curves’ paradigm [3] (Fig. 1). In particular, we argue that there are (at least) 15 NLP problems that need to be solved to achieve human-like performance in sentiment analysis. Such NLP problems are organized into three layers:

Syntactics Layer, Semantics Layer, and Pragmatics Layer. The sequence of the different modules is just indicative and may require re-ordering depending on the specific data or domain that is being processed. For example, named-entity recognition (NER) may have to be performed before lemmatization as some named entities would not be recognizable after lemmatization, e.g., “Guns N’ Roses” versus “gun n rose”.




Fig. 1. Jumping NLP curves. Borrowed from the field of business management and marketing prediction, this paradigm reinterprets the evolution of NLP research as the intersection of three overlapping curves, which will eventually lead NLP research to evolve into natural language understanding.

Syntactics Layer

The main aim of the Syntactics Layer is to pre-process text so that informal text is reduced to plain English, inflected forms of verbs and nouns are normalized, and basic sentence structure is explicated.

Microtext Analysis

The proliferation of Social Web technologies and the increasing use of computer-mediated communication has resulted in a new form of informal written text, termed microtext, which is characterized by relaxed spelling and reliance on abbreviations, acronyms, and emoticons. This is partly a consequence of Zipf’s law, or principle of least effort (for which people tend to minimize energy cost at both individual and collective levels when communicating with one another), and it poses new challenges for NLP tools which are

usually designed for well-written text. The first step in tackling the challenge of developing algorithms to correct the non-standard vocabulary found in microtexts is to realize that the number of different spelling variations may be massive but they follow a small number of simple basic strategies, such as abbreviation and phonetic substitution. While most of the literature on microtext analysis exploits supervised learning, unsupervised approaches have recently gained increasing popularity [1], as microtext evolves too quickly to construct a comprehensive set of training data.

Sentence Boundary Disambiguation

In any document-level NLP task, sentence boundary disambiguation (SBD) is an important subtask that involves deconstructing text into sentences. SBD is particularly tricky when sentence boundary identifiers are not clearly defined, such as in the presence of emoticons. In fact, the accuracy of SBD can be improved after performing microtext normalization. Early works employed decision trees to identify whether the presence of a full stop in text indicates the boundary of a sentence. Later, Reynar et al. [17] employed maximum entropy learning to create a sentence segmentation classifier that considered sentence boundary detection as a boundary disambiguation task, where every token containing ‘!’, ‘.’ or ‘?’ was a potential sentence boundary.

POS Tagging

Part-of-speech (POS) tagging is a fundamental NLP task that labels each word by its part of speech, e.g., adjective, verb, and noun. Most of the existing works consider POS tagging as sequence labeling task. The WSJ-PTB (the Wall Street Journal part of the Penn Treebank Dataset) corpus contains 1.17 million tokens and has been widely used for developing and evaluating POS tagging systems. One characteristic of the POS tagging problem is strong dependency between adjacent tags. With a simple left-to-right tagging scheme, it is possible to model dependencies between adjacent tags only by feature engineering. In an effort to reduce feature engineering, Huang et al. [5] concatenated word embeddings and manually-designed word-level features and employed a bidirectional long short-term memory (LSTM) network to model wider context arbitrarily.

Text Chunking

Text chunking, also known as shallow parsing, is a task that follows POS tagging and that adds more structure to the sentence. The result is a grouping of the words in simple syntactic substructures (chunks), e.g., noun groups or verb argument structure. Unlike full syntactic parsing, which implies building a deeply nested structure, text chunking splits sentences into non-overlapping groups of words that represent a syntactic unit, without attempting to identify their internal structure or identify their relationships between each other. An example of a chunked sentence is: “[NP *European Commission*] [VP *will probably launch*] [NP *a legal case*] [PP *against*] [NP *Poland*]”. Typically, a chunker uses many more superficial text analysis techniques than a full syntactic parser, and thus is much simpler, much more robust, and

less resource-intensive. In particular, the chunks are usually unambiguous and do not depend on the choice of a syntactic formalism. Grouping words into chunks permits further identification of important basic relations between words, such as concept names and basic syntactic roles. Common approaches to chunking include rule-based and machine-learning methods, such as transformation-based learning [16].

Lemmatization

Lemmatization is the process of converting a given word into a base form, i.e., a morphologically correct root form. This is useful for detecting a concept like `eat_burger` from all its possible inflected forms, e.g., `ate_burger`, `eat_burgers`, `eating_burgers`, `eaten_burger`, etc. Unlike stemming, lemmatization is meaning-preserving as it does not chunk away suffixes, regardless of semantics, and it preserves the word's POS tag. While a lemmatizer would reduce plurals to singular and inflected forms of verbs to their infinitive form, a stemmer would brutally remove morphological and inflexional endings from words. For example, it would reduce the words `democrats` (noun), `democratic` (adjective) and `democratize` (verb) to the same root (`democrat`). A recent approach to lemmatization is a deep learning method by Kestemont et al. [6], who applied temporal convolutions to model the orthography of input words at the character level and used distributional word embeddings to represent the lexical context surrounding the input words.

Semantics Layer

The main aim of the Semantics Layer is to deconstruct the normalized text obtained from the Syntactics Layer into concepts, resolve references, i.e., named entities and anaphora, and filter out neutral content from the input in order to improve the accuracy of sentiment classification.

Word Sense Disambiguation

A word can have multiple meanings depending on its context. The task of word sense disambiguation (WSD) consists in recognizing which sense of an ambiguous word is used in the input sentence and it is key to improving the accuracy of concept extraction. WSD can benefit from POS tagging for disambiguating some word senses, e.g., `fine` as noun (penalty) or `fine` as an adjective (good). Sometimes, however, POS tags are just not enough, e.g., the word `train` can be disambiguated as a verb but not as a noun (i.e., `train` as in 'high-speed train' versus `train` as in 'train of elephants'). Lee and Ng [7] evaluated a variety of knowledge sources and supervised learning algorithms for WSD on SENSEVAL-2 and SENSEVAL-1 data. Their training data consisted of sentences having ambiguous words with their sense tagged manually. The knowledge sources used were POS tags of neighboring words, unigrams in the surrounding context, local collocations, and syntactic relations.



Fig. 2. Sentiment analysis' big suitcase of NLP problems. For long, sentiment analysis has been mistaken for the task of polarity detection. This, however, is just one of the many NLP problems that need to be solved to achieve human-like performance in sentiment analysis.

Concept Extraction

Deconstructing text into concepts is key for a semantic-aware analysis of text. Concepts can be either single words or multiword expressions, the latter being semantic atoms that should never be broken down into single words. The concept `pain_killer`, for example, should not be split into `pain` and `killer`, which are two words with completely different semantics (and polarity). State-of-the-art approaches can be classified into two main categories: statistical methods and linguistic rules. The former usually leverage term frequency and word location to calculate term weighting, the latter often use POS tagging and text chunking to extract meaningful n-grams. A recent approach [2] leveraged conceptual primitives automatically discovered by means of hierarchical clustering and dimensionality reduction. Verb concepts like `eat`, `slurp`, and `munch` were represented by their conceptual primitive `INGEST` and noun concepts like `pasta`, `noodles`, and `toast` were replaced with their ontological parent `FOOD`. This way, any verb-noun combination between such concepts, e.g., `eat_pasta`, `slurp_noodles`, or `munch_toast`, would be generalized and extracted as `INGEST_FOOD`.

Named Entity Recognition

NER is a subtask of information extraction that aims to locate and classify named entities into pre-defined categories. CoNLL 2003 is the standard English dataset for NER and concentrates on four types of named entities: persons, locations, organizations and miscellaneous. NER is important for many other NLP tasks involved in sentiment analysis, especially anaphora resolution and aspect extraction. NER is often broken down as two distinct problems: detection of names, typically simplified to a segmentation problem, and classification of the names by the type of entity they refer to. While early approaches mostly leveraged a domain-specific lexicon, recent works mostly employ deep learning on a training set. [8] presented a label embedding method that incorporates prototypical and hierarchical information to learn pre-trained embeddings and adapted a zero-shot learning framework that can predict both seen and previously unseen entity types. NER includes other NLP sub-tasks, such as temporal tagging, where general heuristic rules are usually employed to recognize time expressions in text [20] and which is sometimes more difficult than NER itself because, while standard named entities are usually expressed using formal language, temporal entities are subject to short expressions, small vocabulary, recurrence, and similar syntactic behaviors.

Anaphora Resolution

Anaphora can be defined as the presupposition that points back to a previous item. The pointing back reference is called an anaphor and the entity to which it refers is its antecedent. The process of determining the antecedent of an anaphor is called anaphora resolution, which is still an open NLP challenge that needs to be tackled in many domains, including machine translation, summarization, and question-answering. In the context of sentiment analysis, this task is key to resolve pronouns in text before subjectivity detection or aspect extraction can be applied. The most widespread types of anaphora are: pronominal anaphora, which is realized by anaphoric pronouns; adjectival anaphora, realized by anaphoric possessive adjectives; and one-anaphora, the anaphoric expression is realized by a "one" noun phrase. When resolving anaphora, some constraints must be respected: number agreement (to distinguish between singular and plural references); gender agreement (to distinguish between male, female, and neutral genders); and semantic consistency (it is assumed that both the antecedent clause and the one containing the anaphora are semantically consistent). Grammatical, syntactic or pragmatic rules have been widely used in the literature to identify the antecedent of an anaphor. While early works focused on the use of parse trees, discourse models, POS tagging, and lexical databases, more recent approaches have leveraged genetic algorithms [11].

Subjectivity Detection

Subjectivity detection is a NLP task that aims to remove 'factual' or 'neutral' content, i.e., objective text that does not contain any opinion, from online product reviews. This pre-processing step is crucial to increase the accuracy of sentiment analysis systems, as these are usually optimized for the binary classification task of distinguishing between

positive and negative content. Previous methods used well-established general subjectivity clues to generate training data from un-annotated text. In addition, features such as pronouns, modals, adjectives, cardinal numbers, and adverbs have shown to be effective in subjectivity classification. Some existing resources contain lists of subjective words, and some empirical methods in NLP have automatically identified adjectives, verbs, and n-grams that are statistically associated with subjective language. However, several subjective words such as ‘unseemingly’ occur infrequently and consequently a large training dataset is necessary to build a broad and comprehensive subjectivity detection system. Recently, Chaturvedi et al. [4] proposed a novel framework that exploits the features of both Bayesian networks and fuzzy recurrent neural networks for filtering out neutral content in a time- and resource-effective manner.

Pragmatics Layer

The main aim of the Pragmatics Layer is to extract meaning from both sentence structure and semantics obtained from Syntactics Layer and Semantics Layer, respectively. After performing some kind of user profiling (personality and sarcasm detection), the Pragmatics Layer interprets metaphors (if any) and extracts opinion targets and the polarity associated with each of them.

Personality Recognition

Personality is a combination of behavior, emotion, motivation, and thought pattern characteristics of an individual. The automatic detection of personality traits has many important practical applications. In recommender systems, for example, the products and services recommended to a person should be those that have been positively evaluated by other users with a similar personality type. The Big Five is the most widely accepted model of personality. Mairesse et al. [9] used the Linguistic Inquiry and Word Count (LIWC) dataset and other features, such as imageability, to conduct automated personality detection on an essay dataset. Recently, [10] presented a deep learning method to determine personality type from stream-of-consciousness essays by detecting the presence or absence of the Big Five traits in the author’s psychological profile. For each of the five traits, they trained a separate binary classifier, with identical architecture, based on a novel document representation technique. The classifier is implemented as a specially designed deep convolutional neural network (CNN), with injection of the document-level Mairesse features, extracted directly from the text, into an inner layer. The first layers of the network treat each sentence of the text separately, then sentences are aggregated into a document vector.

Sarcasm Detection

Sarcasm is always directed at someone or something. A target of sarcasm is the person or object against whom or which the utterance is directed. Targets can be the sender,

the addressee or a third party (or a combination of the three). The presence of sarcastic sentences may completely change the meaning of a review, therefore misleading the interpretation of its polarity. While the use of irony and sarcasm is well studied from its linguistic and psychologic aspects, sarcasm detection is still represented by very few works in the computational literature. To date, most approaches to sarcasm detection have treated the task primarily as a text categorization problem. Sarcasm, however, can be expressed in very subtle ways and requires a deeper understanding of natural language that standard text categorization techniques cannot grasp. [12] suggested a theoretical framework in which the context of sentiment words shifts the valence of the expressed sentiment. This is made on the assumption that, though most salient clues about attitude are provided by the lexical choice of the writer, the organization of the text also provides relevant information for assessing attitude. More recently, [14] developed deep models based on a pre-trained CNN for extracting sentiment, emotion and personality features for sarcasm detection.

Metaphor Understanding

Metaphors are commonly used to substitute complex concepts with simple concepts which bear similar ideas but are not literally applicable. In the context of sentiment analysis, metaphor detection and understanding are necessary for aspect extraction and polarity detection. The aspects of an undetected metaphor, in fact, could be classified as off-topic and, hence, reduce the accuracy of sentiment classification. Early works employed hand-crafted rules and knowledge bases for metaphor detection, but they suffered from scalability issues. More recent approaches tried to leverage either conceptual metaphor mappings or selectional preferences. Both require extensive knowledge of the mappings/preferences in question, as well as sufficient data for all involved conceptual domains. Creating these resources is expensive and often limits the scope of these systems. Recently, Hovy et al. [18] proposed a statistical approach to metaphor detection that requires no knowledge of semantic concepts or the metaphor's source domain. The model utilizes the rarity of novel metaphors, marking words that do not match a text's typical vocabulary as metaphor candidates.

Aspect Extraction

In opinion mining, different levels of analysis granularity have been proposed, each having advantages and drawbacks. Aspect-based sentiment analysis focuses on the relations between aspects and document polarity. Aspects are opinion targets, i.e., the specific features of a product or service that users like or dislike. For example, the sentence "The screen of my phone is really nice and its resolution is superb" expresses a positive polarity about the phone under review. More specifically, the opinion holder is expressing a positive polarity about its *screen* and *resolution*; these concepts are thus called opinion targets, or aspects. It is important to identify aspects because reviewers may express opposite polarities about different aspects in the same sentence. Early approaches to aspect extraction included linguistic rules based on statistical observations, pointwise mutual information between noun phrase and product class, topic modeling, and more. Recently, Poria et al. [15] proposed a hybrid approach to aspect extraction

which consisted in the use of a 7-layer deep CNN, for tagging each word in opinionated sentences as either aspect or non-aspect word, in concomitance with aspect-specific linguistic patterns.

Polarity Detection

Polarity detection is the most popular sentiment analysis task. In fact, many research works even use the terms ‘polarity detection’ and ‘sentiment analysis’ interchangeably. This is due to the (limited) definition of sentiment analysis as the NLP task that aims to categorize a piece of text as either positive or negative. In fact, early approaches simply focused on this binary classification (often ignoring the presence of neutral content) by employing knowledge bases, rule-based classification, and supervised learning. Later works focused on a finer-grained categorization that included measuring the intensity of the polarity detected [19]. Recent approaches to polarity detection include deep learning techniques, e.g., CNNs and LSTMs relying on constituency parsing trees, and hybrid frameworks that leverage an ensemble of linguistics, knowledge representation, and machine learning to achieve a better understanding of the contextual role of each concept within the sentence, by allowing sentiments to flow from concept to concept based on dependency relations [13].

Erik Cambria is an assistant professor in the School of Computer Science and Engineering at Nanyang Technological University and founder of SenticNet. Contact him at cambria@ntu.edu.sg.

Soujanya Poria is a research scientist at Temasek Laboratories at Nanyang Technological University. Contact him at sporia@ntu.edu.sg.

Alexander Gelbukh is a research professor at the CIC of the Instituto Politécnico Nacional. Contact him at gelbukh@cic.ipn.mx.

Mike Thelwall is a professor of Information Science at University of Wolverhampton. Contact him at m.thelwall@wlv.ac.uk.

References

1. Thales Bertaglia and Maria das Graças. Exploring word embeddings for unsupervised textual user-generated content normalization. In *COLING*, pages 112–120, 2016.
2. Erik Cambria, Soujanya Poria, Rajiv Bajpai, and Björn Schuller. SenticNet 4: A semantic resource for sentiment analysis based on conceptual primitives. In *COLING*, pages 2666–2677, 2016.
3. Erik Cambria and Bebo White. Jumping NLP curves: A review of natural language processing research. *IEEE Computational Intelligence Magazine*, 9(2):48–57, 2014.
4. Iti Chaturvedi, Edoardo Ragusa, Paolo Gastaldo, Rodolfo Zunino, and Erik Cambria. Bayesian network based extreme learning machine for subjectivity detection. *Journal of The Franklin Institute*, 2017.

5. Zhiheng Huang, Wei Xu, and Kai Yu. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*, 2015.
6. Mike Kestemont, Guy de Pauw, Renske van Nie, and Walter Daelemans. Lemmatization for variation-rich languages using deep learning. *Literary and Linguistic Computing*, page fqw034, 2016.
7. Yoong Keok Lee and Hwee Tou Ng. An empirical evaluation of knowledge sources and learning algorithms for word sense disambiguation. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*, pages 41–48. Association for Computational Linguistics, 2002.
8. Yukun Ma, Erik Cambria, and Sa Gao. Label embedding for zero-shot fine-grained named entity typing. In *COLING*, pages 171–180, Osaka, 2016.
9. François Mairesse, Marilyn Walker, Matthias Mehl, and Roger Moore. Using linguistic cues for the automatic recognition of personality in conversation and text. *Journal of Artificial Intelligence Research (JAIR)*, 30:457–500, 2007.
10. Navonil Majumder, Soujanya Poria, Alexander Gelbukh, and Erik Cambria. Deep learning-based document modeling for personality detection from text. *IEEE Intelligent Systems*, 32(2):74–79, 2017.
11. Ruslan Mitkov, Richard Evans, and Constantin Orasan. A new, fully automatic version of mitkov’s knowledge-poor pronoun resolution method. In *Computational Linguistics and Intelligent Text Processing*, pages 168–186. Springer, 2002.
12. Livia Polanyi and Annie Zaenen. Contextual valence shifters. In *Computing attitude and affect in text: Theory and applications*, pages 1–10. Springer, 2006.
13. Soujanya Poria, Erik Cambria, Alexander Gelbukh, Federica Bisio, and Amir Hussain. Sentiment data flow analysis by means of dynamic linguistic patterns. *IEEE Computational Intelligence Magazine*, 10(4):26–36, 2015.
14. Soujanya Poria, Erik Cambria, Devamanyu Hazarika, and Prateek Vij. A deeper look into sarcastic tweets using deep convolutional neural networks. In *COLING*, pages 1601–1612, 2016.
15. Soujanya Poria, Iti Chaturvedi, Erik Cambria, and Amir Hussain. Convolutional MKL based multimodal emotion recognition and sentiment analysis. In *Proceedings of ICDM*, Barcelona, 2016.
16. Lance A Ramshaw and Mitchell P Marcus. Text chunking using transformation-based learning. In *Natural language processing using very large corpora*, pages 157–176. Springer, 1999.
17. Jeffrey C Reynar and Adwait Ratnaparkhi. A maximum entropy approach to identifying sentence boundaries. In *Proceedings of the fifth conference on Applied natural language processing*, pages 16–19. Association for Computational Linguistics, 1997.
18. Marc Schulder and Eduard Hovy. Metaphor detection through term relevance. In *Proceedings of the Second Workshop on Metaphor in NLP*, pages 18–26, 2014.
19. Mike Thelwall. The heart and soul of the web? sentiment strength detection in the social web with sentistrength. In *Book Cyberemotions Pages*, pages 119–134, 2017.
20. Xiaoshi Zhong, Aixin Sun, and Erik Cambria. Time expression analysis and recognition using syntactic token types and general heuristic rules. In *ACL*, pages 420–429, 2017.