



FreMPHONE: A French Mobile Phone Corpus for Aspect-Based Sentiment Analysis

Sarsabene Hammi^(✉), Souha Mezghani Hammami, and Lamia Hadrach Belguith

ANLP Research Group, MIRACL Lab, FSEGS, University of Sfax, Sfax, Tunisia
souha.hammi@ihecs.usf.tn, lamia.belguith@fsegs.usf.tn

Abstract. Aspect-Based Sentiment Analysis (ABSA) task is one of the Natural Language Processing (NLP) research fields that has seen considerable scientific advancements over the last few years. This task aims to detect, in a given text, the sentiment of users towards the different aspects of a product or service. Despite the big number of annotated corpora that have been produced to perform the ABSA task in the English language, resources are still stingy, for other languages. Due to the lack of French corpora created for the ABSA task, we present in this paper the French corpus for the mobile phone domain “FreMPHONE”. The constructed corpus consists of 5217 mobile phone reviews collected from the Amazon.fr website. Each review in the corpus was annotated with its appropriate aspect terms and sentiment polarity, using an annotation guideline. The FreMPHONE contains 19257 aspect terms divided into 13259 positives, 5084 negatives, and 914 neutrals. Moreover, we proposed a new architecture “CBCF” that combines the deep learning models (LSTM and CNN) and the machine learning model CRF and we evaluated it on the FreMPHONE corpus. The experiments were performed on the subtasks of the ABSA: Aspect Extraction (AE) and Sentiment Classification (SC). These experiments showed the good performance of the CBCF architecture which overrode the LSTM, CNN, and CRF models and achieved an F-measure value equal to 95.96% for the Aspect Extraction (AE) task and 96.35% for the Sentiment Classification (SC) task.

Keywords: Aspect-based Sentiment Analysis · Aspect Extraction · CBCF · FreMPHONE · Sentiment Analysis

1 Introduction

In recent years, online platforms and commercial websites have witnessed a significant increase in the mass of exchanged information. This huge amount of information made the major trading houses scramble to obtain this information and exploit it to achieve their commercial and political interests. In fact, comprehending what customers think about services, products, political personalities, etc., can significantly influence the changing of the course of events for the better, including the enhancement of commercial products or changing of political strategies, etc. With the growing desire to analyze this electronic content, researchers [1] put forward, for the first time in 2004, a new scientific task that

they called sentiment analysis (SA). This task aims to detect from the online available content the sentiment of customers about someone or something [1]. However, the sentiment analysis task may not be enough in some cases, especially with the increase of requirements of companies. In fact, some companies tend to know not only the opinions of customers about their industrialized products but also about the different aspects of these products. From here, came the Aspect-Based Sentiment Analysis (ABSA) task which provides more fine-grained sentiment analysis. This task enables companies to improve the different aspects that customers did not like in a product and thus ameliorate the satisfaction rate. Taking the following review “*I fell in love with this camera*”, the ABSA task aims to identify the positive sentiment expressed for the aspect term “**camera**” of the entity “**phone**”.

Early studies have divided the ABSA task into two sub-tasks: The Aspect Extraction (AE) task and the Sentiment Classification (SC) task. The AE task aims to extract the aspect terms in which customers express their opinion about them (“**quality**”, “**screen**”, “**facial recognition**”, etc.), while the SC tends to assign sentiment polarity (“**positive**”, “**negative**” or “**neutral**”) to the aspects extracted in the AE task.

Since the accomplishment of any ABSA task needs the use of a well-annotated corpus, we found many corpora that have been constructed for this purpose in the English language. Among these corpora, we mention the Amazon Customer Reviews corpus [1], TripAdvisor corpus [2], Epinions corpus [3], SemEval-2014 ABSA corpus [4], etc. However, the number of publicly available French corpora for the ABSA task is still scarce. For this reason, we propose in this study the first available Mobile Phone corpus “FreMPhone” for the ABSA task. FreMPhone corpus serves as a complementary resource to existing French resources and it is composed of 19257 aspect terms. In addition, we introduce a new neural network architecture CBCF to evaluate the reliability of the FreMPhone corpus. This architecture combines the CNN, Bi-LSTM, and CRF models. It uses first the Bi-LSTM and CNN models to select the pertinent features for classification and then employs the CRF model to detect the appropriate label for each word. This architecture surpasses the LSTM, CNN, and CRF models and achieves encouraging results for both AE and SC tasks.

The rest of the paper is organized as follows. Section 2 presents the ABSA-related corpora constructed for the French language. Section 3 describes the process of data collection, the annotation manner, and the characteristics of the created corpus. Section 4 details the CBCF architecture proposed to solve the AE task and SC task on the FreMPhone corpus. Section 5 illustrates the obtained experimental results. Section 6 provides the conclusion and some perspectives.

2 Related Works

Despite the abundance of ABSA-annotated corpora for the English language, there is a scarcity of corpora for the French language. After a survey study effectuated on the existing French corpora for the ABSA task, we concluded that only a little number of them are available. In this section, we present only the available French corpora for the ABSA task. These corpora are CANÉPHORE [5], the SemEval-2016 restaurant corpus [6], and the SemEval-2016 museum corpus [6]. The CANÉPHORE [5] corpus was created in 2015 and it contains 5372 tweets that discuss the “Miss France” event elaborated

in 2012. This corpus was annotated using the annotation tool Brat¹. For each tweet, the annotators identify the entities (subjects, aspects, or markers) and their sentiment polarities which can be positive or negative. As a result of this annotation, a set of 292 aspects was extracted. The SemEval-2016 restaurant corpus [6] was built within the fifth task of the SemEval international workshop. It was collected from the Yelp² website and consists of 457 restaurant reviews (2365 sentences) that reflect the attitudes of consumers concerning the restaurant domain. The annotation of this corpus was released in two steps. For the first time, the whole corpus was annotated by a French native speaker linguistic. After that, the annotation was verified by the organizers of this task. For each review, the annotator identifies the entities and their attributes and then assigns to them a sentiment polarity. The annotation in this work was effectuated based on 6 categories of entities (*restaurant, food, drinks, ambiance, service, location*), 5 categories of attributes (*general, prices, quality, style_options, miscellaneous*), and 3 sentiment classes (*positive, negative, neutral*). As a result of this annotation, a set of 3484 tuples was produced. Within the same task, [6] built the second French corpus for the museum domain. This corpus consists of 162 (655 sentences) French museum reviews crawled from the Yelp website for the ABSA task. This corpus was annotated based on 6 entities (*museum, collections, facilities, service, tour_guiding, location*) and 8 attributes (*general, prices, comfort, activities, architecture, interest, setup, miscellaneous*). Each entity-attribute pair was annotated into positive, negative, or neutral. At the end of the annotation, a set of 891 tuples was obtained.

These presented corpora are freely available online. Table 1 contains the different details of the three presented French corpora.

Table 1. Overview of ABSA available corpora.

Corpus	Size	Entity-Attribute	Number of Aspects	Source
CANÉPHORE	5372 tweets	-	292	Twitter
SemEval-2016 restaurant	2365 sentences	30	3841	Yelp
SemEval-2016 museums	655 sentences	48	891	Yelp

3 Corpus and Annotation

This section presents the created corpus “FreMPhone” (source, structure, statics, etc.) and details the different steps followed in its annotation. In the following paragraphs, we mean by aspect terms both aspects and entities.

¹ <http://brat.nlplab.org/>

² <https://www.yelp.com/>

3.1 Corpus Presentation

The FreMPhone corpus was collected from the Amazon³ website. For this research, Amazon was selected as the data source owing to its substantial volume of available comments. Amazon is one of the big American e-commerce companies that was founded in 1994. It sells several kinds of products (books, dvds, mobile phones, etc.) around the world and it gives its customers, by writing reviews on its website, the opportunity to express their opinions about the products that they purchased. This makes Amazon one of the largest websites that contain the biggest database of customer reviews. In this work, we have only been interested in the mobile phone domain. A corpus of 5651 mobile phone reviews was collected for this purpose from May 2022 until August 2022. To make our corpus more appropriate for the development of the ABSA method, we deleted the reviews that contain senseless sentences and that consisted only of one word. Also, we removed the reviews that are written in a language other than French. The final number of reviews obtained is 5217. These reviews encompass over six brands represented in Table 2.

Table 2. Number of reviews in the corpus by type of brand.

Brand	Samsung	iPhone	Nokia	Huawei	Oppo	Other brands
Number of reviews	2152	1445	589	456	328	274

To collect the FreMPhone corpus, we used the data extraction tool “Instant data scraper”⁴ which is based on AI (artificial intelligence) to predict which data is most relevant on the HTML page.).

3.2 Corpus Annotation

To well annotate our French mobile phone corpus, we studied some of the most referenced annotation guidelines for ABSA corpora such as, [2, 4], and [6]. After that, a subset of 1000 reviews was chosen randomly from the corpus and was annotated by two annotators separately. As a result, a set of 42 aspect terms were found in these reviews. Then, we prepared an annotation guideline, based on the studied guidelines. This guideline contains a brief definition of the ABSA task and rules for the aspect’s polarity classification.

To ensure a better process of annotation, we have divided the corpus into five sets, with each set containing approximately 1000 reviews. Every set was subjected to annotation by a minimum of two annotators, and each one of them worked individually on the same copy of the file. After the annotation of each set, the three annotators discuss the disagreement, divergences, and inconsistencies found in the annotation and fix the

³ <https://fr.wikipedia.org/wiki/amazon>

⁴ <https://chrome.google.com/webstore/detail/instant-datascraper/ofaokhiedipichpaobibbnahnkdoiiah>

misinterpretations. The final decision is taken by the first author. It is important to mention that in certain cases the first author might require to affect a series of improvements and adjustments in the guideline to enhance the process of annotation.

The process of annotation continued for more than 3 months. The annotators labeled approximately 50 reviews each day where each review contain a number of sentences ranging between 1 (at least) and 30 (at most). For each review text in the corpus, the annotators are required to identify the aspect terms and then assign to them a sentiment polarity. This sentiment polarity can be positive (+1), negative (-1), or neutral (0).

3.3 Annotation Scheme

As discussed above, the annotation process takes place in two steps. In the first step, the annotators are charged to identify the different aspect terms in a specific corpus. For this purpose, we have used the IOB [7] tagging scheme. Each identified aspect term is labeled as being at the beginning (B) or inside (I). The other words (non-aspects) in the corpus are labeled as outside (O). The following example represents the annotation manner of a review using the IOB format.

<i>The</i>	<i>iPhone</i>	<i>is</i>	<i>just</i>	<i>perfect</i>	<i>but</i>	<i>I'm</i>	<i>disappointed</i>	<i>by</i>	<i>the</i>	<i>storage</i>	<i>capacity</i>
<i>O</i>	<i>B</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>B</i>	<i>I</i>

As it is shown by the example above, when the aspect is composed of only one term, the tag “B” is given to the aspect term (“**iPhone**”). If the aspect term is composed of more than one term, the annotation scheme gives the first term of the compound aspect the tag “B” (“**storage**”) and gives the other words composing the compound aspect term the tag “I” (“**capacity**”).

After the identification of aspect terms, we will give each term of them a sentiment polarity. At this level, the label “POS” (if the customer expresses a positive sentiment toward the aspect), “NEG” (if the customer expresses a negative sentiment toward the aspect), or “NEUT” (if the customer mentions the aspect term in the review without expressing any sentiment) is assigned for each aspect term. The following example represents the annotation manner of a review using the POS, NEG, NEUT, and O labels.

<i>The</i>	<i>iPhone</i>	<i>is</i>	<i>just</i>	<i>perfect</i>	<i>but</i>	<i>I'm</i>	<i>disappointed</i>	<i>by</i>	<i>the</i>	<i>storage</i>	<i>capacity</i>
<i>O</i>	<i>B</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>B</i>	<i>I</i>
<i>O</i>	<i>POS</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>O</i>	<i>NEG</i>	<i>NEG</i>

3.4 Corpus Statistics

In this section, we provide the results of the statistical study performed on the collected corpus. The FreMPhone corpus contains a total number of words equal to 71461 distributed among 10758 sentences. Statistical information demonstrated in Fig. 1 proves the big variability in the number of words in reviews which ranges between 2 words and 322 words in each review. As it is clearly demonstrated in Fig. 1, the majority of

reviews (38%) are composed of less than 50 words. 29% of reviews contained a number of words ranging between 50 and 100 words. The rest of the reviews are distributed as follows, 19% are composed of a number of words that varies between 100 and 200 words, 9% contain a number of words ranging between 200 and 300 words and only 5% are composed of more than 300 words.

After the annotation of our corpus with IOB schema and sentiment labels (POS, NEG, NEUT), we obtained 19257 aspect terms. Figure 2 shows the distribution of these aspect terms according to the number of sentiment classes. 13259 aspect terms are classified as positive, 5084 aspect terms are classified as negative and 914 aspect terms are classified as neutral.




Fig. 1. Distribution of reviews according to the number of words.

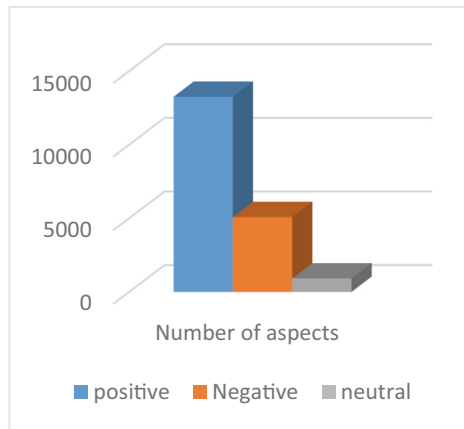


Fig. 2. Distribution of aspect terms according to the sentiment classes.

In addition, we present in Fig. 3 the 10 most used aspect terms in the corpus. As is shown, the most three discussed aspect terms in the corpus are “telephone/*telephone*” (appeared 2146 times in the corpus), “produit/*product*” (appeared 1025 times in the corpus), and “prix/*price*” (appeared 979 times in the corpus).

4 CBCF: A New Architecture for AE and SC

As discussed above, the ABSA task is divided into two subtasks: Aspect Extraction (AE), and Aspect Sentiment Classification (ASC). To solve these subtasks, we propose a new neural network-based architecture “CBCF” which is a combination of Bi-LSTM, CNN, and CRF models. As shown in Fig. 4, the proposed CBCF architecture consists of four main components: the embedding layer, the CNN model, the Bi-LSTM model, and the CRF model. In the embedding layer, the Word2Vec [8] model is used to create word embedding vectors. This model was trained on 20 000 mobile phone reviews collected from Amazon Website. After that, a CNN [9] model is employed to select features useful for the classification. This model is applied due to its capacity to detect proficient local features. However, this model still enables to capture of distant information related to

the target word. For that, we applied the Bi-LSTM model, which is capable of learning long-term dependencies. This model takes as input the features extracted by the CNN model and produces new features as output, incorporating contextual information of the target word. Subsequently, the generated features are incorporated into the CRF [11] model in order to detect the label of a given word. This model takes into account the dependencies between labels, where it exploits the neighboring labels to identify the current label of a target word.




Fig. 3. The top 10 most used aspect terms in the corpus.

Embedding Layer. The objective of this step is to transform the words contained in the corpus into numerical vectors which are comprehensible by the deep learning algorithms. In fact, many methods are used to achieve this purpose in ABSA-related studies such as the One-hot encoding method, the Term Frequency-Inverse Document Frequency (TF-IDF), etc. However, the main issue with such methods is related to their inability to capture contextual information which is important information in the ABSA task. For that, we employed in this study the word embeddings method. This method represents the words that appear in a similar context with word vectors relatively close.

Many models have been developed to create embeddings of words (such as Glove, ELMO, etc.). In this study, we employed the Word2Vec model. Similar to any model that is characterized by neural networks-based architecture, the Word2Vec model needs to be trained on a huge mass of information to give reliable word vectors that take into consideration contextual information. In fact, there are numerous publicly available word embedding models that have been trained on a big number of data collected from several sources (e.g. Twitter, Wikipedia, etc.). However, these models are inefficient in our task since they have been trained on multi-domain corpora and not on a mobile phone corpus. For that, we used 20 000 mobile phone Amazon reviews to train the Word2Vec model for both AE and SC tasks. Each word in the corpus is represented as a 350-dimensional vector.

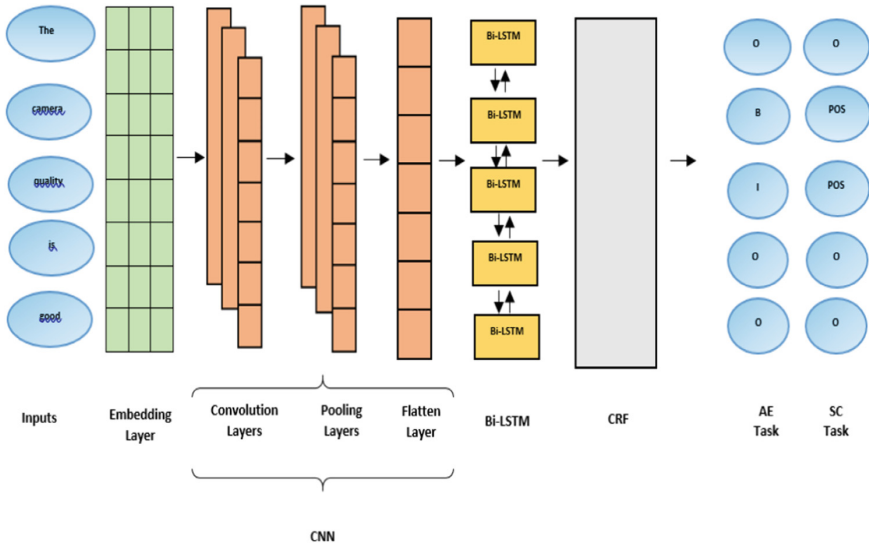


Fig. 4. The CBCF architecture.

Convolution Neural Network. The Convolution Neural Network (CNN) model is a multi-layers deep learning model presented for the first time by [9] to perform the task of forms recognition. In this study, we employed the CNN model in order to extract efficient local features. This model is composed of convolution layer (s), pooling layer(s), and flatten layer. In the first step, by receiving word embedding vectors from the embedding layer, it employs convolution operations to generate feature maps. This operation multiplies (or convolves) the matrix representation “M” by another convolution matrix (or filter) to produce a feature map. After that, the CNN model uses a pooling layer in order to decrease the number of parameters. Finally, the CNN model utilizes a flatten layer to convert the high-dimensional vectors, into one-dimensional vectors.

In this step, we used two convolution layers with 200 feature maps, the activation function ReLU, two pooling layers, each with a pool size equal to 200, and a flatten layer.

Long Short-Term Memory. The Long Short-Term Memory (LSTM) [10] is an artificial neural network model created to treat sequential data. It is an RNN model with a supplementary cell of memory. This model has been designed to solve the problems of vanishing gradients and long-term dependencies learning faced by traditional RNNs. For that, the main advantage of using LSTM lies in its ability to maintain long-term memory and control the flow of information within the network. The LSTM model follows the mechanism of forward propagation and treats the input sequence in one direction from left to right. It considers only the information coming from the previous units to detect the new outputs. In this study, we used a bidirectional LSTM (Bi-LSTM) model which is an enhanced version of the LSTM model. This model predicts the label of a word based on the information coming from the previous units (forward propagation) and the next units (backward propagation).

The Bi-LSTM model is composed of a series of Bi-LSTM neurons, each of which contains a memory cell, and three gates: an input gate, an output gate, and a forget gate. The input gate determines which information must be conveyed to the memory. The output gate chooses the value of the next node. The forget gate decides which information should be deleted.

At instant t , the Bi-LSTM model takes as input the one-dimensional vectors outputted by the flatten layer and produces new vectors of words that take into consideration the long-term dependencies between the words and the local information. These newly produced vectors are incorporated into the CRF model.

In this step, we used a Bi-LSTM model composed of 300 neurons.

Conditional Random Field. The Conditional Random Field (CRF) [11] is a statistical model that is often used for sequence labeling tasks, such as named entity recognition, sentiment analysis, etc. The CRF model assigns a label to each word of a sequence based on a set of parameters, known as transition parameters, to model the probability of a label transition from one word to the other. The CRF model has been used in multiple ABSA-related studies ([12, 13]) and it has achieved good classification results.

At instant t , the outputs of the Bi-LSTM model are passed to the CRF model in order to classify a given word. This model utilizes the interdependence or contextual connection between neighboring tags to predict the final output labels.

5 Experiments and Results

In this section, two experiments were conducted. The first experiment was realized to evaluate the performance of the proposed CBCF architecture for the Aspect Extraction (AE) task and Sentiment Classification (SC) task, as well as the reliability of the FreM-Phone corpus. However, in the second experiment, the performance of the Word2Vec model was assessed in comparison to the FastText model. This experiment aims to detect which model is more suitable to construct word embeddings for AE and SC tasks.

To achieve the experiments, we divided the constructed FreMPhone corpus into 80% for the training phase and 20% for the test phase. The performance of the models was evaluated using three metrics which are Precision (P), Recall (R), and F1-score (F1).

5.1 Experiment 1: Evaluation of the CBCF Architecture

As presented in Table 3, the best classification performances were achieved by the proposed CBCF architecture with an F-measure value equal to 95.97% where it surpassed the LSTM model (91.07%) by approximately 4%, the CNN model (67.14%) by approximately 28%, and the CRF model (72.95%) by approximately 23%.

Table 4 summarized the results obtained by the CBCF architecture and the other deep-learning models for the Sentiment Classification (SC) task on the FreMPhone. As shown in Table 4, we obtained promising results in the SC task where the best F-measure value was obtained after the use of the proposed CBCF (96.35%). This architecture overrode the LSTM model (with approximately 6%), the CNN model (with approximately 7%), and the CRF model (with approximately 40%) for the FreMPhone corpus.

The obtained F-measure (Table 3 and Table 4), using the CBCF architecture significantly surpassed the CNN, Bi-LSTM, and CRF models. So, the combination of CNN, Bi-LSTM, and CRF models gives better results than using each one of them separately and enhances obviously the AE and SC tasks. In addition, the obtained results prove the reliability of our created FreMPhone corpus.

Table 3. Results of the FreMPhone evaluation in the AE task using the CBCF architecture and other models.

		Aspect Extraction Task		
Corpus	Model	Precision	Recall	F-measure
FreMPhone	LSTM	87.76%	94.65%	91.07%
	CNN	76.11%	60.14%	67.19%
	CRF	81.90%	65.76%	72.95%
	CBCF	97.95%	94.08%	95.97%

Table 4. Results of the FreMPhone evaluation in the SC task using the CBCF architecture and other models.

		Sentiment Classification Task		
Corpus	Model	Precision	Recall	F-measure
FreMPhone	LSTM	91.62%	88.78%	90.17%
	CNN	83.81%	95.93%	89.46%
	CRF	61.14%	51.29%	55.78%
	CBCF	93.10%	99.83%	96.35%

5.2 Experiment 2: Comparison of the Word2Vec Model with the FastText Model for AE and SC Tasks

In this experiment, we compared the accuracy of the Word2Vec used model with the FastText [14] model with the aim of determining which model is more suitable for building word embeddings for the AE task and SC task. FastText [14] is a model that utilizes the skip-gram model to generate word embeddings. Each word is transformed into a group of n-grams, and each n-gram is assigned a vector representation. The embedding of a word is obtained by summing up the vectors of all its character n-grams. This method allows FastText to handle out-of-vocabulary words and morphological variations by breaking them down into sub-words and utilizing the information from the sub-words to generate a vector representation for the word. Figure 5 highlights the F-measure results achieved by the CBCF model when employing the Word2Vec and FastText models in the AE task and SC task on the FreMPone corpus. As it is clear in Fig. 5, the FastText model

achieved much better results than the Word2Vec model for the AE task. This result aligns with the nature of the task as it aids in better learning the context information of the aspects within the review. Also, unlike the Word2Vec model, FastText has the capability to create representations for words that did not appear in the training data which can improve mainly the process of aspects detection. However, we remarked that the Word2Vec model achieved better results than the FastText model in the SA task. These findings can be explained for many reasons. Firstly, the Word2Vec model captures the semantic and syntactic relationships between words, which can be beneficial for the detection of opinion words describing the aspect terms, and thus the SC task. Also, Word2Vec’s embeddings are based on the co-occurrence statistics of words in the training data, which can better detect the sentiment of a word based on the words that appear in the same context.

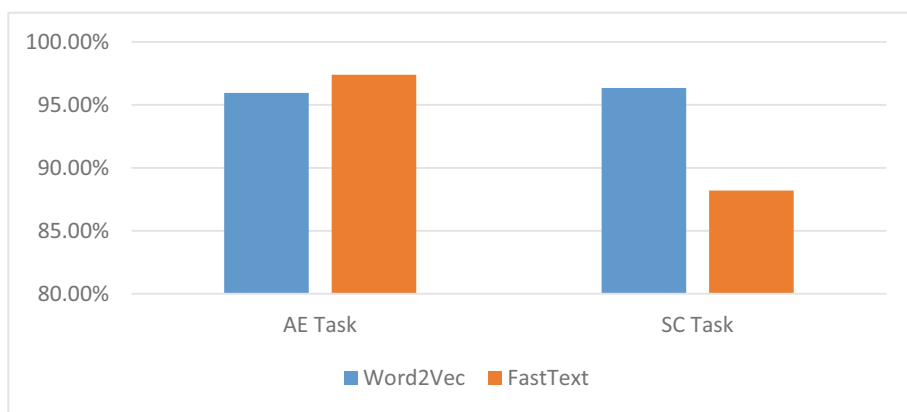


Fig. 5. Illustration of FreMPhone corpus F-measure results using Word2Vec and FastText representation models.

6 Conclusion

In this paper, we presented the French mobile phone corpus FreMPhone constructed for the ABSA task. This corpus was collected from the Amazon website and contains 5217 mobile phone reviews. In this study, we described the process followed in the collection and the annotation of the corpus. Also, we provided a statistical study of the corpus including the number of aspects found, and their distribution according to the positive, negative, and neutral classes. In addition, we proposed a new architecture “CBCF” to evaluate the FreMPhone corpus constructed. This architecture combines the strength of the deep learning models, namely CNN and Bi-LSTM, with the machine learning model CRF. This architecture surpassed the deep learning models and achieved very encouraging results reach to 95.97% for the AE task and 96.35% for the SC task. The obtained results validate the effectiveness and high performance of our proposed architecture in accurately extracting aspect terms and their associated sentiment. Furthermore, they

demonstrate that combining deep learning models with machine learning models yields superior results compared to using each model separately. Deep learning models excel at capturing complex patterns and relationships in textual data, while machine learning models provide a structured framework for labeling and sequence tagging tasks. Moreover, we compared the performance of the word embedding models namely, Word2Vec and FastText. Findings show that the FastText constructs more pertinent embedding vectors for the AE task, while the Word2Vec model ensures a better SA. These outcomes prove that the choice of a suitable word embedding model can have a significant impact on the performance of aspect-based sentiment analysis tasks.

Overall, this research contributes valuable insights into the field of ABSA, providing a comprehensive corpus, an effective architecture, and a comparative analysis of word embedding models. These findings pave the way for future advancements in aspect-based sentiment analysis and hold significant potential for various applications in natural language processing and sentiment mining.

In future work, we will enhance the annotation process and we will annotate the corpus according to the implicit aspect terms and the aspect's categories. Also, we will extend the FreMPhone corpus, with reviews collected from other domains such as restaurants, hotels, etc. Also, we will improve the proposed CBCF architecture by using BERT embedding.

References

1. Minqing, H., Bing, L.: Mining and summarizing customer reviews. In: 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining ACM, pp. 168–177 (2004)
2. Wang, H., Lu, Y., Zhai, C.: Latent aspect rating analysis on review text data: A rating regression approach. In: Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2010), pp. 783–792. Washington, US (2010)
3. Samaneh, M., Martin, E.: Opinion Digger: An unsupervised opinion miner from unstructured product reviews. In: 19th ACM International Conference on Information and Knowledge Management (CIKM 2010), pp. 1825–1828 (2010)
4. Pontiki, M., et al.: Semeval-2014 task 4: Aspect based sentiment analysis. In: ProWorkshop on Semantic Evaluation. Association for Computational Linguistics, pp. 20–31
5. Lark, J., Morin, E., Saldarriaga, S.: CANÉPHORE: un corpus français pour la fouille d'opinion ciblée. In: 22e conférence sur le Traitement Automatique des Langues Naturelles (2015)
6. Pontiki, M., et al.: Semeval-2016 task 5: Aspect based sentiment analysis. In: ProWorkshop on Semantic Evaluation Association for Computational Linguistics, pp.19–30 (2016)
7. Ramshaw, L., Marcus, M.: Text chunking using transformation-based learning. Natural language processing using very large corpora, pp 157–176 (1999)
8. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint [arXiv:1301.3781](https://arxiv.org/abs/1301.3781) (2013)
9. Fukushima, K., Sei, M.: Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition. In: Competition and cooperation in neural nets, pp. 267–285. Springer, Berlin, Heidelberg (1982)
10. Graves, A.: Long short-term memory. Supervised sequence labeling with recurrent neural networks, pp. 37–45 (2012)
11. Lafferty, J., McCallum, A., Pereira, F.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data (2001)

12. Heinrich, T., Marchi, F.: Teamufpr at absapt 2022: Aspect extraction with crf and bert. In: Proceedings of the Iberian Languages Evaluation Fórum (IberLEF 2022), co-located with the 38th Conference of the Spanish Society for Natural Language Processing, Online. CEUR (2022)
13. Lei, S., Xu, H., Liu, B.: Lifelong learning crf for supervised aspect extraction. arXiv preprint [arXiv:1705.00251](https://arxiv.org/abs/1705.00251) (2017)
14. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. arXiv preprint [arXiv:1607.01759](https://arxiv.org/abs/1607.01759) (2016)