# Opinion mining on experience feedback: A case study on smartphones reviews

Laurent Brisson
**Institut Mines-Telecom, Telecom Bretagne**
UMR CNRS 6285 Lab-STICC
Université européenne de Bretagne, Brest, France
Email: laurent.brisson@telecom-bretagne.eu

Jean-Claude Torrel
**EGIIS**
Technopôle Brest Iroise, Brest, France

Email: jtorrel@egiis.fr

***Abstract*—Through the development of electronic commerce, social media and collaborative media, the social commerce appeared. Social commerce, a subset of electronic commerce, is based on social interactions in order to buy and sell goods and services. Nowadays, before buying, people give more importance to the experience feedback they found on internet. However, it is difficult to get an overview of this experience feedback since it is scattered in many online resources, and buyers never have time to read many pages of comments. In this paper, we present an approach which grabs and analyzes experience feedback in order to publish a summary of opinions about a product. We develop this approach with a case study on smartphones and publish a dataset of thousands of comments on a wide range of smartphones. To summarize experience feedback, we use a linguistic appraisal model, based on appreciation, affect and judgement, and we set up an approach using methods and tools from the fields of natural language processing, opinion mining and sentiment analysis.**

## I. Introduction

Electronic commerce is a sector that is still growing rapidly today and have to face new challenges: the consumers are increasingly demanding as they look for the best value for their money and search lots of information on products they want. What has changed fundamentally is the nature of the consumer. Generations Y and Z behave differently with social media from previous generations. Now, social interactions have a strong and immediate impact on purchase behaviour. However, although it is now easy to find a huge amount of experience feedback on many goods and services, the only overall vision that one usually gets is a ranking. Thus, to form an opinion, it is often necessary to go through several pages of comments.

Experience feedback analysis consists in analyzing consumers' comments which may contain emotions and explicit or implicit facts. To deal with such data, while the amount is far beyond the reading capability of a human being, automatic methods to extract and summarize opinions, known as opinion-mining or sentiment analysis, emerged. Opinion mining is the field that deals with the analysis of subjective statements from texts, the identification of opinions, the estimation of their orientation and the extraction of arguments that relate to these opinions. Opinion mining can be performed at various levels of analysis: document level, sentence level and aspect level. Document level opinion mining analysis [1] is beyond the scope of this paper since it aims to give an overall rating of the polarity of a document. Sentence level opinion mining analysis [2], [3] is well-fitted to polarity classification of short texts like tweets or short reviews. However, that kind of analysis does not consider the various aspects of an opinion, that is why lots of works mainly focus on aspect level polarity. Aspect level analysis is often a two-step task: the first one consists in aspect extraction while the second one focuses on polarity detection which requires opinion word extraction.

In this paper, we address a specific challenge related to aspect level opinion mining, building a platform which summarizes and qualifies experience feedback from reviews. As a constraint, these reviews may be from various domains and the whole analysis process must be deployed in a short period. This constraint excludes the use of supervised learning algorithms that often need lots of works to create a relevant learning dataset. The novelty of this approach and its application, relies on the use of a linguistic framework which focuses on appraisal in English [4] and a process that implements this framework in two main steps: the first one assists an expert to quickly create a knowledge base in order to extract relevant opinions and the second one runs syntactic analysis and search strategies. As part of our contribution, we make our dataset and a validation set available to the community in order to allow cross-validation with other approaches [5].

This paper is organized as follows. Section 2 presents related work in the field of opinion mining and sentiment analysis. Section 3 presents the linguistic framework we use. Section 4 presents the overall approach, methods and tools. Section 5 presents appreciation and affect extraction methods. Section 6, presents the platform, the results and discusses future work.

## II. Opinion Mining and Sentiment Analysis

Aspect extraction approaches can be divided into unsupervised approaches, mainly domain and language independent, and supervised approaches that often require manually labeled data to train models. Unsupervised approaches include statistical and rule-based approaches. Hu and Liu [6] apply apriori algorithm to find frequent aspects. Firstly, they apply a Part-of-Speech (POS) tagging algorithm in order to split the text into qualified words (nouns, verbs, adjectives, etc.). Then, after association rules extraction, the most relevant aspects are extracted by pruning the most frequent itemsets which are meaningless or redundant. Using some linguistic rules that correlate adjectives with opinion expression [7], they extract opinion words and evaluate their orientation. Popescu and

Etzioni [8] achieve higher precision using relaxation labeling, an unsupervised classification algorithm, which takes into account the context of a word and some local constraints manually specified (conjunctions, disjunctions, syntactic dependency rule templates, morphological relationships, synonyms and antonyms). Supervised approaches, also called "model-based approaches", aim to improve statistical based approaches with some features - in the machine learning sense - that cannot be detected without training. Kessler and Nicolov [9] use Support Vector Machine and manually labeled data to train a model that ranks extracted aspect/opinion pairs. Moghaddam [10] uses Generalized Sequential Pattern (GSP) mining to find patterns to extract relevant opinions. Zheng and al. [11] uses a variant of Latent Dirichlet Allocation for jointly extract aspect and sentiment words.

Several sources of knowledge can be used to detect polarity: corpus, dictionaries and lexicons. Corpus-based approaches try to find co-occurrences of words or phrases to detect opinions. Dictionary-based approaches use synonyms, antonyms and generally WordNet dictionary [12]. Lexicon-based approaches use lexicons like SentiWordNet [13] or create their own lexicon. In the field of lexicon induction, some works focus on unsupervised induction based on subjective features [14], while others use labeled data [15]. Blair-Goldensohn et al. [16] adopt a semi-supervised approach and use provided labels and domain specific characteristics of service reviews to perform polarity classification. They implement a common approach using a set of seed words (words whose polarity is already known) to infer opinions [17].

Some approaches do not consider aspect-level analysis as a two-steps task and perform a joint aspect/opinion extraction. Miao et al. [15] and Li et al. [18] merge aspect and opinion extraction in a unified process using Conditional Random Fields (CRF). CRFs models are conditional probabilistic sequence models like Maximum Entropy Markov Models (MEMM).

## III. "Ways of Feelings": Linguistic framework and Semantic resources

### A. A linguistic framework for appraisal

People give their feedback about goods and services in many different ways. One can write a well-structured review and be factual while others are subjective and express their feelings. Natural Language Processing (NLP), opinion mining and sentiment analysis fields provide lots of methods and algorithms to analyze, classify and evaluate the polarity of various kind of documents.

To structure and to conduct our computational processes a framework modeling the language seemed us necessary. Since we have to deal with experience feedback, we choose to use the framework proposed by two linguists, James R. Martin and Peter R. R. White [4] which focuses on appraisal in english. Appraisal is a discourse semantic resource which includes attitude (dealing with our feelings), engagement (dealing with the way to express feelings) and graduation (dealing with the way to intensify or to soften feelings). To achieve our goal we decided, in a first approach, to focus on attitude. In their framework, Martin and White, define attitude with three concepts: appreciation, affect and judgement.

*Definition 1 (Appreciation):* Appreciation involves evaluations of semiotic and natural phenomena, according to the ways in which they are valued or not in a given field. [4]

Our appreciation extraction method is an aspect-based opinion mining task. We aim to extract three different expressions of appreciation: explicit qualifier about an explicit aspect *"The battery life is good"*, *"The only problem is the camera"* (here the qualifier could be a noun (NN) or an adjective (JJ)), the existence (or not) of an aspect *"It doesn't have auto-focus"* and the behavior of an aspect *"Touch screen lags a bit at times after an app is closed"*.

*Definition 2 (Affect):* Affect is concerned with registering positive and negative feelings: do we feel happy or sad, confident or anxious, interested or bored ? [4]

Our affect extraction method aims to extract two different expressions of affect: affect as a "mental processes" *"I am enjoying with this phone"*, *"I hate this phone"*, *"I like its display"* and affect as a "quality" *"I'm very happy"*. Consequently we perform an aspect-based opinion mining task. and a sentiment analysis of the sentence.

*Definition 3 (Judgement):* Judgment deals with attitudes towards behavior, which we admire or criticize, praise or condemn. [4]

According to Martin and White, judgement is clearly different from appreciation in the way judgement express feelings about behavior while appreciation express feelings about the value of things. In our context, handling judgement is out of scope but it's a perspective for future work. That is why, in the following, we consider opinion mining from experience feedback as two tasks: appreciation extraction and affect extraction.

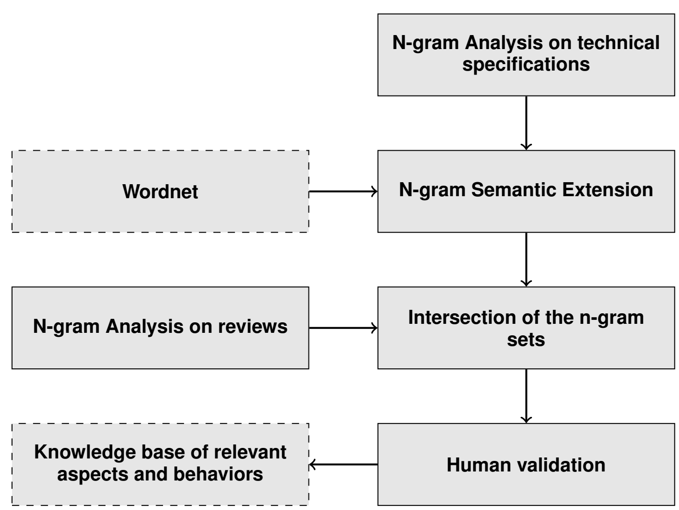


Fig. 1. Knowledge base building process

### B. A knowledge base for semantic relevance

The knowledge base is an important part of our approach since it allows us to define a relevant semantic perimeter for the analysis. However, as one of our constraints is to be able to deploy the whole process of our approach in a short period we have to minimize interactions with an expert to validate the content of this knowledge base. The knowledge base is composed of two resources: an aspect-lexicon which is a list of words and phrases defining the relevant aspects for an opinion mining task and a polarity lexicon (we use Sentic lexicon [19]) which gives us polarity information about words. In the next section we explain how this information help us to extract relevant opinions.

Before analyzing opinions, we have to collect textual resources of the same domain in order to help the expert to build the aspect-lexicon. For example, in our experiment we gathered technical specifications about various smartphones. Then, we performed 1-gram and 2-gram analysis on both technical specifications dataset and user comments dataset. We restricted 1-gram to words which are 'nouns' in their sentence. Then, the 1-gram and 2-gram sets from technical specifications were extended with Wordnet [12] using synonyms, direct hyperonyms and direct hyponyms and we intersected n-gram list from reviews with the extented n-gram list from specifications in order to select relevant aspects. In the last step, an expert validated the list. Here are some recommendations to the expert to quickly select the most relevant terms. First, check the most frequent words which are present in both dataset (for example, in our experiment we considered only words occurring more than 100 times). Second, filter trademarks and proper nouns: these words could be used later for automated context definition. Last, check words added with synonymy, hypernonymy or hyponymy relationships since these word are most likely the less relevant ones.

## IV. APPRAISAL PATTERNS EXTRACTION

### A. Overall approach

Figure 2 gives an overwiew of our approach. First, we perform lexical analysis and lexical transformations on our dataset: character cleaning, full stop detection, tokenization and part-of-speech tagging. Second, we perform syntactic analysis. For these two first tasks we use OpenNLP [20]. Finally, we use a custom knowledge base, in order to extract relevant speech elements related to appreciation or affect in the context of smartphone reviews. This approach is based on two key concepts: syntactic analysis and strategies to find appraisal patterns in parse trees.

Syntactic analysis is performed using constituency-based parse tree which are based on Chomsky language theory [21] and use constituent grammars. A constituency-based parse tree (Figure 4) focuses on the linear structure of the phrase where intermediate nodes are non-terminal grammatical categories and leaves are terminal grammatical categories [22]. Our strategies, to find appraisal patterns, follow three steps: a breadth-first search (BFS) in the parse tree, pattern recognition using automata and rule triggering to extract relevant opinions. During the breadth-first search in the parse tree, each node is submitted to an automaton to match some patterns. We made the choice to use automata because they are a sound and readable method for pattern extraction, but they are not in the core of our approach and could be replaced by another suitable method. When a node of the tree triggers a transition, its sub-trees are handled by the current state of the automaton. When an automaton reaches a final state, a rule is triggered according to some conditions. These conditions are mainly about word polarity or word presence in the knowledge base. They are explained in the next section. For this purpose, all nodes triggering a transition are saved into a variable. Let $N$ be the set of all nodes which trigger a transition toward the state $S$ of an automaton. We define $V_s$ the saved variable for $S$ state, as the union of all leaves of the nodes of $N$. Properties of saved nodes are: $polarity$ (given by Sentic lexicon [19]), chunk or part-of-speech $tag$, and membership to our custom aspect knowledge base $kb$. In our approach, polarity inverter words such as 'not' or 'no' are called modifiers.

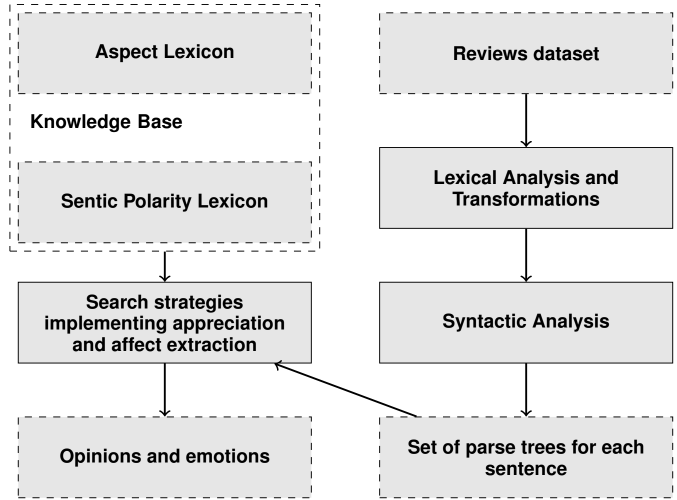


Fig. 2. Overview of the approach

### B. Appreciation Extraction

Appreciation extraction consists in several strategies to extract aspect qualification, aspect presence and behaviour qualification.

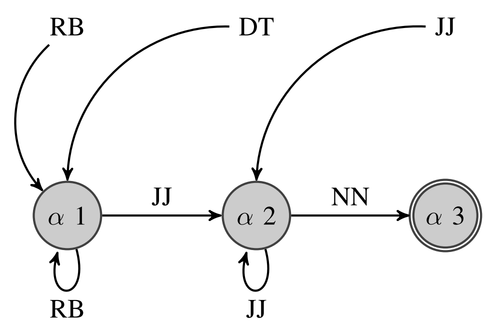


Fig. 3. $\alpha$ automaton is dedicated to extract explicit adjectives on an aspect

*1) Aspect qualification:* We consider two ways to qualify an aspect: using adjectives and using nouns. The $\alpha$ automaton (Figure 3) is designed to extract explicit adjectives about an aspect. For example, let's consider the sentence "What a good phone!" (Figure 4). We define a rule which is triggered if $\alpha 3$ is in the aspect-lexicon and if $\alpha 2$ is a polarized adjective. That rule extracts modifiers from $\alpha 1$ (none in this example since there isn't a negative word in the sentence), qualifiers from $\alpha 2$ ('good') and the aspect from $\alpha 3$ ('phone').

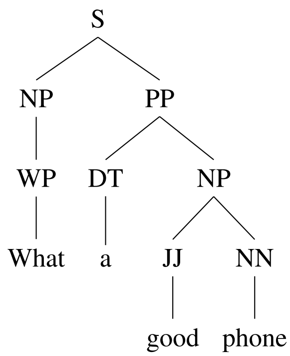


Fig. 4. Parse Tree of *"What a good phone!"*

In order to handle more complex syntactic structures we define the $\gamma$ automaton (Figure 6). It is a hierarchical automaton that handle several rules related to appreciation but also to affect. Sub-automata $\gamma B$ and $\gamma C$ are recursives: all the states of $\gamma B$ trigger $\gamma B$ with a 'VP' node and all the states of $\gamma C$ trigger $\gamma C$ with a 'NP' or 'PP' node. This property allows us to handle the conjunction 'and' since it creates nested 'NP' phrases. Negation words are handled with $\gamma 3$ and $\gamma 4$ states within each rules presented below.

Another way to qualify an aspect is to use nouns instead of adjectives. To understand how we extract nouns that qualify an aspect let's consider the sentence "The camera is the only problem" (Figure 5). A rule is triggered if $\gamma 1$ is in the aspect-lexicon of the knowledge base and if $\gamma 6$ is a polarized noun that is not in the knowledge base. Thus, we extract the aspect 'camera' from $\gamma 1$ (since the word is in the aspect-lexicon) and the qualifier 'problem' from $\gamma 6$ since this word has a negative polarity and is not in aspect-lexicon.

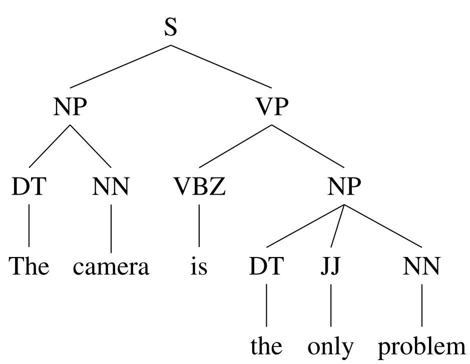


Fig. 5. Parse Tree of *"The camera is the only problem"*

*2) Aspect presence:* Let's consider the sentence "The phone has a good camera" (Figure 7). In order to detect the presence of an aspect we use $\gamma$ automaton and we define a rule that is triggered if $\gamma 1$ and $\gamma 6$ are in the aspect-lexicon and if $\gamma 3$ contains a verb expressing membership. Thus, we extract the aspect from $\gamma 1$ ('phone') and the qualifier from $\gamma 6$ ('camera').

*3) Behaviour qualification:* In order to detect an opinion about the behaviour of an aspect we use $\gamma$ automaton and we define a rule that is triggered if $\gamma 1$ is in the aspect-lexicon and $\gamma 3$ contains a polarized verb. Thus, we extract the aspect

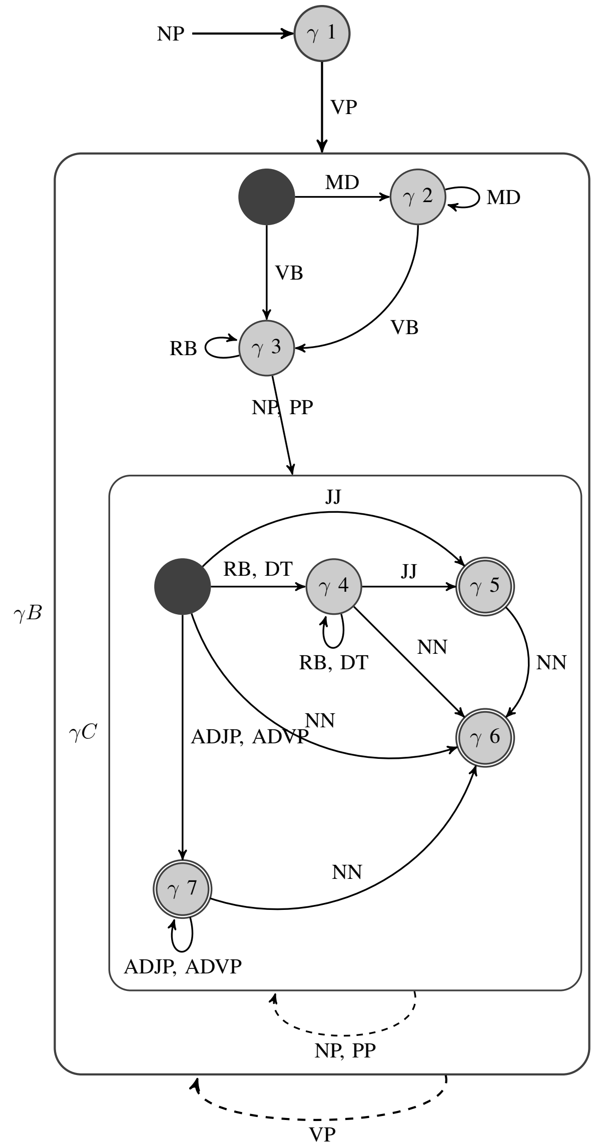


Fig. 6. $\gamma$ automaton is hierarchical: it triggers $\gamma B$ and $\gamma C$ which are recursives. A dashed loop means each state of the automaton can trigger the automaton itself (see section IV-B1).

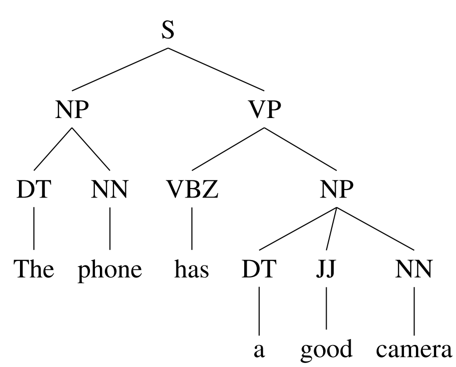


Fig. 7. Parse Tree of *"The phone has a good camera"*

from $\gamma 1$ ('touch screen') and the qualifier from $\gamma 3$ ('lags'). No modifier is extracted from $\gamma 3$ and $\gamma 4$ in this example.

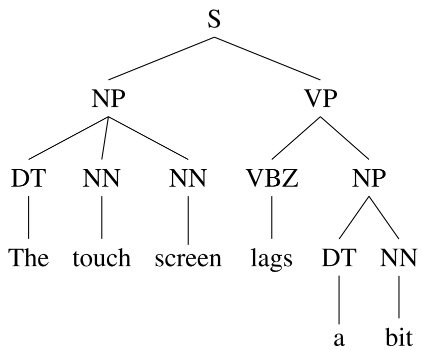


Fig. 8. Parse Tree of *"The touch screen lags a bit"*

### C. Affect Extraction

In order to extract affect expressed like a mental process we use $\gamma$ automaton and we define a rule that is triggered if $\gamma1$ is the personal pronoun 'I', $\gamma6$ is in the aspect-lexicon and $\gamma3$ is a polarized verb. Let's consider the sentence "I hate this phone" (Figure 9). Since the personal pronoun 'I' and the polarized verb 'hate' are in the sentence, we extract the aspect from $\gamma6$ ('phone') and qualifiers from $\gamma3$ ('hate').

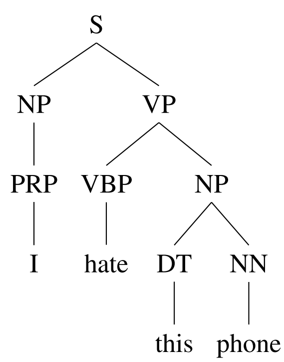


Fig. 9. Parse Tree of *"I hate this phone"*

## V. An Online Platform to Summarize Users Feedback: Results and Discussion

### A. Datasets

We built a dataset about 368 smartphones, containing 40,160 comments and 81,431 sentences for a total size of 8.5 Mo. We also gathered their technical specifications from 8 manufacturer websites. In order to validate our approach we selected a subset of our dataset and we gave the task to 3 people to read and annotate 527 comments (containing 708 opinions) on 6 phones. These people had to find qualifiers about smartphones aspects and to qualify their polarity. We only keep, in the final validation dataset, annotations with a consensus between the three reviewers. These datasets are available online [5].

### B. Results

We assess the performance of our approach considering two tasks: finding relevant appreciation in user feedback and qualifying the polarity of an appreciation. For this evaluation, we use the three following measures: precision, recall and f-measure.

| Dataset | Appreciation Detection | | |
|---|---|---|---|
| | Precision | Recall | F-Measure |
| Galaxy S5 | 0.87 | 0.62 | 0.72 |
| HTC Desire 310 | 0.90 | 0.50 | 0.64 |
| HTC One (M8) | 0.82 | 0.43 | 0.56 |
| iPhone 6 | 0.83 | 0.50 | 0.62 |
| Lumia 1320 | 0.90 | 0.68 | 0.78 |
| XPERIA C | 0.87 | 0.66 | 0.75 |
| Average | 0.87 | 0.57 | 0.68 |

Fig. 10. Appreciation detection performances

| Dataset | Polarity Evaluation | | |
|---|---|---|---|
| | Precision | Recall | F-Measure |
| Galaxy S5 | 0.82 | 0.59 | 0.69 |
| HTC Desire 310 | 0.82 | 0.48 | 0.60 |
| HTC One (M8) | 0.80 | 0.43 | 0.56 |
| iPhone 6 | 0.77 | 0.47 | 0.59 |
| Lumia 1320 | 0.86 | 0.68 | 0.76 |
| XPERIA C | 0.76 | 0.62 | 0.68 |
| Average | 0.81 | 0.55 | 0.65 |

Fig. 11. Polarity evaluation performances

Figures 10 and 11 show performances of our approach over the validation dataset. Average precision to find relevant appreciation is 0.87% and average precision to qualify polarity is 0.81%. Polarity errors are mainly due to semantic confusions since we do not consider the context of the sentence to define the polarity of a word. Average recall to find relevant appreciation is 0.57% and average recall to qualify polarity is 0.55%. These results satisfy our objective to extract relevant comments to understand the pros and cons of each smartphone, however, the recall of our approach have to be improved. There are several issues to explain this low recall that we plan to address in our future work. First, we did not yet implement affect extraction expressed like a mental state. Furthermore, it is difficult to handle complex sentences with constituency-based syntactic analysis. Finally, we do not link several phrases in the same sentence. For example, in the sentence "I have just bought this phone, it lags!", we currently do not consider 'it' as the 'phone' and our behavior qualification rule is not triggered.

### C. Discussion and future work

These results show we address our initial challenge to extract relevant opinions without supervised learning on a specific dataset. In order to improve the number of opinions detected, we want to consider two possible approaches: the syntactic one and the pragmatic one. We can improve syntactic analysis with dependency-based syntactic analysis [23], [24] which provides grammatical relationships between words to handle more complex phrases. Since, interpreting grammatical dependencies is a difficult task, we plan to bootstrap this work using our current approach. The pragmatic approach is to consider contextual information during semantic analysis. This context can be handled within a document or outside a document. To handle context within a document, we plan to build a context-based polarity lexicon. For this purpose, we have to transform our knowledge base into an ontology to be able to manage several semantic domains. To handle context outside a document, we plan to collect data on the use of the platform we deployed in order to gather user feedback about our opinion summaries.

## VI. Conclusion

Social commerce is the new challenge of electronic commerce. Nowadays, before buying, people give more importance to the experience feedback they found on internet than to products' technical specifications. However, although it is now easy to find a huge amount of experience feedback on many goods and services, it is often necessary to go through several pages of comments to have a coherent overview. In this paper, we present a platform that summarizes and qualifies user-experience feedback. We use a sound language framework modeling appraisal in English and two core concepts to model attitude: appreciation and affect. In a first step, we built a knowledge base in a interactive way with a expert in order to define a relevant semantic perimeter for the analysis. Then, we applied a constituency-based syntactic analysis and strategies to find appraisal pattern in parse trees. These parse trees were analyzed using automata and rules in order to extract appreciation and affect expressions. We built two datasets which are available online [5] to allow cross-validation with other approaches, the first one is a full dataset with 368 smartphones, while the second is a validation dataset, annotated manually to evaluate and compare approaches. Our results show good precision performances which guarantees relevant opinion summaries on the platform.

## References


[1] J. Blitzer, M. Dredze, and F. Pereira, "Biographies, bollywood, boomboxes and blenders: Domain adaptation for sentiment classification," in *In ACL*, 2007, pp. 187–205.

[2] N. Farra, E. Challita, R. Assi, and H. Hajj, "Sentence-level and document-level sentiment mining for arabic texts," in *Data Mining Workshops (ICDMW), 2010 IEEE International Conference on*, Dec 2010, pp. 1114–1119.

[3] R. Sharma, S. Nigam, and R. Jain, "Polarity detection at sentence level." *International Journal of Computer Applications*, vol. 86, p. 29, 2014.

[4] P. R. R. White and J. R. Martin, *The Language of Evaluation: Appraisal in English*. London/New York: Palgrave Macmillan, 2005.

[5] L. Brisson and J.-C. Torrel, "Smartphones dataset," 2015. [Online]. Available: http://perso.telecom-bretagne.eu/laurentbrisson/datasets/

[6] M. Hu and B. Liu, "Mining opinion features in customer reviews," in *Proceedings of the 19th National Conference on Artifical Intelligence*, ser. AAAI'04. AAAI Press, 2004, pp. 755–760. [Online]. Available: http://dl.acm.org/citation.cfm?id=1597148.1597269

[7] V. Hatzivassiloglou and J. M. Wiebe, "Effects of adjective orientation and gradability on sentence subjectivity," in *Proceedings of the 18th Conference on Computational Linguistics - Volume 1*, ser. COLING '00. Stroudsburg, PA, USA: Association for Computational Linguistics, 2000, pp. 299–305. [Online]. Available: http://dx.doi.org/10.3115/990820.990864

[8] A.-M. Popescu and O. Etzioni, "Extracting product features and opinions from reviews," in *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, ser. HLT '05. Stroudsburg, PA, USA: Association for Computational Linguistics, 2005, pp. 339–346. [Online]. Available: http://dx.doi.org/10.3115/1220575.1220618

[9] J. S. Kessler and N. Nicolov, "Targeting sentiment expressions through supervised ranking of linguistic configurations," in *3rd Int'l AAAI Conference on Weblogs and Social Media (ICWSM 2009)*, E. Adar, M. Hurst, T. Finin, N. S. Glance, N. Nicolov, and B. L. Tseng, Eds., 2009. [Online]. Available: http://www.cs.indiana.edu/~jaskessl/icwsm09.pdf

[10] S. Moghaddam and M. Ester, "Opinion digger: an unsupervised opinion miner from unstructured product reviews," in *Proceedings of the 19th ACM international conference on Information and knowledge management*, ser. CIKM '10. New York, NY, USA: ACM, 2010, pp. 1825–1828. [Online]. Available: http://doi.acm.org/10.1145/1871437.1871739

[11] X. Zheng, Z. Lin, X. Wang, K.-J. Lin, and M. Song, "Incorporating appraisal expression patterns into topic modeling for aspect and sentiment word identification," *Knowledge-Based Systems*, vol. 61, no. 0, pp. 29 – 47, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0950705114000434

[12] G. A. Miller, "Wordnet: A lexical database for english," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, Nov. 1995. [Online]. Available: http://doi.acm.org/10.1145/219717.219748

[13] A. Esuli and F. Sebastiani, "Sentiwordnet: A publicly available lexical resource for opinion mining," in *In In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC06*, 2006, pp. 417–422.

[14] P. D. Turney, "Thumbs up or thumbs down?: Semantic orientation applied to unsupervised classification of reviews," in *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ser. ACL '02. Stroudsburg, PA, USA: Association for Computational Linguistics, 2002, pp. 417–424. [Online]. Available: http://dx.doi.org/10.3115/1073083.1073153

[15] Q. Miao, Q. Li, and D. Zeng, "Mining fine grained opinions by using probabilistic models and domain knowledge," in *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology - Volume 01*, ser. WI-IAT '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 358–365. [Online]. Available: http://dx.doi.org/10.1109/WI-IAT.2010.193

[16] S. Blair-goldensohn, T. Neylon, K. Hannan, G. A. Reis, R. Mcdonald, and J. Reynar, "Building a sentiment summarizer for local service reviews," in *Proceedings of the WWW2008 Workshop: NLP in the Information Explosion Era*, 2008. [Online]. Available: http://www.cl.cs.titech.ac.jp/~fujii/NLPIX2008/paper3.pdf

[17] A. Andreevskaia and S. Bergler, "Mining wordnet for a fuzzy sentiment: Sentiment tag extraction from wordnet glosses." D. McCarthy and S. Wintner, Eds. The Association for Computer Linguistics, 2006. [Online]. Available: http://dblp.uni-trier.de/db/conf/eacl/eacl2006.html#AndreevskaiaB06

[18] F. Li, C. Han, M. Huang, X. Zhu, Y.-J. Xia, S. Zhang, and H. Yu, "Structure-aware review mining and summarization," in *Proceedings of the 23rd International Conference on Computational Linguistics*, ser. COLING '10. Stroudsburg, PA, USA: Association for Computational Linguistics, 2010, pp. 653–661. [Online]. Available: http://dl.acm.org/citation.cfm?id=1873781.1873855

[19] E. Cambria, R. Speer, C. Havasi, and A. Hussain, "Senticnet: A publicly available semantic resource for opinion mining." in *AAAI fall symposium: commonsense knowledge*, vol. 10, 2010, p. 02.

[20] J. Kottmann, B. Margulies, G. Ingersoll *et al.*, "Apache opennlp. the apache software foundation," 2013.

[21] N. Chomsky, *Knowledge of Language: Its Nature, Origin, and Use*, ser. Convergence. Praeger, 1986.

[22] B. Santorini, "Part-of-speech tagging guidelines for the penn treebank project (3rd revision)," 1990.

[23] M. Marneffe, B. Maccartney, and C. Manning, "Generating typed dependency parses from phrase structure parses," in *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC-2006)*. Genoa, Italy: European Language Resources Association (ELRA), May 2006.

[24] M.-C. de Marneffe and C. D. Manning, "The stanford typed dependencies representation," in *Coling 2008: Proceedings of the Workshop on Cross-Framework and Cross-Domain Parser Evaluation*, ser. CrossParser '08. Stroudsburg, PA, USA: Association for Computational Linguistics, 2008, pp. 1–8. [Online]. Available: http://dl.acm.org/citation.cfm?id=1608858.1608859