Universität Zürich[UZH]

Institut für Computerlinguistik

A verb-centered Sentiment Analysis for French

Eine verb-zentrierte Sentimentanalyse für Französisch

Masterarbeit der Philosophischen Fakultät der Universität Zürich im Studienfach Multilinguale Textanalyse

Referent: Prof. Dr. M. Volk

Betreuer: Dr. M. Klenner

Verfasserin: Susanna Tron

Matrikelnummer 07-428-741

Dörflistrasse 107

8050 Zürich

November 11, 2013

Abstract

Sentiment Analysis is a challenging domain of NLP in many aspects. Words and phrases need to be recognised and marked not only with regard to their grammatical and syntactic role, but also with regard to their polarity, which can be positive, negative or neutral. Furthermore, they cannot be considered in isolation, because they can be modified by other elements of the sentence, which include other polar words, prepositions, adverbs and, most importantly, verbs. The modelling of verbs has been omitted in many state-of-the-art systems and models, where the focus is often put on the noun phrase level. In this master project, we model and implement a verb-centered sentiment analysis system for French. We build a French polarity lexicon resource, where we propose a new polarity tag-set with fine-grained polarity labels, and where we specify the semantic behaviour of a selection of verbs. For the implementation of the system, we adopt a rule-based and compositional approach, where the focus is put on the role and function of verbs. We show how the inclusion of a verb-component increases our model's performance, albeit not as significantly as we expected, due to its dependence on lexicon coverage and parser accuracy. Finally, we conduct a number of empirical analyses in order to build a theoretical model for the assessment of so-called polarity conflicts, in view of possible extensions of the present system.

Zusammenfassung

Sentimentanalyse ist ein anspruchsvolles und sehr vielfältiges Gebiet der Computerlinguisitk. Wörter und Phrasen müssen nicht nur auf Hinblick ihrer grammatikalischen und syntaktischen Funktion erkannt und markiert werden, sondern auch bezüglich ihrer Polarität. Diese kann entweder positiv, negativ oder neutral sein. Solche polaren Wörter können überdies nicht isoliert betrachtet werden, da sie von anderen Satzelementen, wie zum Beispiel von weiteren polaren Wörtern, Präpositionen, Adverbien, oder Verben umgeben und modifiziert werden können. Die Modellierung von Verben wird jedoch in vielen führenden Arbeiten nicht mit einbezogen: Der Schwerpunkt liegt meistens auf der Nominalphrasen-Ebene. In dieser Masterarbeit implementieren wir ein verb-zentriertes Sentimentanalyse-system für Französisch. Wir erstellen ein französisches Polaritätenlexikon, für das wir sowohl ein neues, feineres Polaritäten-Tagset vorschlagen als auch die polaren Rahmen einzelner Verben spezifizieren. Für die Implementierung des Systems wählen wir einen regelbasierten und auf dem Kompositionalitätsprinzip basierenden Ansatz, bei dem der Schwerpunkt auf die Rolle und Modellierung der Verben liegt. Wir zeigen wie das Einbinden der Verbkomponente zu einer Verbesserung unseres Systems führt. Diese ist allerdings von einem kleineren Ausmass als erwartet, da die Leistung des Systems stark von der Lexikonabdeckung und der Präzision des Parsers abhängig ist. Schliesslich führen wir empirische Analysen durch, um ein theoretisches Modell zu erstellen, das im Hinblick auf mögliche Erweiterungen des aktuellen Systems der Verrechnung von sogenannten Polaritätskonflikten dienen soll.

Acknowledgements

I would like to express my gratitude to all the people who helped and supported me during my Master project.

First, I would like to thank Dr. Manfred Klenner for agreeing to supervise my thesis and for allowing me to carry it out in the context of a research project at the Institute of Computational Linguistics. I am grateful for his help and efforts to introduce me to the domain of Sentiment Analysis, which I was rather unfamiliar with at the beginning, and for the constructive and stimulating input he provided me with throughout the project. I would also like to thank Prof. Martin Volk, who encouraged me to take this daring step of carrying out my thesis in an unfamiliar domain and moreover for French, a language that has received less attention than others in the domain of NLP.

I am indebted to the precious help and patience of Michi Amsler, who not only provided me with the basis of my system, but also spent many hours helping me with technical and practical issues. In this sense, I would also like to thank Nora Hollenstein, who also helped me in a number of technical matters. I am also very grateful to my friends Esther Germann and Franziska Tobler for proof-reading my thesis, and to Katrin Rettich for helping me with initial formatting problems.

Last but not least, I would like to thank my family and friends for supporting me throughout my thesis. I would like to address special thanks to my companion, Max Michels, and to my father, Giorgio Tron, who always supported me during my studies. I would lastly like to dedicate this work to my late mother, Alda Muratore.

Contents

List of Figures

List of Tables

List of Acronyms

CONLL	Conference on Computational Natural Language Learning. Refers to the data format for storing sentence structure, more precisely to the format of dependency parsing output.
CG	Constraint Grammar. Used interchangeably with VISL or VISL CG-3, to refer to the CG format or implementation tool of the VISL group.
DG	Dependency Grammar
NEG	Negative
NLP	Natural Language Processing
NP	Noun Phrase
PP	Prepositional Phrase
PoS	Part of Speech
POS	Positive
UZH	University of Zurich
VISL CG-3 or VISL	Refers to the constraint grammar tool of VISL ("Visual Interactive Syntax Learning"), a research and development project of the University of Southern Denmark. The tool itself is called VISL CG-3. We will often use the abbreviation "VISL".

1. Introduction

1.1. Motivation

Web 2.0 and social media are two interrelated phenomena that have emerged around 10 to 15 years ago and have continued to grow and gain in importance ever since. The world wide web has become a platform where users do not only gather but also share information. This information can be objective or subjective. Objective pieces of information are factual explanations or descriptions, such as the description of scientific phenomena on wikis for example. Subjective information includes opinions and evaluations, e.g. in product reviews on an online retailer website, as well as feelings and emotions, for example in blog-entries recounting a particular event or experience. The resulting phenomenon of "big data", i.e. of very large quantities of data collections that cannot be managed and analysed manually anymore, has led to the fast emergence and development of new NLP tasks, among which *sentiment analysis* resides.

In an industrial or commercial context for example, companies are interested in knowing as much and as fast as possible about how popular their products are. Online reviews and customer e-mails, but also people expressing themselves on social media platforms, constitute a very large source of raw information. Sentiment analysis tools allow to quickly and automatically analyse and synthesise this data: they assess whether the opinions expressed in a collection of texts (e.g. reviews or articles) are *positive* or *negative*, and deduct the dominant opinion that is being expressed. This allows the concerned companies to rapidly respond to their customers' feedback and demands. Sentiment analysis is also increasingly used with similar aims in the domain of reputation analysis, for example in the context of electoral campaigns.

Besides its applicability, sentiment analysis is a fascinating topic also from a purely NLP-research perspective. A performant sentiment analysis system is based not only on thorough syntactic and grammatical knowledge, but also on semantic knowledge. Indeed, semantic knowledge is necessary in order to assess whether a proposition is *polar*, i.e. positive or negative, neutral (not polar) or ambivalent (bi-

polar, i.e. both negative and positive). Basic or prior semantic knowledge is induced to the system via a so-called *polarity lexicon*. A polarity lexicon consists of an extensive list of lemmas, i.e. words in their base form, which are marked with both their *polarity* (or *valence*) and word class. This list enables the system to identify a word in a text as polar or neutral. However, polar words cannot be regarded in isolation. They can for example be negated, e.g. "good" vs. "not good", and/or appear together with words of the opposite polarity, e.g. "disappointed hopes". In order to assess the meaning or, more precisely, the polarity of such cases, the system needs additional semantic knowledge. On the one hand, the system has to be able to recognize words such as "not" as a so-called modifier, in this case as a *shifter*. For this, modifiers have to be included in the polarity lexicon. Other types of modifier are *diminishers*, e.g. "less", and *intensifiers*, such as "more".

On the other hand, the system also needs to know how to process such phrases. This can be done via hand-crafted *rules*, which tell the system that, for example, if a positive adjective is modified by a shifter, such as in the case of "not good", the positive polarity of the adjective is to be reversed and to be computed as negative. In the case of noun phrases with two opposing polarities, such as "disappointed hopes", other rules are going to determine that the noun phrase is to be computed as negative. It is however important that the rules consider several layers or levels or *compositionality*, in order to correctly assess a case such as "no disappointed hopes" for example. Similarly, verbs and their arguments have a different connotation or meaning with regard to their composition. The phrase "to encourage charity" for example, has an overall positive polarity (or connotation), whereas the phrase "to encourage criminal intentions" does not. Many state-of-the-art sentiment analysis systems do not include verbs in their models. We believe that the inclusion of verbs is important and productive, as the two aforementioned examples suggest. Moreover, so-called polar verbs can also carry polar meanings by themselves, for example when they are used with neutral objects, e.g. "to encourage a friend/a project". The same is true for intransitive verbs, e.g. "the project failed".

This is why we chose to model a rule-based and verb-centered sentiment analysis system, based on the principle of compositionality. We moreover chose to develop this system for French, because as in many other areas of NLP, little research has been done so far for French sentiment analysis. Available resources (e.g. polarity lexicons or dependency parsers) are also rare, but add to the challenging aspect of the project. Moreover, the choice of French enabled us to carry out the project in a multi-lingual context, namely in the context of the *ARGUS project*. The ARGUS project is a commercial research project resulting from the cooperation between the UZH Institute of Computational Linguistics and the ARGUS company. *ARGUS*

der Presse AG[1] is a Swiss company based in Zurich, which offers information management services. These services include media monitoring, media intelligence and media analysis. The project aims at developing a competitive sentiment analysis system for German, English and French. The project director is Dr. Manfred Klenner. Michael Amsler developed the German prototype of the system, which served as a basis for the other two languages. Nora Hollenstein was responsible for the English version of the system.

1.2. Research Questions

As mentioned above, the basic semantic knowledge that we need to provide our system with consists of the a priori polarity at the word-level, namely the polarity lexicon. We think that distinguishing only between polar categories (positive, negative and neutral) and modifiers (negators or shifters, intensifiers and diminishers) is insufficient, especially with regard to the compositionality principle on which we are going to base our model. That is, words cannot be considered in isolation. The meaning and polarity of the noun phrase "disappointed hopes" differ for example from the ones of the NP "fulfilled hopes". Similarly, we believe that there are different types of negative and positive polarities, which is important on an upper level of compositionality. Compare for example the two negative nouns "victim" and "criminal". The phrase "to support the victims" is going to result in a different overall polarity than "to support the criminals", due to the respective connotation of the object. In other syntactic frames, the subject's polarity might be the determining element of the sentence's overall polarity. Consider for example "The evil plan succeeded" vs. "The (good) plan succeeded."

In other words, we expect polar verbs to behave in specific ways, by having certain polarity *effects* or *expectations* on their arguments. In the case of "to support", we expect the supported object to be positive, in the sense that it deserves support: it is either something neutral that is viewed positively, e.g. "to support an initiative", or something that is agreed upon to be positive by common sense or world knowledge, e.g. "to support charity". If the object carries a negative polarity, a so-called *polarity conflict* can occur. However, whether a polarity conflict occurs depends on the *type* of polarity that is involved, as the aforementioned examples "to support the victims" vs. "to support the criminals" suggest. A fine-grained distinction between polarity types, which enables to differentiate between factual, emotional and moral polarity types, is needed. If a polarity conflict occurs, the

[1] www.argus.ch, last accessed on 20th October 2013.

unexpected negative polarity of the object is going to propagate over to the verb's positive polarity, by reversing it, as well as to the subject. Polarity conflicts can also occur inside a noun phrase, when a positive adjective modifies a negative noun, or vice-versa. Consider for example the noun phrase "an embarrassed smile" vs. "a hypocritical smile". Here too, we believe that the distinction between different polarity types can help to nuance and to determine the overall polarity of the NP with greater precision, which in turn will allow for a more concise assessment of the verb phrase and hence of the sentence.

In summary, the aims of the present project consist of:

- Defining fine-grained polarity labels, which allow to differentiate between different types and/or intensities of negative and positive polarities.
- Building a polarity lexicon, where adjectives, adverbs and nouns are marked with the aforementioned fine-grained polarity labels.
- Extracting the most frequent polar verbs out of the corpus and building a verb-polarity frame resource, where all possible frames, i.e. all possible behaviors or scenarios of the verb, are specified.
- Elaborating and implementing compositional aggregation rules on the noun and prepositional phrase level, as well as on the verb phrase and sentence level. The latter rely on the verb-polarity frame specifications.
- Automatic evaluation of the system at the document-level.
- Conducting empirical analyses of polarity conflict phenomena and deriving theoretical model for conflict resolutions from the results.

1.3. Thesis Structure

In the following chapter, we are going to give an overview of the related literature and the state-of-the-art in the domain of sentiment analysis (2.1), followed by a short account of the theoretical grammatical background and of the tool that we used in the elaboration of our system (2.2). Chapter 3 is dedicated to the data that we used, namely our corpus (3.1), and to the resources that we built, which are the polarity lexicon on the one hand (3.2), and the verb-polarity frame resource on the other hand (3.3). In chapter 4, we describe the core of the sentiment analysis system, namely its pipeline, which can be divided into three steps. First, the linguistic preprocessing and formatting of the input data (4.1), which subsequently allows us to mark words with their priori polarity (4.2) and then to carry out the subsequent,

actual sentiment analysis (4.3). Chapters 5 and 6 are both concerned with evaluation tasks. In chapter 5, we deal with the automatic evaluation of our system at the document-level. Chapter 6 features more manual analyses, namely the empirical analyses of polarity conflicts involved in noun phrases (6.1) and in verb-frames (6.2). The final chapter 7 wraps up the results and insights we gathered from our project, along with a number of concluding remarks.

2. Related Work and Theoretical Background

This chapter is divided into one section on related work and literature, and into one section outlining the theoretical background and tools on which the project is based.

2.1. Related Work

The domain of sentiment analysis has received growing attention in the last 10 to 15 years (Liu, 2012, p. 4). Yet, although it might appear as a young discipline, the amount of related literature and research is significant. The approaches and focuses are numerous and varied, and, for this reason, we would like to begin this chapter by clarifying the approaches, concepts and focuses of the studies.

First, the general term *sentiment analysis* needs to be specified. In related work, the distinction between the terms *opinion mining* (or *opinion analysis*) and sentiment analysis is often made. As the terms themselves suggest, *opinion mining* refers to the expressions of opinions about given entities, e.g. opinions about a product in online reviews (Hu and Liu, 2004). *Sentiment analysis* however, is often employed for the analysis of emotions and feelings in texts, for example people expressing themselves about certain topics in social media or in blogs (Klenner et al., 2009b). The tasks and approaches differ slightly, but the field of study remains the same. Therefore, we use the terms sentiment analysis and opinion mining interchangeably for both the analysis of opinions and feelings, since both are covered in our project. In this chapter, however, we will make the distinction whenever it is relevant to the cited work.

The approaches vary more substantially in the methods they employ, and in their levels and focus of analysis. The employed methods can be broadly divided into machine-learning approaches and rule-based approaches. Machine-learning approaches often feature a "bag-of-word" representation, such as in (Hu and Liu, 2004). The approach often consists of automatically analysing n-grams, looking for recurrent combinations of opinion words. Rule-based approaches on the other hand aim

at capturing and modelling linguistic knowledge in rules. Rules may concern syntax or semantics, or both. The level of analysis can be at the document-level, sentence-level, clause-level, or even at the word-level. Most research combines two or more levels of analysis. The focus of analysis can consist of the type of employed resources (e.g. lexicon) and in the way they are generated; it can also concern specific linguistic phenomena and the way they impact sentiment analysis (e.g. negation).

This section is going to be divided into three sub-sections, which correspond to three substantial topics relevant to our project: first, the generation and properties of polarity lexicons. Second, the compositionality principle in sentiment analysis with regard to the assessment of polarity at NP level and, more importantly, with regard to the role of verbs in sentiment analysis. The third section is going to deal with the complex topic of negation.

2.1.1. Sentiment Analysis Lexicons

As far as the generation of lexicons for sentiment or opinion analysis is concerned, two main methods can be found in related work: manual compilation, such as in (Wilson et al., 2005) and automatic or semi-automatic generations, as in (Esuli and Sebastiani, 2006) and (Neviarouskaya et al., 2009b).

Esuli and Sebastiani (2006) use the free lexical database WordNet[1] to build a lexical resource for opinion mining, called *SentiWordNet*. The procedure consists of a quantitative analysis of the *glosses* (definitions) of each *synset* (i.e. of each set of synonyms). Three numerical scores, *Objective*, *Positive* and *Negative* are associated to each synset. The scores describe how positive, negative or objective the terms of the synset are. The aim is thus to classify the senses of a term, and not the term itself (Esuli and Sebastiani, 2006, p. 417). This is motivated by the assumption that different senses of the same term can have different connotations. A version of WordNet, WOLF (Sagot and Fišer, 2008), is available for French. However, since the synsets were automatically generated, the resource still contains errors and unsatisfying quality when ambiguous words are processed (Sagot and Fišer, 2008, p. 18). Hence, to our knowledge, no similar resource to *SentiWordNet* is available for French opinion analysis. This is why we chose to compile our French polarity lexicon manually (as described in section 3.2.2).

Out of the *SentiWordNet* resource, several authors built a lexicon for their specific research purposes, as for instance in (Fahrni and Klenner, 2008). Neviarouskaya et al. (2009b) build and automatically expand a topical polarity lexicon, called *SentiFul*. The lexicon is topical because it only contains polar words describing or

[1] http://wordnet.princeton.edu, last accessed on 13th October 2013.

relating to one specific topic, namely emotions. Moreover, it is restricted to content words, i.e. adjectives, adverbs, noun and adverbs (Neviarouskaya et al., 2009b, p. 363). The automatic expansion of the lexicon is done via the synonymy relations available in *SentiWordNet*, and the derivation and scoring of morphologically modified words. The authors further stress the importance of scalability (or intensity), which is expressed via "emotional vectors" ranging from 0.0 to 1.0 (Neviarouskaya et al., 2009b, p. 363). The aim is to assess a term's affiliation to one or more emotions (e.g. the adverb "pensively" is affiliated to both the emotions "sadness" and "interest"), and whether it is positive, negative or both (Neviarouskaya et al., 2009b, p. 364). Polarity scores and weights determine if a word is more or less positive or negative. Since the object of analysis are emotions, the specificity of the lexicon (i.e. the topical focus on emotion words), as well as the concept of scalability, suit the needs of the project well.

However, as pointed out in other works (Polanyi and Zaenen (2006); Wilson et al. (2005); Moilanen and Pulman (2007); Klenner et al. (2009a) and (2009b)), lexicons ought to include also so-called *modifiers*, such as shifters, intensifiers and diminishers. This is especially important if a compositional approach is adopted (see section 2.1.2). From a less technical or structural point of view, and from a more semantic perspective, a purely polar lexicon, i.e. a lexicon where only the distinction between subjective (positive or negative) and objective (or neutral) is made, is insufficient. A common method is to add intensity values to the polarity of words. Yet, this does not enable a semantic insight into the nature of the evaluative language that is being used. As outlined by Whitelaw et al. (2005), interest in *non-topical* sentiment analysis has been growing. In other words, sentiment analysis that is applied to texts of which the topic or domain is not known a priori. The results of the sentiment analysis are thus also aimed at revealing what *type* of subjectivity we are dealing with. In fact, when dealing with subjective or evaluative language, the type of *attitude* expressed is relevant: whether an emotion, a factual opinion or a judgment (of moral or social dimension) is expressed gives a different weight and/or character to the analysed stance or text.

In order to assess this finer granularity of subjective language, a number of projects, for example (Neviarouskaya et al., 2010a) and (Zhang et al., 2012), lean on the *Appraisal Theory*. The Appraisal Theory is a relatively recent linguistic theory, elaborated by the linguists Martin and White (2005), about evaluative language in English. The authors define appraisal as "one of three major discourse semantic resources construing interpersonal meaning (alongside involvement and negotiations)" (Martin and White, 2005, p. 34). Appraisal itself is divided into three interacting categories, *attitude*, *engagement* and *graduation*:

> **Attitude** is concerned with our feelings, including emotional reactions, judgements of behaviour and evaluation of things. **Engagement** deals with sourcing attitudes and the play of voices around opinions in discourse. **Graduation** attends to grading phenomena whereby feelings are amplified and categories blurred. (Martin and White, 2005, p. 35, emphasis in original)

Graduation is subdivided into *Focus* and *Force*. It can apply to both Attitude and Engagement and is concerned with "gradability" (Martin and White, 2005, p. 37), which is comparable to the concept of intensity or scalability, as found elsewhere in related work. Engagement concerns the way the speaker positions himself in the texts and towards the expressed opinions ((Martin and White, 2005, p. 36); (Zhang et al., 2012, p. 284)). Attitude is the "framework for mapping feelings as they are construed in English texts" (Martin and White, 2005, p. 42), and is the category or concept of Appraisal Theory which has received most of the attention in the domain of sentiment analysis.

In Appraisal Theory, the category of Attitude is itself divided into 3 subcategories: *affect*, *judgment* and *appreciation*. **Affect** describes negative or positive feelings and emotions, such as "happy" or "grief" (Martin and White, 2005, pp. 35, 42, 45 et seq.). **Judgment** is "concerned with resources for assessing behaviour according to various normative principles": it deals with "attitudes toward behaviour which we admire, or criticise, praise or condemn" (Martin and White, 2005, pp. 35, 42, 52 et seq.). In other words, the category of Judgment contains a moral and/or social dimension, directed towards human behaviour. The category of **Appreciation** includes words or expressions used to evaluate "things" or phenomena, according to "the ways they are valued or not in a given field" (Martin and White, 2005, pp. 36, 43, 56 et seq.). An illustration for Appreciation would be, for example, the language used in an evaluation of a product and its properties in a costumer review.

As already mentioned, past studies are based on Appraisal Theory, especially as far as the lexicon's properties are concerned. Neviarouskaya et al. (2010a) expand their *SentiFul* lexicon (Neviarouskaya et al., 2009b) and the coverage of their model to fine-grained types of attitude. In addition to their previous nine affect (or emotion) categories, four polarity labels for judgement and appreciation (positive and negative respectively), as well as a "neutral" label, are included. The authors follow a compositional approach in which they include verbs (cf. 2.1.2.2) and their interaction with different types of attitude words. As far as the differentiation between different types of attitude is concerned, the authors obtain satisfying results, with a slight improvement in overall accuracy when the nine emotion labels are contracted into two encompassing labels, i.e. positive and negative affect (Neviarouskaya et al.,

2010a, p. 86).

Zhang et al. (2012) describe a compositional model for French sentiment analysis. The model is inspired by linguistic approaches, but the authors regret the current absence of an "unified theory for the notion of evaluation" in French linguistics, especially as far as the description of different types of attitude or evaluation is concerned (Zhang et al., 2012, p. 284). Nevertheless, the authors find the principles of Appraisal Theory from Martin and White (2005) to be applicable for French. However, although they give a detailed account of the three categories of Appraisal Theory (attitude, engagement and graduation), the model only includes the component of graduation and does not appear to feature a lexicon with fine-grained polarity labels as described in (Neviarouskaya et al., 2010a).

In our project, we lean on the Appraisal Theory for the definition of the polarity labels of our lexicon as well. More precisely, we lean on the category of Attitude and its taxonomies (i.e. affect, judgment and appreciation), albeit with a number of adaptations and expansions, which we are going to describe in section 3.2.1.

2.1.2. The Compositionality Principle in Sentiment Analysis

According to Polanyi and Zaenen (2006), early works concentrate only on the opinion or sentiments conveyed by individual words. When implemented directly, e.g. when the method consists of assessing opinion in a text by simple counting of positive and negative instances, the results are often wrong or misleading (Polanyi and Zaenen, 2006, p. 1). As in many other NLP tasks, words, in this case opinion words, cannot be considered in isolation. This is best illustrated and demonstrated with the principle of compositionality, as defined in (Dowty et al., 1981, p. 8):

> *The meaning of the whole is a function of the meaning of the parts and their mode of combination.*

In other words, the meaning of an expression, be it at phrase or sentence level, is determined by its structure (syntax) and/or the meanings of its constituents.

2.1.2.1. Basic Models and Implementations

In the context of sentiment analysis at the phrase-level, the polarity of a noun phrase is indeed determined by the polarity of its constituents: a neutral noun can be modified by a polar adjective and become polar itself (e.g. "sad news"), or a polar noun can be modified by an adjective that has an opposed polarity (e.g. "deceived

hopes"). Klenner et al. (2009a and 2009b) develop a tool for German sentiment analysis, called *PolArt*. The authors implement a cascade of transducers in order to compute the polarity on NP level. The system is based on the assumption that compositionality takes place in a straightforward and canonical way, and can therefore be captured by a limited number of hand-crafted rules. If a NP is composed of a positive adjective and a negative noun for example, as in "a perfect misery", the NP polarity is negative (Klenner et al., 2009b, p. 181). Regularities for NP-PP composition are also formulated in rules and integrated in the cascade of transducers. The system receives as input chunked text in a flat structure; simple rules take effect first, moving from word to sentence level. The system copes well with negative classifications and moderately well with positive classification, but poorly with neutral classification. This is, however, due to the evaluation method and not to the adopted approach itself (Klenner et al., 2009b, p. 183). Petrakis et al. (2009) adapted the PolArt system for French. Although the rules are domain-independent, they are language-specific, especially with regard to NP-PP rules, where the preposition needs to be specified in each rule. To the date of the paper, only 10 rules were elaborated (for 70 rules for the German system) and the performance for the French system was not evaluated.

A similar approach to the one in (Klenner et al., 2009a and 2009b) can be found in (Moilanen and Pulman, 2007). Rules of *sentiment propagation* (a neutral polarity is overridden by a non-neutral one), *polarity conflict* and *conflict resolution* (in the case of compositions with conflicting polarities) and *polarity reversal* are applied to dependency structures (see section 2.2.1). As in (Klenner et al., 2009b), so-called modifiers, which can pertain to any word class, are considered to impact the intensity of a words' valence (or polarity), or the valence itself (e.g. "good" vs. "very good" vs. "not good"). From a compositional perspective, it is crucial to include modifiers and their impact in an opinion mining model. Earlier works focusing on modifiers from a compositional perspective often consider shifters to invert polarity and not to affect (if present) the intensity. As illustrated in (Polanyi and Zaenen, 2006, p. 2), if the word "clever" has the positive intensity value of 2, its negated (or shifted) form "not clever" is reversed to -2, i.e. to negative with intensity value 2.

In French sentiment analysis, a few interesting hypotheses have been formulated concerning modifiers in a compositional approach. Especially in the case of shifters or, more broadly, of negation (section 2.1.3), Zhang et al. (2012) argue that simple valence inversion is inaccurate if more than one modifier is involved and if the intensity of polarity is to be assessed. The opinion expression "not extremely good" for example, cannot be equated with "extremely bad" (Zhang et al., 2012, p. 285). A rule-based prototype model was implemented to compute intensity values of opinion adjectives with regard to the impact of modifiers (i.e. shifters, intensifiers

and diminishers). Due to data sparseness, however, the model's potential could not be assessed completely (see section 2.1.3 for further details).

As far as the basic implementation of the compositionality principle in sentiment analysis is concerned, namely the impact of modifiers on polar words, and the NP and PP composition, we adopt a similar approach to the one reported in (Klenner et al., 2009a and 2009b) and (Moilanen and Pulman, 2007). This is described in sections 4.3.1 and 4.3.2. However, we take into consideration the concept of polarity intensities only to a small extent in our model (cf. section 3.2.1). Moreover, we do not consider polarity intensities with regard to the impact of modifiers, as it is the case in (Klenner et al., 2009b) or (Zhang et al., 2012). Here, the authors claim that e.g. "not extremely good" has a different polarity intensity than "not good" (Klenner et al., 2009b, p. 180), or that "not very good" does not equal "very bad" (Zhang et al., 2012, p. 285). Although we deem the two latter approaches as preferable (cf. also section 2.1.3), we stay closest to the approach that is presented in (Polanyi and Zaenen, 2006), namely the straight-forward inversion of the concerned polarities.

2.1.2.2. Compositional Models featuring Verb-Components

A fully compositional approach of course presupposes that verbs are also taken into account. Moilanen and Pulman (2007) also include verbs into their compositional model. Since they base their theory on dependency grammar (see section 2.2.1), the assumption is that in the clausal head, the verb is generally stronger than its complements. This suits best the aforementioned rules of *propagation* (e.g. "failed all his exams"), of *polarity conflicts* (e.g. "spoiled the beautiful ceremony") and of *reversal* (e.g. "prevented the war") (Moilanen and Pulman, 2007, p. 3). Subject and objects are only stronger than their predicator if they are ranked as shifters (e.g. "nobody died"). In a similar manner, Neviarouskaya et al. (2009a) elaborate rules for the compositional assessment of affects on sentence-level. After an analysis on the word level, the "affective content" of phrases is processed (Neviarouskaya et al., 2009a, p. 279). In verb-phrases, copula verbs inherit the polarity of adjectives; in the case of transitive verbs, the verb's valence is also considered to be dominant, also if the noun phrase is of opposite valence. Experiments show that taking into consideration all grammatical levels of a sentence sensibly increases precision and recall, and thus also averaged accuracy.

In (Neviarouskaya et al., 2009c), the authors expand their approach by proposing a rule-based method where rules for semantically distinct verb classes are elabo-

rated. The authors base their investigations on the VerbNet resource[2], a verb lexicon where verb classes are organised according to the syntactic and semantic coherence among members of a class. 73 verb classes were found to be relevant for the task of opinion analysis and were further contracted into 22 classes, which are "differentiated by the role that members play in sentiment analysis and by rules applied to them" (Neviarouskaya et al., 2009c, p. 31). The first verb class for example, *psychological state or emotional reaction*, is subdivided into three subclasses: *object-oriented emotional state* (e.g. "to appreciate"), *transitive subject-driven change in emotional state* (e.g. "to charm") and *intransitive subject-driven change in emotional state* (e.g. "to appeal to"). In other words, verbs are attributed to classes and subclasses according to their subcategorisations and to the semantic functions or impacts of their arguments. For each verb class, a set of rules is developed for the task of sentiment analysis on phrase level. Most verbs are annotated with prior polarity; the sentiment features (i.e. the polarity) of the phrases are however context-sensitive and need to be defined by means of specific rules. In the case of an object-oriented class for example, not only the polarity of the object (with respect to the prior polarity of the verb) is taken into account but also the distinction between interior and exterior perspective. In the case of the phrase "he admires the mafia leader" for example, the interior view is going to be computed as positive, but the exterior as negative (Neviarouskaya et al., 2009c, p. 32).

These rules are implemented in (Neviarouskaya et al., 2010a), together with a fine-grained polarity lexicon and other compositional rules. The system achieves satisfying results (as mentioned in section 2.1.1 above), but the focus is put on the granularity of the attitude labels. Whether the inclusion of rules for semantically distinct verb classes also had a positive impact on the performance of the system is not known. In (Anand and Reschke, 2010) a similar approach is adopted. The authors suggest that verbs should be analysed as "evaluative functor classes", i.e. as mappings to an evaluative value and based on their lexical entailments about their arguments. They identify three types of *entailment*, where a change of state is operated: *possession* by a participant, *existence* of a participant and *affectedness* of a participant (Anand and Reschke, 2010, p. 1). These three change of state classes are divided into two positive and negative sub-classes respectively: "gain and loss", "creation and destruction" and "benefit and injury". In the case of the *existence entailment* for example, the assumption is that if something negative exists, the evaluation is negative. If something negative is destroyed, i.e. if its state is changed into non-existent, the evaluation is positive, as in the case of "destroy weapons" (Anand and Reschke, 2010, pp. 4-5). An annotation study on constructed sentences

[2]http://verbs.colorado.edu/ mpalmer/projects/verbnet.html, last accessed on 13th October 2013.

and a corpus investigation attest the predictions of the theory (Anand and Reschke, 2010, pp. 5-6). No implementation is proposed, but the authors conclude with an interesting remark by stressing the importance of domain evaluation. By referring to the Appraisal Theory, they claim that, for example, the gain of something aesthetically valuable may not give rise to the same evaluative intuition as the gain of something ethically valuable (Anand and Reschke, 2010, p. 6). Thus, further investigations need to show how mapping across finer-grained domains correlates with the semantic aspects of verbs.

For French, no work focusing on or at least including the modelling of verbs is known, except for (Paroubek et al., 2010). The paper presents a project which aims at defining and implementing a model for opinion mining in an industrial context. An extensive review is given about the state of the art in sentiment analysis. However, the authors do not reveal the methods employed in the project, but only the linguistic theories which formed the basis for the model. Those are, mainly, the Appraisal Theory from Martin and White (2005), and the work of Yannick-Mathieu (1991). Yannick-Mathieu's linguistic work deals with the classification of French "sentiment verbs". It is rooted in the field of lexico-grammar, but with a number of substantial differences. Indeed, the theory is aimed at being transposed into a NLP tool able to process sentences which contain "sentiment verbs". While in a lexico-grammatical approach lexical items are usually classified according to their characterising syntactic features, Yannick-Mathieu claims that especially verbs expressing feelings or evaluative stances (which she labels as *verbes de sentiment* "sentiment verbs") should not only be classified according to their syntactic structures, but also and more importantly according to their semantic field (Yannick-Mathieu, 1991, pp. 9-11). The verbs are divided into homogeneous semantic classes, and further refined into classes which share common syntactic properties. Indeed, this classification is considered to be useful because especially in the case of sentiment verbs, the assumed systematic correspondence between syntactic function and semantic interpretation is often problematic. Consider the following examples from (Yannick-Mathieu, 1991, p. 10):

(2.1) *Ce roman passionne Marie.* "This novel fascinates Marie."

(2.2) *Ce roman plaît à Marie.* "This novel pleases Marie."

(2.3) *Marie aime ce roman.* "Marie likes/loves this novel."

As the examples show, it is important to take both syntax and semantics into account, in order to assess what kind of emotion is directed towards which argument of the phrase, and from what argument it originates. The hypothesis is that there is a semantic link between verbs that do not necessarily share the same syntactic

properties (Yannick-Mathieu, 1991, p. 11). This is a very pertinent hypothesis and we are going to re-consider and discuss it in section 6.2.3.

Whereas the approaches described in (Moilanen and Pulman, 2007) and (Neviarouskaya et al., 2009a) constitute in our opinion a simplification of the role and function of verbs in sentiment analysis, the last works that we described above (Neviarouskaya et al. (2009c, 2010a); Anand and Reschke (2010); Yannick-Mathieu (1991)) share a number of commonalities with our approach on verbs. Indeed, they are all based on the same assumption, namely that both syntax and semantics of verbs need to be taken into account. More precisely, all elements in a sentence and their polarity are important and need to be considered with respect to the clausal head (i.e. the verb) in a bottom-up and compositional manner. We are going to classify verbs into individual classes as well, based on similar criteria as described in (Neviarouskaya et al., 2009c) (cf. section 3.3 and 4.3.3). Hand-crafted compositional rules are going to determine the verbs' semantic behaviour, as well as the effects and/or expectations on their respective arguments (section 4.3.3). We are not going to distinguish between exterior and interior view, but, rather, between default and non-default cases. Non-default cases can result in polarity conflicts: we assume, similarly to Anand and Reschke (2010), that the eventual polarity values (and intensities) that are assessed on clause and sentence-level are dependent on the involved polarity types (cf. 6, in particular section 6.2). To resolve polarity conflicts inside verb phrases, we are moreover going to consider the hypothesis by Yannick-Mathieu on potential semantic links between verbs that do not share syntactic properties (section 6.2.3).

2.1.3. Negation

As we mentioned in the previous section, negation, although it is closely related to the compositionality principle, is a topic for itself in the domain of sentiment analysis. Wiegand et al. (2010) present an encompassing survey on negation in sentiment analysis. The authors provide a chronological description of the aspects and methods for negation modelling. Those are broadly divided into two categories: methods used in the context of "bag-of-words" representations, where no prior knowledge of polar expressions is required, and models which do include prior knowledge of polar expressions. The latter typically follow a rule-based approach (at least partly) and include a polarity lexicon in order to obtain a prior knowledge of polar expressions (Wiegand et al., 2010, p. 62). In this case, Wiegand et al. distinguish between models where negation is modelled via *contextual valence shifters* and where it is encoded as *features* operating on expression-level polarity.

The first model corresponds to the one proposed by Polanyi and Zaenen (2006). Based on a lexicon, words are marked as positive or negative. Modifiers, which include shifters, intensifiers and diminishers, come to modify the valence of words and expressions (hence the term *contextual valence shifters*). Shifters are considered to be negators that simply invert the polarity of the word they modify, as in *not beautiful*. As mentioned above, this can be regarded as insufficient. The second approach described by Wiegand et al. (2010) corresponds to the model implemented by Wilson et al. (2005). Here, negation is encoded in features and combined with supervised machine learning. Wilson et al. (2005) distinguish three features: *negation features*, *shifter features* and *polarity modification features*. The *negation features* scan the four words preceding a polar expression for a negation expression. This makes it possible to find negated subjects such as "no politically prudent Israeli could support either of them" (example taken from (Wiegand et al., 2010, p. 63)). The *shifter features* basically function as in (Polanyi and Zaenen, 2006), with the difference that they include any type of word class and so-called negative and positive shifters, which only reverse a particular polarity type (e.g. the positive shifter "to abate" only modifies negative polar expressions, e.g. "abate the damage" (Wiegand et al., 2010, p. 63)). The *polarity modification features* describe cases where polar expressions modify or are modified by other polar expressions, as for example in *deceived hopes* (Wiegand et al., 2010, p. 63).

Besides presenting other approaches to negation in sentiment analysis, Wiegand et al. (2010) also address two pertinent facts concerning negation. The first is that a recurrent or standard approach is to "consider a negated polar expression, such as *not bad*, as an unnegated polar expression with the opposite polarity, such as *good*" (Wiegand et al., 2010, p. 66, emphasis in original). From a pragmatic perspective and as argued above, we agree that this is questionable. The second point is that most negation models are elaborated for the English language. However, even in the case of rather closely related languages, such as English and German for example, the structural differences are significant enough that they do not allow a simple transposition and application of the negation model from one language to the other (Wiegand et al., 2010, p. 66). Thus, negation in sentiment analysis is not only a phenomenon for itself, but also a language-specific phenomenon.

In French sentiment analysis, two works deal extensively with negation to our knowledge: (Benamara et al., 2012) and (Zhang et al., 2012). As mentioned above in section 2.1.2.1, Zhang et al. study the impact of negation on both polarity and intensity. Negation is operated by so-called valence shifters, which consist of negatives, intensifiers and diminishers. Considering other works on negation, the authors claim that current models are still incomplete. This is partly due to the little attention that has been paid to intensifiers and diminishers (Zhang et al., 2012, pp.

285): many works do not take into account that those types of valence shifters can belong to any type of word class; moreover, the "calculation of polarity modified by negatives or intensifiers/diminishers has often been done separately" (Zhang et al., 2012, pp. 285). Zhang et al. refer to a preferable approach presented in (Whitelaw et al., 2005), where, based on the taxonomies of the Appraisal Theory (Martin and White, 2005, cf. 2.1.1), the authors choose to reverse both polarity and intensity in cases such as *not very good*. Based on this, Zhang et al. claim that in French, the analysis of negation needs to be based on the previous analysis of the *force level* of an expression, in other words, the negation analysis needs to be based on a previous *graduation analysis*.

Focusing on adjectives, and taking Appraisal Theory as a basis, Zhang et al. (2012) follow the division of graduation into the two dimensions of force and focus. Force is concerned with the intensity of an expression or word. It includes any type of word class. The authors consider five different values of force: **low** – *un peu* "a little", **moderate** – *moyennement* "fairly", **standard**, **high** – *très* "very" and **extreme** – *extrêmement* "extremely". Adjectives without a priori intensity label are classified as "standard" (Zhang et al., 2012, p. 285). Four rules for negation are elaborated, where both intensity and polarity are taken into account. For example, if an extreme value is negated, e.g. *not extremely good*, the polarity is preserved and the intensity is lowered, based on the assumption that *not extremely good* signifies *a little good* and not *extremely bad*. Focus is also included in the model with modifiers that can either *sharpen* or *soften* the focus, such as *vraiment* "really" or *plutôt* "rather".

Zhang et al. (2012) claim that it is important to consider such modifiers in French negation because negated phrases with both focus and force modifiers can actually occur, such as in *les résultats ne sont pas vraiment très bons* "the results are not really very good" (albeit in our opinion, such cases rather occur in informal and/or spoken language). In such combinations, rules prescribe that focus modifiers are to be regarded as force modifiers. In the example above, *pas vraiment très bons* is then computed as *assez bons*, i.e. "a little good". However, although the method appears promising, it remains unclear whether it is effective. This is due, as the authors acknowledge themselves, to the restriction to adjectives and to the small applicative corpus used for evaluation: only 0.3% of the extracted sentences contain at least two modifiers. Even with regard to the corpus size, this suggests that the described phenomenon occurs only sparsely (Zhang et al., 2012, p. 289).

Benamara et al. (2012) adopt a more structural and grammatical approach. The authors focus on the phenomenon of "multiple negatives" in French. Based on a number of recent linguistic works, three types of negation are distinguished to describe this phenomenon (Benamara et al., 2012, p. 12):

- *Negative operators*, which always appear alone in a sentence and cannot be combined with each other, e.g. *sans* "without"
- *Negative quantifiers*, which express both a negation and a quantification, e.g. *aucun* "none"
- *Lexical negations*, which are implicit negative words, such as *carence* "deficiency"

Unlike negative operators, negative quantifiers can also occur in positive sentences, such as interrogatives. In negative sentences, they always co-occur with the adverb *ne*, and they can be combined with each other as well as with other negative operators (e.g. *nul ne doute de son innocence* "nobody doubts his innocence"). Lexical negations can also be combined with each other, as well as with the first two types of negation (Benamara et al., 2012, p. 12). The authors then formulate several possible combinations of negation types (negation categories) and their effect on polarity. Multiple negatives for example, usually preserve polarity, such as in *nul n'est venu* "nobody came" for instance. One exception is when for example a lexical negation co-occurs with negative operators as in *il ne manque pas de goût* "he does not lack (good) taste". Here, the polarity of the lexical negative *manquer* is canceled due to the presence of the negative operator *ne ... pas*.

The frequency of negative polarity words and multiple negatives is assessed in a corpus of French movie reviews with the aid of a manually built subjectivity lexicon from a previous project. 26% of the extracted reviews contain negative words and/or multiple negations. Benamara et al. further develop four hypotheses concerning the effect of each negation type (Benamara et al., 2012, p. 13):

1. a) The negation always reverses the polarity of an opinion expression, e.g. *exceptionnel* "exceptional" vs. *pas exceptionnel* "not exceptional".

   b) The strength of the opinion expression is lowered.

2. If an opinion expression is negated by a negative quantifier, the strength is greater than if it is negated by a negative operator, e.g. *jamais exceptionnel* "never exceptional" > *pas exceptionnel* "not exceptional"

3. A lexical negative has the same effect as a negative operator, e.g. *manque de goût* "lack of taste" = *pas de goût* "no taste"

4. When multiple negations are present, the scope of the negation has a greater impact, e.g. *plus jamais bon* "no longer ever good" > *plus bon* "no longer good"

The hypotheses are empirically evaluated against human subjects. One protocol tests the first hypothesis (both a) and b)) and a second one tests hypothesis 2 to 4. The protocols contain a set of questions based on the most frequent features found in the corpus. In the first protocol, each question presents two versions of a sentence, once negated and once affirmative. The questions concern the intensity of the opinion expressions. In the second protocol, each question contains a pair of sentences: one containing a negative operator, the other either a negative quantifier or a lexical negation. Here, the participants are asked to compare the intensity of the sentences in each pair. The results validate all four hypotheses about negation. Thus, Benamara et al. (2012) come to the same conclusion as Zhang et al. (2012), namely that negation affects both polarity and intensity. Additionally, they have identified the impact of negative quantifiers and multiple negations on the strength of the negation. This provides a good basis for a later implementation of a computational model for assessing sentiment orientation at the sentence or clause level (Benamara et al., 2012, p. 17).

As far as our model is concerned, we include negation only at the noun phrase level. Our model resembles the one described by Polanyi and Zaenen (2006), with the difference that, similarly to Wilson et al. (2005) and Benamara et al. (2012), we also consider lexical items besides modifiers in our rule-based implementation (cf. section 4.3.1.3). However, unlike Zhang et al. (2012) and Benamara et al. (2012), we do not take into account intensity in the case of multiple negations or in the case of negators combined with one or more modifiers: here, we assume a simple inversion of the polarities as well (cf. section 4.3.1.2) .

2.2. Theoretical Background and Tools

In this section, we would like to give a brief account on the theoretical background on which our sentiment analysis model is based, namely dependency grammar (2.2.1). Indeed, the first step in our pipeline consists of pre-processing the input texts via dependency parsing (section 4.1). The system and the tools used in it are then based on the dependency structure of sentences. Indeed, the dependency structure suits a compositional and verb-centered approach best. We think that this short overview will help to familiarise the reader quickly with the tool that we used, VISL CG-3, described in 2.2.2, which will in turn allow us to describe more accurately how we programmed our model (section 4.2 and 4.3 in particular).

2.2.1. Dependency Grammar

The starting point of the modern tradition of dependency grammar is often considered to take place with the work of French linguist Lucien Tesnière, *Eléments de syntaxe structurale*, which was first published in 1959 (Nivre, 2005, p. 2). Dependency grammar aims at defining and describing the syntactic structure of a sentence. Theories on dependency grammar present a large and diverse set of formalisms and approaches. However, they share a number of basic assumptions, in particular the assumption that "a syntactic structure consists of lexical elements linked by binary asymmetrical relations called dependencies" (Nivre, 2005, p. 2).

In fact, the concept of dependency is based on the idea that the words of a sentence are all connected through binary asymmetrical relationships among each other. In a sentence, words are no longer regarded as individual items: they are interconnected. These structural connections establish dependency relations between words and their totality forms the sentence ((Tesnière, 1959, pp. 11-13) quoted in (Nivre, 2005, p. 2)). A dependency relation usually describes the connection between a superior and inferior term. Other terms for the superior term are *head* or *governor*, and the terms *dependent* or *modifier* are synonyms for the inferior term.

As already mentioned, theories on dependency grammar vary, and so do the criteria for establishing dependency relations. The dependency relations of our parser (section 4.1) are basically established on the sentence and clause level. The head of a sentence is always the main verb, and its direct dependents are the heads of its arguments. The arguments are themselves composed of one head and one or more modifiers. In *L'étudiant lit un livre intéressant* "The student reads an interesting book" for example, the head or root of the sentence is the verb *lire* and its dependents are the heads of the subject-NP and of the object-NP. The subject-NP is composed of the noun *étudiant* (head) and of the definite article *l'* (dependent). The object-NP is composed of the head noun *livre*, which is modified by the indefinite article *un* and by the adjective *intéressant*. This is illustrated in Figure 1 below.

The reason why we choose to base our model on the principles of dependency grammar is that it fits well with the compositionality principle described in section 2.1.2 above: in a sentence, words are interrelated and together form the structure and eventually the meaning of the sentence. Moreover, the dependency structure posits the main verb as the root or the core of the sentence, which suits the verb-centered approach of our project well. In other words, dependency structures will allow us to describe and model compositional phenomena in the context of sentiment analysis, with the aid of the VISL CG-3 tool.

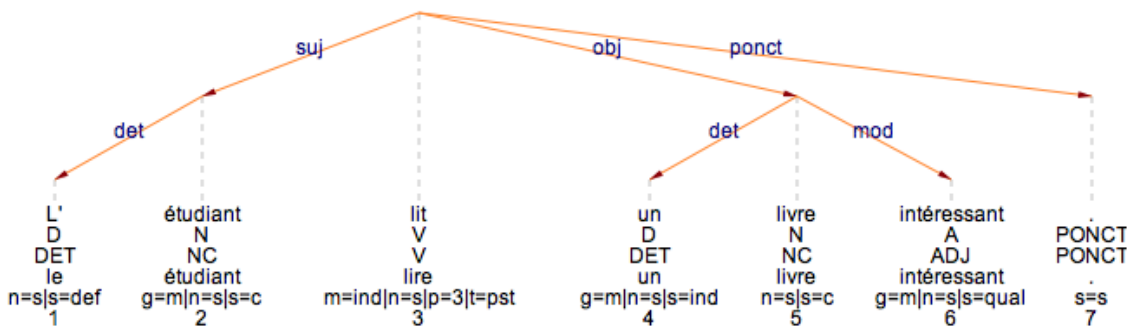

Figure 1.: Dependency Tree for the sentence "*L'étudiant lit un livre intéressant*".

2.2.2. The VISL CG-3 Tool

VISL stands for *Visual Interactive Syntax Learning* and is a research and development project of the University of Southern Denmark (SDU). One of their applications consists of a free constraint grammar compiler and parser, called VISL CG-3. It relies on a language independent CG formalism and is mainly used for context-dependent disambiguation of texts at a morphological and/or syntactic level. A word in a morphologically analysed text for instance, can generate a number of morphological analyses, as illustrated with the verb form *aime* "love(s)" in (2.4):

(2.4)
```
"<aime>"
    "aime"   aimer   V   1SG   PST   IND
    "aime"   aimer   V   2SG   PST   IND
    "aime"   aimer   V   3SG   PST   IMP
```

Example (2.4) illustrates a potential analysis of the verb form *aime* in VISL CG-3 cohort format. A single *cohort* consists of the word form displayed in one line, followed by at least one analysis line immediately below.[3] The VISL CG-3 tool enables us to formulate rules for disambiguation of, in this case, morphological analyses. For example, the rule `SELECT (1SG) IF (-1 (PRO) LINK 0 (je));` means "select the analysis line containing 1SG if at position -1 a pronoun and at the same position the lemma *je* are found". Or, in a more fluent formulation: "if the verb is preceded by the pronoun *je*, select the analysis first person singular (1SG)."

In other words, so-called constraint rules can be formulated by defining contexts and/or conditions regarding the relations between the individual cohorts. Indeed, VISL CG-3 also allows for a representation of dependency relations. Consider the

[3]The cohorts produced by our system always contain only one analysis line (see Figure 4 in section 4.1). Thus, in the next chapters, we will use the term "cohort" for the (single) analysis line.

two cohorts for the noun phrase *belle journée* "beautiful day" in (2.5):

(2.5)
```
"<belle>"
    "belle"  beau  C_ADJ  ADJ  DEP_mod  #4->5
"<journée>"
    "journée"  "journée"  C_NC  NC  DEP_ats  #5->2
```

The noun *journée* is the head of the noun phrase; the adjective *beau* modifies the head and is its direct dependent. This dependency relation can also be expressed in terms of parent-child relation: in this case, the noun is the parent node and the adjective its child. Constraint rules can also be written in terms of parent-child relations in VISL CG-3. This is a very useful component, which is extensively used in our sentiment analysis model, as described in chapter 4.

Operations in VISL CG-3 do not necessarily have to be formulated in terms of relations between cohorts. Simple tagging processes can also be performed via target definition, as for example in the prior polarity marking stage of the pipeline described in section 4.2. Tagging (or marking) can be done via two further features of VISL CG-3 that we extensively use in our implementation, namely the `ADD` and `SUBSTITUTE` rules. They basically function in the same way as the `SELECT` rules: the target, i.e. the word that is to be marked, is defined by restrictions or conditions concerning its properties or context, e.g. a specific word class is targeted or, on the contrary, excluded. If the defined target is found, the rule is applied: a user-defined tag is added, or substituted by another tag. A further useful feature of VISL CG-3 consists of the feature "`LIST`", which allows to merge a number of, for example, PoS tags into one category. For instance, if we would like to target all possible verb forms, we can, instead of enumerating each PoS tag several times, put them into one encompassing category: `LIST verbs = V VS VINF VPP VPR VIMP;`.

A detailed documentation of all VISL CG-3 functionalities and properties can be found on http://beta.visl.sdu.dk/.

3. Data and Resources

In this chapter, we introduce the data and resources of our project. We shortly describe our French news texts corpus, which is the data that we used for several empirical tasks (section 3.1). The largest part of the chapter is dedicated to the resources that we built in the context of our project: on the one hand the French polarity lexicon (section 3.2), and on the other hand the verb-polarity frame resource (section 3.3).

3.1. Corpus

We chose a corpus-driven approach for a number of empirical tasks. The tasks are listed at the end of this section. The corpus was kindly provided by Anne Küppers (Catholic University of Leuven, Belgium), who used it for a study on subjectivity in French news, in the context of the Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA) 2010 (Küppers and Ho-Dac, 2010).

Our corpus consists of two data sets. The first data set is a collection of articles published in one of the principal and most read Belgian newspapers in French, *Le Soir*. The articles were either collected from the printed version or from the online version of the newspaper. They cover general news and cultural topics. The second data set is composed of articles collected from the French online news platform *AgoraVox*. This data set covers a larger range of categories, and further differs from the *Le Soir* sub-corpus in that besides journalists, any subscribed user can contribute and revise articles. The articles of both sub-corpora were published between 2005 and 2009. The corpus was provided in a uniform text format:

1. Newspaper name
2. News section
3. Author(s)
4. Date

5. Heading(s)
6. Article

The *AgoraVox* sub-corpus is the larger one in terms of number of words; *LeSoir* is the larger data set considering the number of articles. The exact layout of the data sets is given in Table 1 below. Since the corpus was originally used for the assessment of the presence of the author and the varying use of subjective language in different French media, we believe it suits our project's needs well.

	Le Soir	**Agoravox**	**Total**
Words	2'744'270	3'281'208	6'025'478
Headings	715	896	1'611
Articles	5'873	4'368	10'241

Table 1.: Corpus properties, taken from (Küppers and Ho-Dac, 2010, p. 26).

We used the corpus for several tasks: first, it was used to expand the general polarity lexicon (see section 3.2.2.4). Second, it served for the extraction of the most frequent verbs (section 3.3.2). Third, a number of articles were used to increase the evaluation corpus provided by ARGUS (section 5.1). Finally, it was used for empirical analyses of the usefulness of fine-grained polarity labels in the context of polarity conflicts in noun and verb phrases (chapter 6).

3.2. Polarity Lexicon

The basis for a sentiment analysis system is the polarity lexicon, which contains word lemmas and their attributed polarities. The polarity lexicon for this project consists of a text list containing nouns, adjectives and adverbs and is going to be integrated into the system in the second step of its pipeline (section 4.2). Verbs are dealt with in a different way, as described in the sections 3.3 and 4.3.3.

The next section (3.2.1) outlines the set of the employed polarity tags along with explanations of their meaning. The labels were jointly established and annotated by the four participants of the ARGUS project, based on an available polarity lexicon for German from a previous UZH project (Clematide and Klenner, 2010). Based on this, we manually composed the French lexicon. This is described in section 3.2.2.

3.2.1. Polarity Labels

Many related projects (see section 2.1) use a basic set of polarity tags: positive, negative and neutral, e.g. for words such as "beautiful", "ugly" and "usual" respectively. Additional tags, such as shifter (e.g. "not"), intensifier (e.g. "more") and diminisher (e.g. "less") are often included as well. In order to nuance polarity strengths, a number of studies also include an intensity grade in the lexicon, which is then computed with other polarity values. For the ARGUS project, we agreed that we would use a new range of fine-grained polarity tags, based on an adaptation of Martin and White's (2005) Appraisal Theory (see section 2.1.1).

As illustrated in the examples below, a differentiation between different kinds of positive and negative polarity is useful if the aim is to get closer to the semantics of an utterance:

(3.1) headache *vs.* murder *vs.* hate

(3.2) smile *vs.* loyalty *vs.* love

In (3.1), the listed nouns are negative, and positive in (3.2). However, in direct comparison, it quickly becomes apparent that we are dealing with different kinds of positive connotations and with different kinds of negative connotations. Whereas "headache" and "smile" are of a factual nature, "murder" and "loyalty" contain a further social or moral connotation of appraisal. The nouns "hate" and "love", however, describe an emotion or feeling. Thus, based on the attitude categories from Martin and White's Appraisal Theory, "appreciation", "affect" and "judgment" (cf. section 2.1.1), we divided the two basic polarities, negative and positive, in 3 subpolarities respectively, and formulated 6 fine-grained polarity labels, as listed in Table 2 below.

Subpolarity	Negative	Positive
Appreciation	A_NEG	A_POS
Affect	F_NEG	F_POS
Judgment	J_NEG	J_POS

Table 2.: New fine-grained polarity labels, based on the Appraisal Theory.

As can be seen in Table 2, we adopted the same category names from Martin and White's Appraisal Theory for the subpolarities. The letter "A" in the polarity labels `A_NEG` and `A_POS` refers to the original attitude category "appreciation". For the "affect" category, the label `F_NEG` or `F_POS` are used. The abbreviation "F" stands for "feeling". The third label is `J_NEG` and `J_POS` respectively, where the

letter "J" stands for "judgment".

However, our polarity labels are a slightly modified version of the attitude taxonomies of the Appraisal Theory (cf. section 2.1.1). In Martin and White's Appraisal Theory, appreciation "involves evaluations of semiotic and natural phenomena, according to the ways in which they are valued or not in a given field" (Martin and White, 2005, p. 43). In other words, appreciation is concerned with aesthetics and/or (target-specific) values, according to specific criteria of evaluation. Words pertaining to the appreciation category express or refer to evaluations of things or phenomena, e.g. "captivating", "good" or "fake" (Martin and White, 2005, p. 45 and 56). Our appreciation subpolarity category includes this type of evaluative terms as well. In addition, we also include words that are not necessarily evaluative, but simply positively or negatively connotated, as for example "convalescence", or "indebted". Moreover, as the above definition suggests, target-specific evaluation terms also pertain to the appreciation category in Martin and White's Appraisal Theory. We exclude target-specific and/or highly ambiguous words, such as "woolly" for instance, which can mean "unclear" or "fuzzy" (e.g. "woolly film plot", "woolly Wi-Fi reception"), as well as "made of wool". Indeed, since we deal with non-topical sentiment analysis, i.e. sentiment analysis on texts whose domain or context is unknown, we avoid such terms in our polarity lexicon.

In other words, in addition to evaluative terms, our "appreciation" subpolarity category includes factual and objective terms that are either positively or negatively connotated. The characteristics of the two other polarity sub-types, affect and judgment, correspond to the original attitude categories defined in the Appraisal Theory: "affect" includes feelings or emotions, e.g. "happiness" or "ashamed". "Judgment" includes any type of judgment of behaviour or actions, with moral or social dimensions, for example "heroic" or "corruption". The "judgment" labels are used also for words describing a certain character trait of persons or of their behaviour, for example "laziness" or "charitable". The remaining polarity labels in our lexicon, i.e. the labels for modifiers, are the same that can be found in the mentioned related works in section 2.1.1. The complete list of labels, together with a short explanation of their label name and a few examples, is given in Table 3.

In order to assess the applicability of the labels, the four ARGUS project participants separately annotated German polarity words. After the annotation of the first few hundred lemmas, we agreed that further differentiations would be interesting. There are cases where the polarity intensity is different. If we consider two cases of `A_NEG`, *Kopfschmerzen* "headache" and *Kalamität* "calamity" for example, it can be easily agreed that there is a difference in intensity. Therefore, we introduced a specifying tag, "`strong`". Moreover, we became aware that there are cases where the polar situation or condition is intentionally and actively attained, whereas in

Tag	Meaning	Examples
A_POS	Appreciation Positive	optimisation, beautiful, productive
F_POS	Affect Positive	sensitive, happiness, love
J_POS	Judgment Positive	charity, fidelity, charming
A_NEG	Appreciation Negative	illness, unstable, loss
F_NEG	Affect Negative	hatred, mourn, afraid
J_NEG	Judgment Negative	corrupted, dictator, torture
DIM	Diminisher	less, decreasing
INT	Intensifier	more, vast
SHI	Shifter	not, absence of

Table 3.: Complete list of polarity labels with examples.

other instances the opposite is the case. For example, when comparing the `A_NEG` adjectives *verlassen* "abandoned" and *rebellierend* "rebelling", we can argue that we are dealing with two different kinds of negative polarity as well. For this reason, we chose to introduce the further differentiation between `active` and `passive`. The following example shows how this distinction can be useful in French:

(3.3) a. *irritable* "that can be irritated" – A_NEG, passive
b. *irritant* "irritating", "that irritates" – A_NEG, active
c. *irrité* "irritated" – A_NEG, passive

The motivation for this additional polarity information is a more accurate calculation and assessment of the compositional polarity values on phrase level. The `strong` label will enable to scale the polarity's intensity. The differentiation between `active` and `passive` will also help to assess the intensity on the one hand, and on the other hand, we believe to improve the performance of the system in cases such as the following:

(3.4) ***Les otages* (A_NEG passive)** *ont été libérés* (verb_pos).
"The hostages were released."

(3.5) ***Les terroristes* (J_NEG active)** *ont été libérés* (verb_pos).
"The terrorists were released."

Since the overall polarity of the sentence is dependent upon the interaction between the verb (here the a priori positive verb *libérer*) and its dependents, the sentence in (3.4) results in an overall positive polarity, whereas the sentence in (3.5) assumes an overall negative polarity. This shows that a differentiation between negative (or

positive) `active` and negative (or positive) `passive` is potentially useful in cases of polarity conflicts in a bottom-up and verb-centered analysis. The utility of fine-grained polarity labels in the context of polarity conflicts is further analysed in chapter 6.

The polarity labels established for German were found to be language-independent and thus applicable to both English and French. The English lexicon was generated via an automatic translation of the German lexicon and manual correction.[1] In the following section, we describe how we composed the French polarity lexicon.

3.2.2. Composing the French Polarity Lexicon

3.2.2.1. Source

For German, an existing polarity list was provided from a previous UZH project (Clematide and Klenner, 2010). The list was then annotated and revised by the four participants of the ARGUS project. Resources for polarity lexica or simple word lists are scarce for French. Thus, to build up the polarity lexicon for French, we were provided with an automatically translated version of the German polarity lexicon from the UZH *PolArt* project (Klenner et al. (2009a,b), cf. section 2.1.1), which was also used in (Petrakis et al., 2009). The German lexicon entries were automatically translated into French via the platform http://dict.leo.org. The initial version contained a total of 7108 entries pertaining to any word class (nouns, adjectives, adverbs and verbs). The given polarity labels were POS, NEG, DIM, INT or SHI, which were always paired with an intensity value (either 1 or 0.7). The translated polarity lexicon did not contain any entries labelled with the neutral polarity (NEUT), although the original German polarity lexicon did. An illustrative lexicon entry is shown in (3.6).

(3.6) abîmé NEG=0.7 (”damaged”)

3.2.2.2. Corrections

Because of the automatic translation, the initial lexicon contained a high number of mistakes or unusable entries. The neutral polarity entries were lost and transferred into either positive or negative, as the example in (3.7) shows.

(3.7) anonyme NEG=0.7 (”anonymous”)

[1] Nora Hollenstein generated the English lexicon through automatic translation with *Google Translate*.

The lexicon also contained numerous, often non-ambiguous cases, where the polarity was simply reversed. One salient example is shown in (3.8):

(3.8) assassin POS=1 ("murderer")

Other kinds of erroneous entries included typos, mainly missing accents or missing vowels at the end of a word (for example *brav* instead of *brave*), and words that were not translated at all and simply transferred to the French lexicon (e.g. erleuchtet POS=1 ("enlightened")). The most recurrent kind of mistakes was false translations. Those included not only false transfers of polarities such as in (3.8), but also double or multiple entries of the same word such as in (3.9).

(3.9) frapper NEG=0.7 ("to beat, to hit")
frapper NEG=1
frapper POS=1

One of the main sources of mistakes is that the platform http://dict.leo.org always provides multiple possible translations. The translations can be rather loose (yet acceptable) or multi-word expressions. Thus, all possible translations and their single components were transferred to the target lexicon (Petrakis et al., 2009, p. 2), which explains the large number of non-polar words, among which were also modal and auxiliary verbs. However, another source of mistakes was that, among the correct translations, a significant number of inaccurate translations was found, like in (3.10).

(3.10) *recherchable POS=1

The intended German correspondent is the adjective *erstrebenswert* "desirable", but "recherchable" is a simple derivation of the verb *rechercher* "to look for", "to search for" rather than an existing lexical entry.[2] The same is true for the words in (3.11) and (3.12).

(3.11) **hospitalisable* for the German *krankenhausreif* "who needs to be hospitalised"

(3.12) **importable* for the German *untragbar* "unbearable"

Although the derivations are morphologically correct, the semantic meaning and accuracy are questionable, and could moreover not be found in the consulted dictionaries.

Besides these kinds of false entries, many other entries had to be filtered out

[2]We consulted the dictionary *Le nouveau Petit Robert de la langue française* (2008) and the online dictionary available at www.larousse.fr.

from the original lexicon. On the one hand, ambiguous words had to be removed. As mentioned above, this is part of the annotation guidelines that we agreed upon for the ARGUS project: if a word is too ambiguous with regard to its polarity and/or its meaning, it has to be excluded from the lexicon in order to avoid false counts in the later computation of polarity values. A few examples are given below.

(3.13) *louer* – 1. "to hire", "to rent"; 2. "to laud, to praise"

(3.14) *sommeil* – 1. "sleep"; 2. "sleepiness", "state of being sleepy, tired"

On the other hand, the original lexicon contained many entries pertaining to spoken, colloquial language (3.15), as well as to regional dialects (3.16). We judged in particular the dialectical entries to be unlikely to appear in our data (cf. section 3.1) and thus discarded them.

(3.15) adjective *benéf*, short for *bénéfique* "beneficial"

(3.16) noun *bouffiole* "blister, mosquito bite"

The most recurrent types of mistakes in the initial version of the French polarity lexicon are summarised in Table 4 below.

Type	**Example**
False transfer of polarities:	
- *Neutral polarity values transferred into POS or NEG*	anonyme NEG=0.7 ("anonymous")
- *Reversed polarity values*	assassin POS=1 ("murderer")
Typos:	
- *Missing accents*	*desespoir* instead of *désespoir* ("despair")
- *Missing vowels at the end of a word*	*brav* instead of *brave*
Overgenerating translations:	
- *Double or multiple entries of the same word, with different polarities*	frapper NEG=0.7 ("to beat, to hit")
	frapper NEG=1
	frapper POS=1
- *Double or multiple entries of the same non-polar word, due to multiword expressions*	mettre NEG=0.7 ("to put")
	mettre NEG=1
	mettre POS= 0.7
	mettre POS= 0.7
Inaccurate translations:	
- *Non-existing lexical items*	**recherchable*, **hospitalisable*, **importable*
- *No provided French translations*	*erleuchtet POS=1 ("enlightened")

Table 4.: Mistake types in original polarity lexicon due to automatic translation.

We opted for the correction and annotation of an available French translation of the German *PolArt* lexicon (Klenner et al. (2009a,b), cf. section 2.1.1), because we believed that the amount of inaccurate or unusable entries would be by far less significant. However, considering the amount of mistakes and of other types of entries that we discarded (cf. also the total numbers in Table 5, section 3.2.2.5), another lexicon source and approach would have been preferable. As for the English polarity lexicon for the ARGUS project, we could have opted for an automatic translation of the German polarity lexicon that was compiled for the project, with a different translation device or platform, which provides only one-word translations, e.g. *Google Translate*.

3.2.2.3. Annotation and Lexicon Format

Unlike in the German polarity lexicon, a differentiation between word class is necessary, since both nouns and adjectives are lowercased. The specification of the word class is important for cases where nouns and adjectives are homographs with a difference in meaning and/or polarity (3.17), but also in cases of conversions (zero derivation), where a double entry in the lexicon needs to be ensured (3.18).

(3.17) adjective *complice* "of complicity" (A_POS) vs. noun *complice* "accomplice (in a criminal matter)" (J_NEG)

(3.18) adjective *prodige*: *un enfant prodige* "a child prodigy", and noun *prodige*: *un prodige* "a marvel"

For this reason, we manually added the part of speech tags to the entries. We then automatically split the lexicon entries into four different files according to their word class: adjectives, adverbs, nouns and verbs, see [A1][3], and then proceeded to the manual annotation or correction of the polarities. For this, we had to look up many lexicon entries, because they were either unknown to us and/or in order to determine their polarity as correctly as possible. Moreover, we agreed on an uniform lexicon format for the project, which we also added manually. The lexicon format is displayed in example (3.19). The lemma is followed by the polarity tag, the optional `strong` and `active/passive` tags, and the word class (in German). Each slot, whether it is filled or not, is separated by a comma. All lexicon files were automatically verified for correct format [A2].

(3.19) `intelligent,J_POS,,,adjektiv`

[3]Henceforth, each time we refer to a programming script or filename, we provide a cross-reference to appendix A. In the appendix, the scripts and filenames are listed according to the chapters they are referenced in.

3.2.2.4. Addenda

We added many entries during the annotation of the lexicon and during the specification of the verb-polarity frames (section 3.3), as well as during the programming of the CG files (section 4.3). Moreover, we automatically expanded the lexicon by extracting the most frequent lemmas of the respective word classes from our corpus [A3]. A more detailed description of the method is given in section 3.3.2.

A substantial addendum was the adding of ca. 370 verb lemmas (or base forms) to the adjective lexicon, as illustrated in (3.20).

(3.20) corrompu, J_NEG strong
corrompre, J_NEG strong

The base form was added because we found out that the parser, when encountering adjectives derived from verb participles, provides the infinitive form of the verb as the lemma form, such as illustrated in the CONLL[4] parser output sample below:

(3.21) `corrompu corrompre A ADJ g=m|m=part|n=s|s=qual|t=past`

Although this base form is morphologically more accurate than the participle form, most adjectives that are derived from verb participles are usually lexicalised in their participle form. Indeed, this problem does not persist for the nouns: for nouns derived from past participles (e.g. *bien-aimé* "the beloved"), the parser provides the participle form as the lemma form. Thus, we expanded the adjective lexicon with the infinitive verb forms in order to ensure the functionality of further components of the system, in particular the component "polmarker" (section 4.2).

For the expansion of the adjective lexicon, we adopted a morphological approach [A4]. The concerned adjectives can either end in "i" (*moisi* "rotten"), "is" (*conquis* "conquered"),"é" (*éclairé* "enlightened") or "u" (*corrompu* "corrupted"). We made a consciously broad simplification as far as their lemmatisation is concerned. The participle endings "i" and "is" are replaced by the infinitive ending "ir"; the participle ending "é" by "er", and "u" is replaced by "re". This method is limited in many respects, for example when it encounters irregular verbs, as is often the case with participles ending in "u" and "is" (cf. examples (3.22) and (3.23)).

(3.22) entretenu → entretenir (not *entretenre)

(3.23) compromis → compromettre (not *compromir)

[4]CONLL refers to an unified format for parser output. See for example the definition by Joakim Nivre at http://nextens.uvt.nl/depparse-wiki/DataFormat (last accessed 27th October 2013).

Moreover, besides wrong lemmatisations, false positives are created. The adjectives *flou* "blurred" or *courtois* "courteous" for example, happen to end in "u" and "is", but do not derive from past participles.

The employed technique works better and almost without errors for the adjectives in "é", except for adjectives starting with the suffix "in", such as *inapproprié*, *inadapté*, *inalteré*, etc. Those are derivations from participles themselves, not from verbs (e.g. the verb *inaltérer* does not exist). However, the method enabled to find all potential cases. After manual corrections, we added 368 of the 445 filtered adjectives in their base form to the adjective polarity lexicon.

3.2.2.5. Total Numbers

The total number of lexicon entries and their classifications are displayed in Table 5. The verbs are excluded from the table since they are being processed in a different way in the system. The table illustrates the distributions among the coarse-grained polarity tags. No distinction between the fine-grained negative and positive polarities (cf. section 3.2.1) is made. As already mentioned, the lexicon does not contain any entries marked with the neutral polarity. This is also the case for the English polarity lexicon that was compiled for the ARGUS project. In fact, the system does not necessarily require neutral polarity entries: per default, unmarked words are considered as neutral. The French lexicon, moreover, does not feature a multi-word list such as the German lexicon. A multi-word list contains lexicalised polar expressions, as for example *schwarz arbeiten* "to work illegaly" (lit. "to work black"). A last important point with regard to the French polarity lexicon is that, unlike for the German lexicon, only one annotator was involved. Hence, no inter-annotator agreement could be assessed as it was the case for the German annotation. As a result, erroneous or ambiguous entries might still persist in the French lexicon.

	Negative	**Positive**	**Diminishers**	**Intensifiers**	**Shifters**	**Total Lemmas**
Adjectives	1550	858	3	34	5	**2450**
Adverbs	35	45	1	35	16	**132**
Nouns	1332	508	1	10	5	**1856**
ALL	2917	1411	5	79	26	**4438**

Table 5.: Properties of the French polarity lexicon.

3.3. Verb-polarity Frame Resource

We compiled a verb-polarity frame resource for the project. The purpose of a resource with fully specified verb-polarity frames is an easier and faster modelling of the verbs in VISL (see section 4.3.3). Moreover, the resource can serve as a basis for further developments of the system. We automatically extracted the most frequent polar verbs from our corpus (section 3.3.2) and manually selected 322 polar verbs, which we specified in individual sheets. Guidelines were discussed in several project meetings but no fixed agreement was made on many points concerning the specification format, since the focus was to provide coherence and compatibility with the parser's output. A detailed description of the specification procedure is given in section 3.3.3. A precise account of the term and concept "verb-polarity frame" is given in section 3.3.1 below.

3.3.1. The Concept Verb-polarity Frame

By "verb-polarity frame specification" we mean the description of a verb's behaviour (or verb-frame) in terms of:

a. The verb's subcategorisation frame, i.e. what kind of syntactic constructions the verb can take (reflexive, transitive, etc.);

b. The expectations, effects and/or propagations of polarity, i.e. how the verb imposes polarity values on its arguments, given a certain syntactic construction.

In other words, the term "verb-polarity frame" encompasses the traditional syntactic frame (a) and, most importantly, the *polarity frame* of a given polar verb (b): in our verb-centered approach on sentiment analysis, we posit that each polar verb has at least one polarity *effect* or one polarity *expectation* on at least one of its arguments, depending on its subcategorisation frame. Whether a verb triggers a polarity effect or expectation on one of its arguments depends, however, not only on its syntactic frame, but also on its semantics: a polarity **effect** is a consequence of polar nature that results from an argument's involvement in a given event that is signified by the verb. A polarity **expectation** is, as opposed to the polarity effect, a presupposition or anticipation with regard to an argument's polarity, given a verb's semantics. Consider for example sentence (3.24) below:

(3.24) The bank failed to inform their costumers.

The verb "to fail" is a priori negative. The verb's arguments, i.e. the subject "bank" and the object "to inform their costumers" do not carry any prior polarity. However, in composition with the verb "to fail", they do gather a polar value: the subject receives a negative effect due to its role in the negatively connotated act of failing. As far as the direct object is concerned, especially if it has a neutral polarity (as in sentence (3.24) above), we assume that it is perceived as something valuable or desirable, since its failed achievement or fulfillment is connotated as regrettable by the verb's semantics. Thus, the described object or situation is expected to be positive, since it is contextualised as.[5]

The concept of polarity effects and expectations does not have correspondents in related work to the best of our knowledge (see section 2.1.2.2). On the one hand, the principal aim of polarity effects is to develop a model where a target-specific analysis is possible. In the above sentence (3.24), although the sentence's overall polarity is negative, only the subject receives a negative effect. In the context of a target-specific sentiment analysis, where "the bank" would be the target for example, a negative value would be added its statistics (or counts) after the sentiment analysis of sentence (3.24). The intended utility of polarity expectations is, on the other hand, rooted in the context of polarity conflicts. In the case of the positive verb "admire" for example (cf. sentence (3.25) below), no effect is produced on the subject, but the object gathers a positive expectation.

(3.25) He admires his country's political leader.

If, however, this expectation is not fulfilled, i.e. if the object has a negative polarity, we are likely to deal with a polarity conflict, as we are going to analyse and discuss in section 6.2. Of course, the distinction between polarity effect and expectation has heuristic reasons and, depending on certain verbs' semantics, it can become less clear-cut. Considering sentence (3.25) above for example, one could argue that the verb not only triggers a positive polarity expectation on its object, but also that the object is to receive a positive effect, because of its being admired. However, for implementational reasons (cf. section 4.3.3), we assume that an argument can only receive one effect or one expectation at a time. This is also the reason why in the context of polarity conflicts, we will also consider verbs that predict a polarity effect on their object (see section 6.2).

Thus, the purpose of the individual verb-polarity frame specification is to cover (predict) as many constructions as possible where a given polar verb can trigger,

[5] As a consequence, a direct object with a negative polarity (e.g. "to betray their costumers") would yield a polarity conflict, because the expectation has been violated. The subject effect would remain unchanged, but the overall sentence-polarity would change due to the violated polarity expectation.

including the expected default and non-default scenarios in terms of polarity effects and expectations. An example with the French verb *battre* "to beat", "to fight" will illustrate the purpose and utility of verb-polarity frame specification. The following types of frames were considered and specified:

(3.26) *Il se bat contre le cancer depuis longtemps.*
"He has been fighting against cancer for a long time."

(3.27) *Elle se bat pour les droits des femmes.*
"She fights for women's rights."

(3.28) *Il bat sa femme.*
"He beats his wife."

The verb *battre* can be used as a reflexive verb with attached prepositional phrases, introduced by either *pour* "for" or *contre* "against". It also functions as a non-reflexive transitive verb. As far as the verb's polarity is concerned, we consider it to trigger as a negative verb in (3.26) and (3.28), but as a positive verb in (3.27).

The examples above prove the utility of frame specification because many verbs can trigger both polarities, depending on their syntactic constructions. On a fine-grained compositional level, the syntactic frame also entails different expectations or effects on the different syntactic components. In (3.28), the subject gets the effect "j_neg active" (`J_NEG aktiv`) and the direct object gathers a simple negative effect. However, lexical disambiguations need to be taken into account. In the sports domain for example, we are likely to find sentences such as *L'athlète bat son record* "The athlete beats his record". Here, the verb polarity and the effects on the subject and direct object need to be reversed. Similarly, the default case in (3.27) can be violated, as for example in (3.29):

(3.29) *Ils se battent pour une dette.* "They fight because of a debt."

The specification of verb-polarity frames allows us to take most scenarios and their exceptions into account. The exceptions can be included in the VISL scripts in a second stage (see section 4.3.3.5) and increase the accuracy of the sentiment analysis.

3.3.2. Extraction of Verbs

In order to compile a meaningful list of polar verbs, we generated a list of the most frequent verbs from the parsed version of our corpus, and compared it with the verb list [A5] that we extracted from the original polarity lexicon (cf. 3.2.2). In the following, we describe the method that we employed [A3].

The method can be divided into two steps. First, the most frequent verb lemmas in the two parsed data sets are counted and summed. Thus, a list of the 1000 most frequent verb lemmas in the corpus is generated and automatically written to an output file [A6]. This high number is necessary because of parser-errors. The lemma *étayer* for example, is a wrong lemmatisation of *étaient* (lemma *être* "to be") and listed as the 34th most frequent verb. The form *ait* is at position 150; however, it is not a lemma but the subjunctive form of the verb *avoir* "to have". Moreover, we correctly expected non-polar verbs, such as auxiliaries, modals, verbs of movement or speech, etc., to be very frequent (see Figure 2).

Figure 2.: 15 most frequent verb lemmas extracted from the corpus.

The second step consists of the comparison of the list of the most frequent verbs in the corpus with the verb list from the polarity lexicon. If a lemma of the top 1000 list is found in the lexicon verb list, it is automatically written into another output file [A7] with a corresponding notice. If the lemma is not found, it is written into the output file as well, with another related remark. Indeed, we considered both cases, matching and non-matching entries, to be relevant. The first case provides the information that the concerned verb is already in the original list of polar verbs, and in the case of a non-match, there are two possibilities:

a. either the lemma is absent from the list because it has no polarity, or

b. it is polar and needs to be added.

With this method, the lexicon verb list was increased by ca. 100 verbs[6], and we also retrieved the information about the most frequent polar verbs in our corpus.

[6]We used this method to expand the 3 other lexicon files as well.

3.3.3. Specification Procedure

In total, we selected and specified the 322 most frequent polar verbs (see previous section 3.3.2) in separate sheets. We then eventually modelled them in VISL for the sentiment analysis (section 4.3.3). We limited the selection to 322 polar verbs because their frequency rapidly decreases in the corpus. While the most frequent polar verb, *permettre* "to allow", "to enable", occurs only 1402 times in the corpus (rank 22), the verb *aimer* "to love", number 11 among the polar verbs, occurs 616 times (rank 59) in the corpus. Polar verb number 300, *détester* "to hate", is verb number 758 in the top 1000 list with only 35 occurrences.

The specification sheets have an uniform layout, which consists of 8 columns. Their content is listed along with explanations and examples in Table 6. These categories all originate from the discussions and hypotheses formulated during the meetings in the context of the ARGUS project. The categories marked in italics were not consequently filled out, because their role and function were not fully discussed and established, and therefore not integrated in the system. The tags in the specification of the verb-polarity frame, henceforth "verb-frame tags", are language and parser dependent, although a few overlaps exist with the German specification. The tags that we used in the specification of French verbs are displayed and described in Table 7.

Rahmen	Syntactic frame, e.g. "subject + direct object"
Verbpolarität	Verb polarity, e.g. "POS"
Relation	Relation between the different components or agents of the frame, e.g. "subj->objdir=neg" (the relation between subject and object is negative)
Erwartungen und Effekte	Expectations and effects, e.g. "suj=neg(eff), objdir=pos(exp)" (the subject gathers a negative effect, the object is expected to be positive)
Quellen und Ziele	Source and target, e.g. "verbQ, sujZ" (the verb is the source of polarity, the subject the target)
Propagierung	The propagation of the polarity, especially if it does not fulfill the expectations of the default case. E.g. "objdir=>verb, suj": the polarity of the direct object has repercussions on the verb's and on the subject's polarity.
Beispielsatz	Example sentence
Bemerkungen	Remarks, comments

Table 6.: Categories for Verb-polarity frame specification.

As can be seen in the rightmost column of Table 7, the main parser tags do not

Verb-frame Tag	**Meaning**	**Parser Tag**
suj	subject	suj
objdir	direct object as in "Elle lit **le journal**."	obj
objdir[PRF]	reflexive object (pronoun) as in "Il **se** réjouit."	aff or mod
de_obj	an object (noun phrase) introduced by *de* such as in "Il se réjouit **des vacances**."	de_obj
sub_que	subordinate clause introduced by *que*, such as in "Il annonce **que tout s'est bien passé**."	obj
sub_de	subordinate clause introduced by *de*, such as in "Il accepte **de changer de travail**."	obj
objp[*prep*]	prepositional phrase; the preposition is specified in the square brackets. If more than one preposition can occur, no specification is given.	mod, a_obj, p_obj
mod	any type of modifier, mostly adverbs, e.g. "Il a été viré **hier**."	mod
*	indicates optionality, e.g. objp*	

Table 7.: Tags for the French verb-polarity frame specification.

differ much from each other. Therefore, we chose to coin our own verb-frame tags for better comprehension and overview. The restrictive parser tag-set is also the reason why the verb-frame tag `mod` can also denote prepositional phrases most of the time: they are tagged by the parser in the same way than other types of phrases. The verb-frame tag `mod` is however used restrictively and when it is expected not to influence or not to be influenced by the verb polarity. For example, one of the frames for the verb *perdre* "lose" is `suj, objdir[PRF], mod*`. This can stand for the following constructions (among others):

(3.30) *Il se perd.* "He gets lost."

(3.31) *Il se perd souvent.* "He gets lost often."

(3.32) *Elle s'est perdue dans la foule.* "She got lost in the crowd."

The frame entails a negative polarity for the verb (`verb_NEG`) and a negative effect on the subject (`subj=neg(eff)`). Whether the slot `mod` is filled or not, and whether it is filled with an adverb or a prepositional phrase, does not change the predicted verb polarity or the effect on the subject.

Similarly, most verbs that contain the simple frame specification `suj` or `suj, mod*` imply that the passive form can also be modelled in the constraint grammar. Compare the two sentences in (3.33):

(3.33) a. *Il a déçu.* "He has disappointed." [active mode, past tense]
b. *Il est déçu.* "He is being disappointed." [passive mode, present tense]

The auxiliary tags differ according to their use (`aux_tps` stands for tense auxiliary and `aux_pass` for passive auxiliary), but the dependency structure does not differ. As long as the presence of auxiliary verbs is allowed and modelled (see section 4.3.3.5), composed past forms (such as the tense *passé composé*), as well as passive forms of the concerned verb will trigger and do not need further differentiation. This is why these syntactic frames are not explicitly specified in the verb-polarity frames.

The verb-polarity frame specifications serve as guide-lines for the implementation of the verb-component in the system (see section 4.3.3). Moreover, and as the polarity lexicon (section 3.2), they constitute an expandable resource of the project.

4. The French Sentiment Analysis System

The French sentiment analysis system consists of a processing pipeline which can be divided into three main steps. The first step consists of the linguistic pre-processing of the data: the input text is parsed for dependency structures and then converted into VISL-readable format (4.1). In the second step, the words are marked with their prior polarity on the basis of the polarity lexicon (4.2). The third step consists of the actual sentiment analysis (4.3), where the polarity values are assessed in a bottom-up sequence, from word to clause level.

4.1. Dependency Parsing and Conversion into CG Format

As discussed in chapter 2, the approach that suits best our model is the principle of compositionality, which is described and modelled most efficiently on the basis of dependency grammar (2.2.1). Hence, a linguistic pre-processing of the input texts, where the sentences are analysed for their dependency structure, is required.

The dependency parser in the French system is an adapted version of the statistical Berkeley Parser (University of California, Berkeley) for French.[1] It is called "Bonsai Parser" and is a product of the cooperation between the University Paris Diderot and the *Institut national de recherche en informatique et en automatique* (INRIA). The Bonsai Parser relies on an extended tagset of the French Treebank[2]. Unfortunately, the parser proves to be very slow, especially with large amounts of data. Moreover, the parsing is often inaccurate, especially when it deals with long sentences and/or unknown words. Recurring cases of parsing mistakes or inaccuracies are described in section 5.3.1 and 5.3.2.

[1] http://alpage.inria.fr/statgram/frdep/fr_stat_dep_bky.html, last accessed on 12th October 2013.

[2] http://alpage.inria.fr/statgram/frdep/Publications/FTB-GuideDepSurface.pdf, last accessed on 12th October 2013.

```
1     Le      le      D     DET    g=m|n=s|s=def   2     det    _     _
2     maire   maire   N     NC     g=m|n=s|s=c     3     suj    _     _
3     craint  craindre      V      V       m=ind|n=s|p=3|t=pst    0     root   _     _
4     pour    pour    P     P      _       3     mod    _     _
5     l'      le      D     DET    n=s|s=def       6     det    _     _
6     image   image   N     NC     g=f|n=s|s=c     4     obj    _     _
7     de      de      P     P      _       6     dep    _     _
8     sa      son     D     DET    g=f|n=s|s=poss  9     det    _     _
9     ville   ville   N     NC     g=f|n=s|s=c     7     obj    _     _
10    .       .       PONCT PONCT  s=s     3     ponct  _     _
```

Figure 3.: The Bonsai Parser output in CONLL format for the sentence *Le maire craint pour l'image de sa ville.* ("The major fears for his city's image.").

```
"<Le>"
 "Le" le C_DET DET DEP_det #1->2
"<maire>"
 "maire" maire C_NC NC DEP_suj #2->3
"<craint>"
 "craint" craindre C_V V DEP_root #3->0
"<pour>"
 "pour" pour C_P P DEP_mod #4->3
"<l'>"
 "l'" le C_DET DET DEP_det #5->6
"<image>"
 "image" image C_NC NC DEP_obj #6->4
"<de>"
 "de" de C_P P DEP_dep #7->6
"<sa>"
 "sa" son C_DET DET DEP_det #8->9
"<ville>"
 "ville" ville C_NC NC DEP_obj #9->7
"<.>"
 "." . C_PONCT PONCT DEP_ponct #10->3
```

Figure 4.: VISL format for the sentence *Le maire craint pour l'image de sa ville.* after the conversion.

Once the input text has been parsed for dependency structures, the parser's CONLL output is automatically converted into a VISL-readable CG format (cf. Figure 4) [A8]. As illustrated in Figure 3, in the CONLL output, each token of a sentence is written in one line, along with 10 tab-separated "fields", which contain grammatical and syntactical information about the token. Of these 10 fields, 7 are kept for the output in VISL format:

1. Token
2. Lemma
3. General part-of-speech tag

4. Fine-grained part-of-speech tag
5. Dependency label
6. Word-ID
7. Head

The VISL format is constituted of cohorts (see Figure 4 and section 2.2.2 for details). It is eventually used for the prior polarity marking (section 4.2) and the actual sentiment analysis (section 4.3).

4.2. Prior Polarity Marker

The second step in the pipeline of the sentiment analysis system consists of marking the words in the newly parsed and converted text (see previous section 4.1) with their prior polarity. The process of the prior polarity marker, or "polmarker" [A9] can be divided into two steps: first, the polarity lexicon is converted into VISL-readable lemma lists (section 4.2.1). Second, tagging rules are established, which include the previously generated lemma lists; the words in the input text are marked with polarity tags (cf. section 3.2.1), according to their word class (section 4.2.2).[3]

4.2.1. Building VISL-readable Lemma Lists

For the first step of the prior polarity marker, we copied and concatenated the adjective, adverb and noun lists of the polarity lexicon (cf. section 3.2) into one single text file [A10]. This new polarity lexicon file serves as input for the conversion into VISL-readable lemma lists [A11]. The conversion consist of the iteration through the polarity lexicon file entry by entry: the lemmas are progressively inserted into lists, according to the following criteria:

1. The polarity
2. The tag `strong`
3. The tag `active` or `passive`
4. The word class

[3]The polmarker scripts were all written by Michael Amsler. We made a number of language-specific adaptations, which we will mention in the following.

If in the polarity lexicon an adjective was annotated with the tags `A_NEG strong` for example, it is put into the list of `A_NEG` adjective lemmas as well as in the list of `strong` adjective lemmas. The criterion of the word class is an adaptation that we made for the French polmarker. The original method was developed for the conversion of the German polarity lexicon. The word class was not taken into account, since the distinction between nouns and adjectives relies on the distinction between upper- and lowercase entries. In French, nouns as well as adjectives are lowercased (cf. examples (3.17) and (3.18) in section 3.2.2.3). Hence, if the word class is not specified, homographs are going to be considered as redundant entries and only the second occurrence is going to be kept. To avoid this, a "redundancy tag" is added whenever a lemma is encountered for the second time. This tag is removed after the lemma has been put into its respective list.[4]

Once the lists are filled with all the lexicon entries, they are converted in VISL-readable syntax, henceforth "VISL lemma lists" [A12]. The VISL lemma lists now contain all the lexicon entries and will be included in the next step of the polmarker, which will carry out the eventual prior polarity tagging (see following section).

4.2.2. Generating the Prior Polarity Marker

In the next step of the prior polarity marking process [A13], rules for the actual polarity tagging are written in VISL syntax [A14]. These rules will carry out the eventual prior polarity marking. The previously compiled VISL lemma lists [A12] are automatically included and accessed for the marking.

The polarity tagging rules consist of simple `ADD`-rules (cf. section 2.2.2), where the target is defined in terms of word class and lemma list(s).[5] If a noun in the VISL input text (cf. section 4.1) is contained in the list `a_positive_lemma_noun` and `strong_lemma_noun` for example, the tags `A_POS` and `POL_strong` are added to its cohort, as illustrated below.

```
"<perfection>"
  "perfection" perfection C_NC NC DEP_root A_POS POL_strong
```

In summary, the prior polarity marking process in the sentiment analysis system is executed in one step ([A9]) on an input text in VISL format: the polarity lexicon is converted in VISL readable lemma lists, which are referenced in the prior polarity marking rules. The rules eventually carry out the prior polarity marking on the input text.

[4]The same adaptation was also made for the English polmarker.

[5]The word class as part of the target definition is also an adaptation for the French polmarker.

4.3. Sentiment Analysis

Once words have been marked with their prior polarities, the actual sentiment analysis can be carried out. Since the sentiment analysis is based on the principle of compositionality and on dependency grammar, the order of the processing steps, a bottom-up sequence (cf. Figure 5), is crucial. First, the shifters and their effects on their heads or dependents are modelled (section 4.3.1). The next step consists of the assessment of NP, PP, and coordination polarities (section 4.3.2).[6]

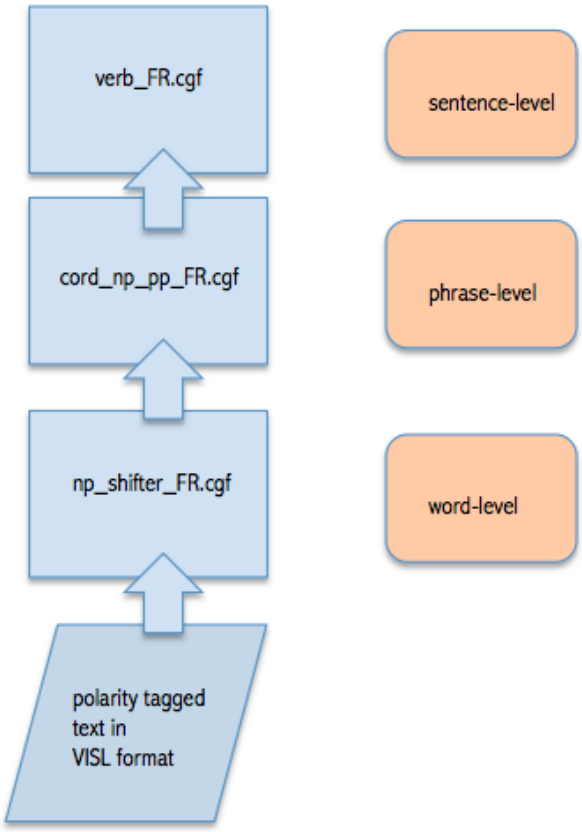

Figure 5.: Bottom-up Procedure of the Sentiment Analysis System.

Example (4.1) and (4.2) below show how this bottom-up sequence is important.

(4.1) *Il travaille avec enthousiasme.* "He works with enthusiasm."

(4.2) *Il travaille sans enthousiasme.* "He works without enthusiasm."

The prior polarity of the noun *enthousiasme* is `A_POS`. The presence of the shifter

[6]Henceforth, we will use the terms "noun phrase" (NP) and "prepositional phrase" (PP). We are aware that the terms are inappropriate in the context of DG; we will, however, use them for simplicity's sake whenever we intend a nominal or prepositional head and its dependents.

sans in (4.2) reverses the polarity to `A_NEG`. Thus, thanks to the pre-processing of the polarity on the word-level, the chunks (or phrases) are assessed correctly as a positive PP in (4.1) and as a negative PP in (4.2). On this basis, the next and final layer, the verb layer, is then analysed (section 4.3.3).[7]

4.3.1. NP Shifter

In the first layer of compositionality [A15], the shifters are first defined in a list. This list includes any word that has been previously marked with the tag `SHI` (cf. Table 3, section 3.2.1) by the prior polarity marker, as well as a number of prepositions. We had to add the prepositions manually, since they are not included in the polarity lexicon (cf. section 3.2). The same applies for combinations of adverbs, or combinations of adverbs and prepositions, which are lemmatised with an underscore, as for example `pas_encore` "not yet" or `loin_de` "far from".

The rules in the NP shifter layer are mainly substitution rules. They consist of replacing the polarity tags with their opposite polarity, along with the additional tag `SHIFTED`. The substitution rules can be divided into three categories or steps:

1. Direct dependency relations (section 4.3.1.1)
2. Indirect dependency relations (section 4.3.1.2)
3. Special cases and a priori disambiguations (section 4.3.1.3)

We adapted the first step from the original German version with almost no changes, except for the parser dependency tags. Step two and three are language specific and do not derive from the original German rules.

Once all the substitution rules have been applied, the shifters are "fired": `ADD`-rules mark the shifter with the tag `FIRED` if its parent or child is a shifted polar word. The purpose of marking shifters once they have triggered, is to avoid redundancy on the one hand, and for possible evaluation purposes on the other hand.

4.3.1.1. Direct Dependency Relations

Direct dependency relations are cases where the shifter is either the parent or the child of the polar word (e.g. *pas beau* "not beautiful"). In this case, the polar tag is substituted with its opposite tag and with the tag `SHIFTED`. The substitution

[7] All steps are implemented in VISL constraint grammar scripts, which are an adaptation of the German prototype. Michael Amsler developed and implemented the bottom-up structure for the sentiment analysis in VISL for the German system; he provided us with the initial versions of his VISL scripts, which we adapted or rewrote for the French system.

is done if the parent or the child of the polar word is a shifter, and if it has not yet been shifted[8]. The aim of this additional condition is, again, to avoid multiple triggering. For cases where a shifter is followed by a noun phrase composed of two polar words, e.g. *loin de la beauté* (`A_POS`) *parfaite* (`A_POS`) "far from the perfect beauty", additional rules ensure that both components are shifted and marked as such.

4.3.1.2. Indirect Dependency Relations

It can happen that polar words do not have a direct dependency relation with a shifter, for example in prepositional phrases introduced by *de*, or in cases where a shifter is followed by another adverb before the polar word:

(4.3) *pas si cher* "not so expensive"

(4.4) *pas assez intelligent* "not intelligent enough"

(4.5) *pas très beau* "not very beautiful"

Here, more complex conditions, with lexical specifications, are necessary. For the case in (4.3) for example, we elaborated a rule with the condition that the polar word has as a child the adverb *si*, which is preceded by a shifter. Similarly, for a negated prepositional phrase such as *pas d'ambitions* "no ambitions", we added a rule where we specify that the parent of the polar word is the preposition *de*, and that the preposition is preceded by a shifter.

4.3.1.3. Special Cases and a priori Disambiguations

We needed to model a few specific cases separately. Certain combinations of polar words and modifiers function differently from what their prior polarities prompt or from what the shifter rules expect. We identified 3 cases:

a. The intensifier *trop* "too" + any positive polar word does not function as an intensifier, but as a shifter. In *trop enthousiaste* "too enthousiastic" for example, we claim that the positive polarity of the adjective *enthousiaste* needs to be shifted to `A_NEG`.

b. The expression *pas mal de* + noun, is composed of the shifter *pas* "not" and the negative adverb *mal* "bad", but corresponds to the English "quite a lot of" or "quite many", e.g. *pas mal de difficultés* "quite a lot of difficulties". In this

[8]This is verified via the presence or absence of the tag `SHIFTED` in the cohort.

case, the polar word *mal* is marked as an intensifier with the tag `EXCEPTION: INT` and the shifter tag is replaced with the tag `NO_SHI`.

c. Certain nouns that are coded as negative in the lexicon function as shifter when they modify another noun: e.g. *absence* "absence" vs. *absence d'optimisme* "lack of optimism". These cases are modelled with rules containing lexical specifications.

For **a.** and **b.**, an a priori disambiguation is necessary. The concerned rules are applied first, since all the rules follow a cascading principle. The additional rules for the cases described in **c.**, however, are added at the bottom.

4.3.2. Coordination and NP-PP Composition

The second layer that is concerned with phrase level analysis [A16] consists of two types of analysis. First, coordinations are analysed for polarities: they include polar or partly polar coordinations, such as "an ugly, horrible coat" or "a blue and horrible coat", and bi-polar coordinations, such as "a beautiful, but dangerous place" (section 4.3.2.1). Coordination structures are considered independently from simple NPs, i.e. nominal heads and their direct dependents (articles and/or adjectives). This is why the assessment and marking of NP and PP heads and their respective children (section 4.3.2.2) is carried out after the one of coordinations .

The rules for the coordinations and for the NP-PP composition are the ones that we had to expand and adapt the least, since their basis is a more syntactic one and thus less lexical and language-specific.

4.3.2.1. Coordination

For the coordination, so-called coordination structures are defined first. In other words, three possible syntactic structures of coordinations are defined, via dependency relations and dependency tags. They include coordinations with both contrasting and non-contrasting conjunctions (e.g. *et* "and" or *mais* "but"), as well as coordinations without lexical conjunctions (e.g. "a beautiful, new coat")[9]. In the Bonsai parser output (cf. section 4.1), the head of a coordination is always the first head noun in the coordination structure; the other coordination elements, i.e. the conjunctions, adjectives or possible other nouns, are all the dependents of the first head noun. The defined dependency structures serve as an instantiation for the marking rules.

[9]In this case for example, the comma is marked as conjunction by the parser.

The heads and children of coordinations are marked via ADD-rules. For example, if the cohort of a word contains a positive tag and it is moreover identified as the head of a coordination structure (as defined previously), the tag COORD_HEAD_POS is added to its cohort. If only the syntactic structure corresponds and no polarity is found, the head is marked as neutral (COORD_HEAD_NEUT). SUBSTITUTE-rules analyse the polarity of the marked heads' children, and determine whether a head keeps its plain polarity or whether it gets the tag COORD_HEAD_BIPOLAR. This happens if a head is marked as a positive coordination head for example, but has one or more negative coordination children. Similarly, if a head is marked as neutral, its tag is substituted by COORD_HEAD_BIPOLAR if it has both negative and positive children. If it has only negative or only positive children, its neutral coordination tag is changed to COORD_HEAD_NEG or to COORD_HEAD_POS respectively.

A slightly more complex case of bipolarity can be encountered in concessive phrases, such as "a beautiful, but dangerous place". From a compositional perspective and for heuristic reasons, we decided that the polarity of the second (concessive) element, in this case the negative polarity, dominates. Thus, the previously attributed substitution, COORD_HEAD_BIPOLAR, is again replaced by the dominating polarity, COORD_HEAD_NEG. In the case of the example "a beautiful, but dangerous place" *un endroit beau, mais dangereux*, this analysis is ensured by a rule where the condition is the presence of the conjunction *mais*, followed by a negative polar word. Once the heads' polarity is determined, the children's polarity is also marked via simple dependency relation conditions. Additional restrictions as to the children's word class (noun or adjective) and dependency relation (e.g. optional presence of a coordination tag) are added in order to ensure more precise marking.

4.3.2.2. NP and PP Composition

The polarity of noun phrases is assessed after the analysis of coordinations. The approach is similar to the marking of the coordination heads. A head is marked according to its own polarity and whether it has been marked as a NP head yet or not.[10] Also, it cannot be part of a bi-polar coordination. Non-polar (or neutral) heads are marked according to their child's polarity: e.g. in the case of the noun phrase *une belle journée* "a beautiful day", the non-polar head *journée* is marked as POS-NP-POLCHUNK-HEAD because of its positive child *belle*.

In cases where NP components have opposing polarities, we decided to set the negative polarity as dominant in any type of opposing polarities combination. In

[10]The cohort of a negative noun (for example) can contain both a head-of-coordination tag (COORD_HEAD_NEG), as well as a head-of-noun-phrase tag (NEG-NP-POLCHUNK-HEAD).

other words, whether a positive noun is modified by a negative adjective (4.6), or a negative noun is modified by a positive adjective (4.7), the NP polarity is always computed as negative.

(4.6) *un sourire hypocrite* "an hypocritical smile"

(4.7) *défaite triomphante* "a triumphant defeat"

Empirical analyses in section 6.1 will assess whether this generalisation is productive. The NP children are, similarly to the children in the coordination structures, also marked (e.g. with the tag `POS-NP-POLCHUNK-CHILD`), again via dependency conditions in `ADD`-rules.

On the basis of the marked noun phrases, prepositional phrases are also tagged. Their heads, the prepositions, are set as target in the `ADD`-rules; if the preposition has as a direct child a marked NP, the prepositional phrase is going to be marked correspondingly. The PP *dans la bonne humeur* "in a good mood" for example, contains the positive NP *bonne humeur*. Since the positive NP head is the direct child of the preposition *dans*, the preposition is going to receive the tag `POS-PP-POLCHUNK-HEAD`. The PP children are then again marked on the basis of dependency relations. At this stage, the analysis and marking on the phrase level is completed.

4.3.3. Verb Analysis

The verb analysis [A17] is the largest and most important component of the sentiment analysis process. It consists of the assessment of polar verbs and their polarity effects and expectations on their arguments (cf. section 3.3.1). The verb analysis is the innovative part of the project, hence its importance. It is moreover parser-specific (because of the provided dependency structures and tags) and, above all, language-specific. This is why the French verb-component strongly differs from the German one, except for the general architecture of the method: polar verbs are attributed to so-called frame-classes (section 4.3.3.1), according to their previously specified verb-polarity frames (cf. section 3.3). For each frame-class, corresponding syntactic frames need to be defined (section 4.3.3.2). On this basis, i.e. on the basis of the defined frame-classes and syntactic frames, rules which carry out the marking of verbs and their arguments are defined (section 4.3.3.3). Copula verbs, verbs with no a priori polarity, and so-called diminisher and intensifier verbs need a different treatment, which is described in section 4.3.3.4. Last but not least, we modelled a few exceptions or special cases, e.g. cases where the verb's polarity is reversed in case of an "unexpected" object polarity (section 4.3.3.5). The modelled verb-polarity frames are all based on the specifications described in section 3.3.

4.3.3.1. Verb Class Definitions: Frame-classes

Based on the verb-polarity frame specifications (3.3), 322 polar verbs, which correspond to the same selection described in section 3.3.3, are grouped into so-called frame-classes. The frame-classes consist of lists of verbs, whose names carry information about the verbs' subcategorisation and polarity frame. The frame-class `fclass_subj_neg_objdir_eff_verb_neg` for instance, denotes the syntactic frame `subject + direct object`. The verb is negative (`verb_neg`) and the direct object gathers a negative effect (`neg_objdir_eff`). Verbs belonging to this frame-class list are, for example, the verb *critiquer* "to criticise" or *dénoncer* "to denounce". For technical reasons, a syntactic component can only be marked with one effect or expectation at a time: a direct object for example, cannot have both a positive expectation and a positive effect, but only one of both.

In total, we defined 72 frame-classes. An encompassing list, `verb_frames`, contains all the corresponding frame-class tags, as well as the verb-polarity tags `verb_POS` and `verb_NEG`.

4.3.3.2. Syntactic Frames

Each frame-class posits a specific syntactic structure. In other words, a polar verb triggers a verb-polarity and certain polarity effects or expectations given a specific syntactic structure. Thus, all required syntactic structures need to be defined in order to eventually establish the verb-marking rules (cf. the following section, 4.3.3.3).

In total, we defined 18 syntactic frames for the French verb analysis. This turned out to be a difficult task because of the parser's restricted tagset. The parser tag `mod` (`DEP_mod` in VISL format) for instance, is used for adverbs, personal or relative pronouns, as well as for prepositions. Hence, many lexical specifications are needed. If the syntactic frame is aimed at matching sentences with subordinate clauses introduced by *de* or *que* for example, an additional condition with a lexical specification must be included in the definition of the concerned syntactic frame. Other potential, but unwanted child nodes need to be negated (excluded) in order to avoid unwanted triggering of the rules. However, the opposite problem, namely to set syntactic frames that are too restricted, also exists.

4.3.3.3. Verb, Effect and Expectation Marking

The verbs are marked via ADD-rules, which set the following conditions:

a. No marking on the verb has been done yet

b. A specific syntactic frame

c. The verb is contained in the specified frame-class list

If all the conditions are fulfilled, two tags are added to the verb's cohort: on the one hand, a frame-class tag, which corresponds to a frame-class tag listed in verb_frames (cf. section 4.3.3.1), and on the other hand, a verb-polarity tag, i.e. verb_POS or verb_NEG. The frame-class and verb-polarity tags are separated for evaluation purposes on the one hand (i.e. for the assessment of the verb's polarity, regardless of its frame), and for the marking of the effects and expectations on the other hand. In this case, the frame-class tag is considered.

The polarity effects and expectations that are generated by the verbs are in fact marked with a similar method to the marking of verbs, namely with ADD-rules. The target of the rules are, depending on which type of frame-class is concerned, the subject, the object(s), the clause, or the prepositional phrase of the sentence. For some types of targets, e.g. direct object, several tags or tag-sets need to be defined in the rules. The main condition of the ADD-rules is the presence of the concerned frame-class tag in the cohort of the parent node. If a noun is a direct object in a sentence and its parent node is a verb marked with the frame-class tag FCLASS_subj_neg_objdir_eff for example, it is marked with the tag EFF_NEG-POLCHUNK-HEAD. The children of the so-called polarity chunk heads are consequently marked as well (e.g. with EFF_NEG-POLCHUNK-CHILD), by means of dependency relations .

4.3.3.4. Copula, NoAPriori, Intensifier and Diminisher Verbs

A number of verbs or verb constructions need a different processing in terms of verb, effect and expectation marking. There are two types of verbs whose polarity is dependent on their object's or modifier's polarity: on the one hand, copula verbs, and on the other hand, verbs that belong to the frame-class that we labelled NoAPrioriPol. In the case of copula verbs, the object or modifier in the sentence determines the verb's polarity and the effect on the subject. In sentence (4.8) or (4.9) for example, the verb is marked as negative and the subject receives a negative effect.

(4.8) *Il est triste.* "He is sad."

(4.9) *Il a la poisse.* "He has bad luck."

Four `NoAPrioriPol` frame-classes function in a similar manner. However, the syntactic frame needs to be specified, because it plays an important role as to which component in the sentence determines the verb's polarity. In the case of the frame-class `fclass_NoAPrioriPol_subj_eff_reflobj_dep` for instance, the reflexive verb's polarity and the effect on the subject are determined by the modifier (labelled "`dep`" for "dependent"). In sentence (4.10), the verb and the subject are marked as negative because of the negative adjective *tendu*.

(4.10) *La situation s'annonce <u>tendue</u>.* "The situation is expected to be tense."

We added the copula and `NoAPrioriPol` frame-classes because we expected this to contribute to a better assessment on document-level.

A number of intensifier and diminisher verbs are processed in a similar way to the `NoAPrioriPol` verbs. In the case of the intensifiers, the verb is marked as negative if something negative is intensified, as example (4.11) shows.

(4.11) *Les actes de violence ont augmenté.* "Acts of violence have increased."

We added three lexical items that are not a priori negative to the conditions of the rule processing intensifier verbs: *prix* "prices", *taxe* "tax" and *facture* "bill". Thus, the verb in (4.12) is going to be marked as negative too.

(4.12) *Les prix ont augmenté.* "Prices have risen."

As far as the subject is concerned (in this type of frame), the tag `INTENSIFIED` is appointed, or `A_NEG INTENSIFIED` in the case of the specified lexical items. For the positive marking of intensifier verbs, no lexical items are added to the conditions: here, the only condition is the presence of a positive child node, which will then be assigned the tag `INTENSIFIED`.

In the case of diminisher verbs, e.g. *diminuer* "reduce, decrease", we decided to treat them as verbs with a sort of shifting effect: if something negative is diminished, the verb is marked as positive, and vice versa. As far as the polarity of the subject or object in question is concerned, however, we decided not to reverse the polarity but to add the tag `DIMINISHED`. Consider the opposite of sentence (4.11), *Les actes de violence ont diminué.* "Acts of violence have decreased.": the act of decreasing can be considered to be of a positive nature, but the acts of violence themselves cannot. We anticipated the addition of these types of frame-classes to be useful for the assessment of the polarity on document-level as well.

4.3.3.5. Non-default Cases

So far, we have only dealt with default cases, i.e. verbs in simple affirmative sentences in active voice, with a default behaviour of the polarities in terms of effects and expectations. Thus, we decided to add passive voice constructions and compound tenses, as well as a number of non-default cases.

The frame-classes that include only the subject (with or without a modifier) are also marked in the passive voice. This is enabled via a syntactic frame (cf. section 4.3.3.2) that accepts an auxiliary and an optional preposition. In the case of the verb *accepter* "to accept" for example, which belongs to the frame-class `fclass_subj_verb_pos`, the following forms are marked:

(4.13) *Le comité accepte.* "The committee accepts."

(4.14) *Le comité a accepté.* "The committee has accepted."

(4.15) *La résolution est acceptée.* "The resolution is accepted."

(4.16) *La résolution a été acceptée.* "The resolution has been accepted."

(4.17) *La résolution est/a été acceptée par le comité.* "The resolution is/has been accepted by the committee."

The verb is marked as positive in each case. In example (4.17), the subject is additionally marked with a positive effect. The condition for the marking of a positive subject effect is the presence of, in this case, the preposition *par*. We opted for this restriction because although there are distinctive tags for tense auxiliaries and passive auxiliaries, the parser often has difficulties with the disambiguation, particularly in cases such as (4.16), or (4.14) as opposed to (4.15). In fact, unlike in English for example, the auxiliary *être* "to be" can be used for past tense as well as for passive voice. Compare the two sentences *Il est tombé.* "He fell/has fallen." vs. *Il est consolé (par quelqu'un)* "He is cheered up (by someone)". In order to avoid ambiguities, we added specific prepositions to the conditions that have to be fulfilled in order for the sentence to be processed and marked as a passive construction that has an effect on the subject, as in (4.17).

The aim of the verb-centered verb analysis is moreover to assess the sentence polarity more accurately by taking into account the polarity of the different agents and patients of a polar verb (see chapter 6.2). We identified a few cases where the default case is not applicable or where it needs to be reversed (i.e. the polarity of the verb and its effects and/or expectations need to be reversed).

1. Verbs belonging to the frame-class `fclass_subj_pos_objdir_eff_verb_pos`

are not marked if the object has a negative polarity (this is ensured with an additional condition in the verb-marking rule). This is based on the assumption that the negative object polarity conflicts with the polarity and effect of the verb, as sentence (4.18) shows for example. A number of lexical items are excluded from this, as the noun *problème* for instance (cf. (4.19)).

(4.18) *Ce genre de politique encourage le crime organisé.*
"This type of politics encourages organised crime."

(4.19) *Ils ont résolu le problème.* "They solved the problem."

2. Verbs of the frame-class `fclass_pos_subj_eff_verb_pos` are not marked if their subject is negative. A subsequent `ADD`-rule targets the concerned frame-class and marks the verb as negative if the subject is negative. Thus, the following sentence (4.20) is going to be computed as a negative sentence.

   (4.20) *La violence est cultivée dans ce milieu.*
   "Violence is culitivated in this milieu."

3. Similarly, the negative verbs *détruire* "to destroy" and *échouer* "to fail" are excluded from the default marking of their frame-class `fclass_neg_subj_eff_verb_neg` if their subject is negative: in this case, the verb is marked as positive instead of negative. The effect on the subject, however, is omitted. Thus, a more accurate differentiation between the polarities of sentences with the two afore-mentioned verbs is possible, such as in (4.21) and (4.22), where the former is negative and the latter positive.

   (4.21) *Les négociations de paix ont échoué.* "Peace negociations have failed."

   (4.22) *L'attentat terroriste a échoué.* "The terrorist attack failed."

4. As illustrated in the frame-specification examples (3.26) to (3.28) in section 3.3.1, the verb's polarity can vary according to the syntactic frame and/or type of object it takes. The verbs *lutter* "to fight" and *protester* "to protest" are marked as negative if they modify a direct object or a prepositional phrase, which is usually introduced by *contre* "against". If, however, the concerned verbs are the heads of a prepositional phrase introduced by *pour* "for", the verb is marked as positive, and the prepositional phrase is marked with the positive expectation tag `EXP_POS-POLCHUNK-HEAD`. Of course, this is a generalisation, but we estimate that in most cases, the constructions "to fight (or protest) for" and "to fight (or protest) against" have opposed polar connotations. The same procedure is chosen for the verb *battre* "to fight" when it is used in its reflexive form and with the preposition *pour*.

5. The verb *battre* calls for one more special treatment, when it is used with a direct object. Indeed, in the sports domain, it often occurs with *record* "record, best performance" or *adversaire* "opponent". In this case, a positive event is described. Hence, a substitution rule replaces the negative verb and subject effects with positive ones whenever the verb *battre* occurs with *record* or *adversaire*.

6. A last non-default case is analysed via substitute rules: in the case of the verbs *interdire* "to forbid", *se débarrasser* "to get rid of", *empêcher* "to prevent, to stop" and *arrêter* "to stop", the negative tag `verb_NEG` is substituted by its opposite `verb_POS` if the object is negative. Thus, the sentence in (4.23) will be analysed as negative, whereas the one in (4.24) will, more accurately, be assessed as positive.

   (4.23) *La commune a interdit les jeux dehors.*
   "The local authority has forbidden out-door activities."

   (4.24) *Une nouvelle loi interdit la fumée dans les endroits publics.*
   "A new law forbids smoking in public areas."

The first two non-default cases (**1.** and **2.**) outlined above are subject to broad generalisations. More precisely, they ought to be considered in the context of polarity conflicts in verb-frames, i.e. cases where the verb-polarity and the polarity effects or expectations should be modified according to the argument's polarity. While the polarity of the sentences in (4.18) and (4.20) is certainly negative, the question remains open whether a negative argument systematically results in a polarity conflict (or non-default case) with any verb pertaining to the mentioned frame-classes. We are going to answer this question with an empirical analysis in section 6.2.

We eventually decided not to include the negation of verbs. Due to its complexity (as alluded to in section 2.1.3), it is likely to constitute a project of its own. The implemented verb-component that we described here already comprises around 180 `ADD`- and `SUBSTITUTE`-rules for 72 frame-classes, which suggests a voluminous amount of rules for the modelling of verb negation as well.

The sentiment analysis is herewith completed. Its performance is assessed in the evaluation in chapter 5, as well as in chapter 6, where we are going to conduct empirical analyses on polarity conflict phenomena.

5. Evaluation

The evaluation of our sentiment analysis system was conducted at the document-level, although we focused on the phrase and sentence-level during the development and implementation phases of our system. Indeed, the evaluation system was developed for the needs of the ARGUS project in the first place. ARGUS focuses (among other things) on the analysis of media texts and on the subsequent summary of the retrieved analyses. The texts usually concern a specific topic, be it a product, a person or a company, e.g. in the context of reputation analysis. For this reason, and also in order to provide ARGUS with an overview of the model's potential and performance at an advanced stage of the project[1], a language-independent evaluation system at the document-level was implemented.[2] In the following, we are going to describe the evaluation material (section 5.1) and provide a technical description of the method (section 5.2). In section 5.3, we present and discuss the scored results.

5.1. Evaluation Corpus

We were provided with 53 French annotated articles in XML format by the ARGUS department for media analysis. The XML structure is composed of 7 tag pairs, as illustrated below:

```
<Artikel>
<Id>1219015</Id>
<Sprache>2</Sprache>
<Titel> ... </Titel>
<Text>
...
</Text>
<Coder>8</Coder>
```

[1] As described in section 5.3, the prototypes were tested several times. We made two evaluation runs before we actually finished developing our model, and one final evaluation run at the end.

[2] The evaluation system was written and implemented by Michael Amsler.

```
<Bewertung>3</Bewertung>
</Artikel>
```

The file ID corresponds to the file name, which moreover includes the rating code ("`Bewertung`"), as well as a language code ("`Sprache`"). The complete file name of the example above is `1219015_3_f.xml` for instance.[3] This information is used in the evaluation process. Tables 8 and 9 list the rating and language codes respectively.

Code	Polarity ("Bewertung")
1	Negative
2	Neutral
3	Positive
4	Ambivalent

Table 8.: ARGUS rating codes for the evaluation input files.

Code	Language ("Sprache")
1	German
2	French
3	Romansh
4	Italian
5	English

Table 9.: ARGUS language codes for the evaluation input files.

We estimated that 53 news articles were insufficient for the evaluation. Moreover, the articles were distributed unevenly among the different possible ratings: there were only 5 positive articles, but 14 negative, 15 neutral and 19 ambivalent articles. Thus, we annotated 47 additional articles that we selected from our news corpus. We deliberately increased the two polar categories, positive and negative, more than the other two because we wanted to assess the system capability of polarity classification in the first place. The rating of neutral texts is a challenging aspect of sentiment analysis in itself and is not covered in our project. The category "ambivalent" is, in our opinion, a relatively vague category and a category where inter-annotator agreement is even more important.[4] An overview of the properties of our final evaluation corpus is given in Table 10.

[3]The code "f" stands for "French", and differs from the code used inside the file (cf. Table 9).

[4]The assessment of a text's polarity is a subjective task. Ideally, at least 3 annotators should be involved. However, only one annotator was involved in the annotation of the ARGUS material and in the annotation of the additional material respectively.

Coding	Number of Articles
Negative	32
Neutral	17
Positive	32
Ambivalents	19
Total	**100**

Table 10.: Evaluation corpus properties.

5.2. Evaluation Method

The evaluation process [A18] is launched together with the path to one input folder [A19], which contains the coded XML files for evaluation (as described in 5.1 above), and the paths for two output folders. The first output folder, henceforth "overview output folder" [A20], will contain an overview of the collected values during the different stages of the evaluation process. The second output folder, henceforth "evaluation output folder" [A21], will contain all the files that will be generated during the evaluation process.

The rating codes (henceforth "*coded ratings*") are extracted from the input filenames. The file names are written together with their coded rating to the terminal screen, as well as to an output file [A22]. An overview of the evaluation corpus is also displayed in the terminal screen, indicating how many files are going to be processed in total, together with the amount of the respective coded ratings. A so-called "subprocess" [A23] is then launched, together with the path of the evaluation output folder. Each input file, i.e. each coded file from the input folder [A19], is converted into a temporary file, where the XML tags and the meta-data are removed and stored in variables. The temporary files are then processed by our French sentiment analysis system (as described in chapter 4). For each step of the system's pipeline, a corresponding version of the original input file is generated and stored in the evaluation output folder.

The rating of the newly analysed documents (henceforth "*calculated ratings*") is outsourced to another subprocess [A24]. The document rating is calculated on the basis of the ratio of prior polar words and of the resulting sentiment analysis: the percentage of positive and negative verbs and heads (NP, PP, effect and expectation) with respect to the total number of tokens determines whether a text is positive, negative, neutral or ambivalent. Verb heads have a greater weight than noun heads: if possible, the rating is based on the verb values only. The title of the document

(the first sentence of the document), is rated separately. The calculated document and title ratings are written progressively into a corresponding output file [A25]. Moreover, a statistics file is created for each document and stored in the evaluation output folder, featuring all the analysis information that was collected during the rating process.

The final VISL version that was generated by the sentiment analysis is moreover used as input for another subprocess [A26]. Here, the VISL format is converted into a complex XML tree, and the metadata, which now additionally contains the newly calculated document rating, is included once more. This XML structure is important for two tasks. First, it is used to produce a HTML version that provides a visualisation of the input and output of the system. The text is displayed once in its unmarked and immediately below in its marked version. The a priori polar words, as well as the words or chunks that were marked by the sentiment analysis system, are highlighted in different colors and fonts, whose meaning is explained in a caption below the marked text, next to a statistics table. The statistics table is filled with information extracted from the corresponding statistics file and summarises the polarity values of the text before and after it was processed by the sentiment analysis system, as well as the calculated document rating. Below, two more boxes are destined to display other content, such as the dependency structure of a sentence for instance.[5] Two shortened HTML versions are provided in appendix B, together with their respective initial XML version.

The second task for which the final VISL version is taken as input, is the generation of a new, final XML version of the article, which is analogue to the input file. It contains all the meta information of the original XML file, as well as the ratings and annotations that were produced by the sentiment analysis in XML syntax. If a noun was marked as a `A_NEG` and/or `NEG-NP-POLCHUNK-HEAD` by the sentiment analysis system for example, it is now enclosed by the tags `<negative></negative>`. This final XML file can be used for potential further evaluation or summary tasks. Once all the input files have been processed and rated, the actual evaluation process begins.

For this, a last subprocess is launched [A26]. The values in the files with the coded ratings [A22] and the calculated ratings [A25] are accessed, and an automatic comparison of the values of each pair of documents is made. If two documents have the same rating, a corresponding message is written to the terminal screen; if the ratings do not match, the coded and the calculated ratings, along with the concerned file name are displayed. According to the outcomes of the comparisons, the respective counters of false and true negatives, and of false and true positives

[5] For now, no content is displayed; a default error message appears.

are incremented for each type of document-rating, i.e. positive, negative, neutral and ambivalent document rating. Based on these counts, the recall, precision and F-scores for all the document-rating types are calculated, as well as the overall accuracy of the system. The calculated title ratings are moreover compared to the coded and calculated document ratings, in order to assess whether the titles have a predictive value as far as the document polarity is concerned. All the calculated values of the evaluation process are printed to the terminal screen as well as written in a summarised form to a third output file [A28].[6]

In summary, the evaluation process consists of the following automatic steps[7]:

1. The coded ratings of the input files are extracted and stored.
2. The input files are processed by the French Sentiment analysis system.
3. The polarity of the newly analysed documents is calculated (calculated ratings) and stored.
4. An HTML version of the calculated documents is generated. The HTML version provides a visualization of the sentiment analysis and of the resulting calculated document rating (see appendix B).
5. The coded and calculated ratings of each pair of document are compared. On the basis of these values, the precision, recall and F-score for each type of document-rating (positive, negative, neutral or ambivalent) are calculated, as well as the overall accuracy of the system.

The results of the evaluation are presented and discussed in the following section.

5.3. Results and Discussion

5.3.1. First Evaluation Runs and Corrections

As mentioned above, we conducted a number of test-runs with the initial set of evaluation data that was provided by ARGUS, before we finished the implementation of our model. This enabled us to correct mistakes or omissions in our lexicon and system. Table 11 shows the results of the very first evaluation run.

[6] All the overview files, i.e. the files summarising the coded, the calculated and the evaluated ratings are now all stored in the overview output folder.

[7] As mentioned above, the evaluation method was written and implemented by Michael Amsler. No language-specific adaptations were necessary.

Evaluation Test Run 1	**Recall**	**Precision**	**F-score**
Positive	1.0	0.16	0.27
Negative	0.36	0.45	0.4
Neutral	0.07	1.0	0.13
Ambivalents	0.21	0.44	0.29
Overall Accuracy:	**0.28**		

Table 11.: Results of the first evaluation run (53 articles - 14 negative, 15 neutral, 5 positive and 19 ambivalent).

While the overall accuracy rate is conspicuously low, the recall for positive document ratings and the precision of neutral document ratings are salient because of their high value: indeed, they reach the maximum value of 1.0. The counts of false positives reveal that many documents were erroneously rated as positive: out of 38 false positives, 27 concern positive document calculation. Hence, even documents that were coded as negative were calculated as positive by our system.

We manually analysed the calculated documents with the aid of the HTML versions (cf. section 5.2 and appendix B). This enabled us to quickly detect the reasons for the unsatisfying results of the first evaluation run:

- Numerous entries in the polarity lexicon, which were coded as positive, turned out to be too ambiguous, e.g. *justice* 1. "justice, fairness" 2. "law", or to work as intensifiers rather than as positive words, e.g. *considerable* or *expansion*. We manually deleted and modified around 70 lexicon entries.

- Many negative lexical items, particularly domain-specific words, were found to be missing in our polarity lexicon. We added approximately 30 lexicon entries, among which *faillite* "insolvency" and *endettement* "indebtedness".

- Shifters involved in indirect dependency relations (as described in section 4.3.1.2) had not been included in the model yet, as well as the necessary a priori disambiguations described in section 4.3.1.3 (a. and b.).

- The positive adverb *bien* "well" often occurs in constructions where it loses its positive polarity, e.g. in *bien des gens* "(quite) a lot of people" or *aussi bien que* "as well as".

- Similarly, the positive adjective *bon* "good" and the negative noun *coup* "blow, knock" loose their a priori polarity in specific constructions, e.g. *une bonne partie de* "much/many of", or *un coup de main* "a bit of help, a helping hand".

- Last but not least, we detected one parser error that we deemed important to correct manually: the shifter *sans* "without" is not always correctly lemmatised by the parser, especially if it is at the beginning of a sentence. Hence, it is not marked as shifter, which has consequences on the sentiment analysis.

We included the necessary disambiguations concerning the three last analysis errors in the NP shifter layer of our sentiment analysis system (cf. section 4.3.1, and 4.3.1.3 in particular), and updated our polarity lexicon and prior polarity marker (section 4.2). The results of the second evaluation test are displayed in Table 12 below.

Evaluation Test Run 2	**Recall**	**Precision**	**F-score**
Positive	1.0	0.33	0.5
Negative	0.5	0.35	0.41
Neutral	0.13	1.0	0.24
Ambivalents	0.42	0.5	0.46
Overall Accuracy:	**0.42**		

Table 12.: Results of the second evaluation run, after corrections and modifications (53 articles - 14 negative, 15 neutral, 5 positive and 19 ambivalent).

Except for the negative document rating precision, all values have increased. The positive document rating precision and F-score have almost doubled. The neutral document rating recall is still very low, which means that the system still has difficulties in rating neutral documents, but it has improved, as the F-score also suggests. The only value that does not show any improvement is the negative document rating precision, which has decreased. That is, especially documents that were coded as ambivalent are calculated as negative by the system. Nevertheless, as the overall accuracy score suggests, we managed to improve the system's performance through the corrections and additions outlined above.

5.3.2. Final Evaluation Runs

As described in section 5.1, we expanded the evaluation material, and we finished the implementation of our system. Hence, we implemented the last few verb-frame classes and the non-default verb cases described in section 4.3.3.5 and updated our polarity lexicon with a few last entries. Table 13 lists the results of the final evaluation run.

The scores of the positive document rating recall and of the neutral document rating precision have slightly decreased: they do not reach the maximum value any-

Final Evaluation Run	**Recall**	**Precision**	**F-score**
Positive	0.84	0.75	0.79
Negative	0.78	0.63	0.69
Neutral	0.24	0.8	0.36
Ambivalents	0.42	0.42	0.42
Overall Accuracy:	**0.64**		

Table 13.: Results of the final evaluation run (100 articles, cf. Table 10).

more, but score 0.84 and 0.8 respectively. This is however compensated with an increase of the respective precision and recall values, which results into a higher F-score for both rating categories. The F-score of the positive document rating has increased from 0.5 to almost 0.8, and the neutral document rating F-score has also risen from 0.24 to 0.36. In other words, the system now copes slightly better with neutral document classification. As far as the positive document rating is concerned, the system has become more precise by no longer over-classifying texts as positive. In direct comparison, the negative document classification also scores more satisfying results, but with lower scores. We believe that this is due to the fact that the system does not feature the negation of verbs yet. The ambivalent document rating category has actually scored lower results than in the second evaluation test run (cf. Table 12). As already mentioned, we think that this due to the vague definition of the category and to the lack of inter-annotator agreement in the coding of the evaluation corpus in the first place.[8] All in all, we think that the evaluation delivered satisfying results, with an overall accuracy score of 0.64.

Last but not least, we wish to assess whether the inclusion of verbs increases the performance of our sentiment analysis system. For this, we run the evaluation on the same corpus but without the verb-component in the sentiment analysis system [A29]. Table 14 displays the results that were scored without the verb-component. As estimated, a lower overall accuracy was scored. However, the difference is not as significant as we expected, and is moreover mainly due to a decrease of the scores in positive document rating. The precision score for negative document classification has noticeably increased for instance. The same is true for the recall of the neutral document classification, although the precision has decreased drastically. The scores for ambivalent document rating changed the least, but they also result in an increased F-score.

[8]There are, for example, a number of texts that we would have coded differently. However, we left them unchanged.

Final Evaluation Run, without Verbs	**Recall**	**Precision**	**F-score**
Positive	0.63	0.65	0.63
Negative	0.75	0.73	0.74
Neutral	0.41	0.47	0.44
Ambivalents	0.47	0.43	0.45
Overall Accuracy:	**0.60**		

Table 14.: Results of the final evaluation run, without verbs (100 articles, cf. Table 10).

We conclude that the inclusion of verbs improves the system's performance in classifying positive texts but, considering the ambivalent and neutral classification scores, it appears to slightly mislead the system's assessment. Balancing the verbs' weight in the automatic document rating is likely to improve the system's performance, but we did not verify this in the context of this evaluation. For the negative document rating, we still believe that the inclusion of negated verbs (besides the inclusion of affirmative negative verbs) would help to improve the system's performance. Moreover, if we consider the proportions of negative and positive lemmas in our lexicon (cf. Table 5, section 3.2.2.5), the evaluation values appear to be more comprehensible.

Of course, the hypothesis on the potential impact of negated verbs needs verifications that cannot be carried out in the context of the present project. We are moreover aware that our evaluation corpus is relatively small and lacks an inter-annotator agreement of the codings. These are also factors that have an effect on the evaluation results. One last explanation for the unexpectedly low impact of the verb-component is the parser's performance. Already during the implementation phase, we noticed that the parser encounters difficulties when coping with long and/or complex sentences. Subordinate clauses, especially those introduced by the preposition *de*, are often wrongly attached to the direct object or other elements of the sentence instead of the verb, as for example in *Ils accusent leurs concurrents de tricherie/d'avoir triché* "they accuse the other competitors of cheating". As described in section 4.3.3.2, the definition of templates, i.e. of dependency structures, is already a difficult task in itself, because of the parser's restricted tag-set. This in combination with wrong phrase attachments and lemmatisations of the parser is likely to prevent a significant number of verbs from triggering. Of course, this claim also needs quantitative verifications.

All in all, we think that the system's performance is satisfying, with an overall accuracy of 0.64, especially if we consider its potential for amelioration. The polarity lexicon can still be expanded and corrected, which would lead to a better coverage of a priori polar words. The same is true for the verb-component, where more verbs can still be added to the existing frame-classes, as well as the aforementioned negation of verbs.

We did not include the results of the title ratings in comparison with the coded and calculated document ratings because we could not draw any pertinent conclusions. The best results were scored without the verb-component, where 38 of the 100 title ratings matched the calculated document ratings, and 37 matched the coded document ratings. We think that additional data is necessary in order to evaluate the utility of separate title rating.

6. Empirical Analyses of Polarity Conflicts

In this chapter, we are going to discuss the empirical analyses that we conducted in order to verify the two main hypotheses that we formulated concerning the fine-grained polarity categories. The first hypothesis concerns conflicting polarities on the NP level (section 6.1). The second deals with conflicting polarities inside verb-polarity frames. We begin each section with outlining our hypotheses about the concerned conflict phenomenon, followed by a technical description of the method that we used for the empirical verification of our hypotheses. The third and last subsection discusses the results and whether the hypotheses could be verified or not. The discussions also aim at formulating theory drafts for possible future projects.

6.1. Polarity Conflicts of NP Constituents

6.1.1. Hypotheses

As described in section 4.3.2.2, we chose to systematically compute a NP with two conflicting polarities, i.e. composed of a positive adjective and a negative noun (or vice-versa), as a negative noun chunk (`NEG-NP-POLCHUNK`). This is based on the assumption that negative polarities are generally more dominant than positive polarities. More precisely, we believe that positive adjectives tend to affirm or intensify the negative polarity of their head-noun, whereas negative adjectives tend to reverse or negate the positive polarity of the positive noun that they modify. A second and related point of verification concerns the distinction between different fine-grained types of polarity. Assuming that the above hypothesis is tenable, we would like to verify whether we can determine the fine-grained polarity type at the NP-level as well. In view of the analysis of polarity conflicts in verb-frames (section 6.2), we need to assess, in the case of a noun phrase composed of a `A_POS` adjective and a `J_NEG` noun for example, whether the polarity type of the NP will be `A_NEG`,

`J_NEG` or whether it remains undetermined (or undeterminable).

In summary, we are going to verify the 2 following hypotheses:

1. Conflicting polarities in a noun phrase always result in a negative NP because
   a) a negative adjective reverses or negates the positive polarity of the noun that it modifies.
   b) a positive adjective functions as intensifier of the negative polarity of the noun that it modifies.
2. The distinction between different types of polarity allows for an automatic mapping of a specific polarity type on the NP level.

6.1.2. Method

In order to verify the above propositions, we took the readily analysed versions of our data sets as input material[1]. The analysis of the VISL files consists of the extraction of noun phrases that are composed of one adjective and one noun of opposing (conflicting) polarities [A30]. Two output files are generated: one contains all the cases which were found, i.e. the part-of-speech and polarity tag combinations, together with their respective lexical items [A31]. The second output file [A32] contains the list of all the types of part-of-speech and polarity tag combinations that were extracted, together with their number of occurrences.

First, the VISL files need to be pre-processed. We would like to be able to iterate through one sentence at a time, and inside the sentence, through each cohort, as well as through all the tags of each cohort. In order to do this, we processed the data in the following way: we excluded the cohorts that contain only the word form (see section 2.2.2) and generated a nested list from the analysis cohorts. In the VISL format, the sentences are separated by a blank line. We used these blank lines as a delimiter between sentences. Hence, we transformed our VISL data set into a list where each sentence is itself a list of lists. That is, each word (cohort) of a sentence is a list, where each item corresponds to a tag of the analysis line. The generation of a nested list out from the VISL data was quite challenging. However, the nested list format enabled us to easily carry out our analysis of NPs.

Indeed, it enabled us to adopt a look-ahead approach for our analysis in order to extract the following combinations:

1. A positive adjective preceding a negative noun (POS adj. + NEG noun)

[1] Cf. section 3.1 for the data sets description and chapter 4 for the sentiment analysis.

2. A positive adjective *following* a negative noun (NEG noun + POS adj.)
3. A negative adjective preceding a positive noun (NEG adj. + POS noun)
4. A negative adjective *following* a positive noun (POS noun + NEG adj.)

In French, the adjective can either precede or follow the noun that it modifies, which is why we want to take into account both types of NP constructions.[2] We minimised the number of iterations by merging the fine-grained polarity tags into two global tag-sets, positive and negative. In order to describe the extraction of NPs with conflicting polarities, we assume that the script is currently looking for the first combination that we listed above, namely "POS adj. + NEG noun".

For each sentence in the nested list, an iteration is made through each cohort. If an adjective is found in the cohort, a second iteration through each element of the current cohort takes place, looking for a positive tag.[3] If a positive tag is found, a kind of look-a-head procedure is applied: the index of the current cohort is put into a variable with regard to its position in the sentence. Based on this, the index of the next cohort is determined and the same procedure is repeated for the second constituent of the NP: if a noun is found in the next cohort, the iteration is repeated through the same cohort, looking for a negative tag. If a negative noun is found, the part-of-speech and polarity tag combination (henceforth "tag combination"), as well as their corresponding tokens are written into the first output file ([A31]). The tag combination is moreover put into a list. Thus, a list of tag combinations is generated for each data set. The two lists are merged, and the occurrences of each tag combination are counted and sorted. The total number of combinations and the number of different types of combinations are eventually written to the second output file ([A32]). The results are presented and discussed in the following section.

6.1.3. Results and Discussion

With the method described above, 3033 cases of noun phrases with conflicting polarities were found in the corpus, for 40 different types of combinations. Six types of combinations contained an analysis tag instead of a polarity tag (e.g. `NEG-NP-POLCHUNK-HEAD (NC)`[4] `+ A_NEG (ADJ)`). We assumed this to be owing to parsing or tagging errors, and removed them from the list. Moreover, several pairs of

[2]The construction noun + adjective actually occurs more often than its opposite version. This is in any case what our results suggested before we contracted the two construction types into one category (cf. the appendix frequency tables in B.1).

[3]Due to the different tagging processes of the sentiment analysis, the tags are not always at the same position in the cohort. This is why an iteration through the whole cohort is necessary.

[4]The tag "`NC`" stands for "*nom commun*" (proper name).

combination types needed yet to be contracted into one: the types were in fact also differentiated according to the positions of their components. `A_POS (ADJ) + A_NEG (NC)` and `A_NEG (NC) + A_POS (ADJ)` were considered as two distinct categories for instance. We contracted such pairs of types into one category type. Table 15 and Table 16 present a synthesis of the results after the aforementioned modifications. The complete frequency tables before and after the modifications can be found in appendix C.1.

Total Number of NPs with Conflicting Polarities	3027
Number of different combinations	18

Table 15.: Total number of noun phrases containing conflicting polarities.

	Tag Combination	**Number of Occurrences**
1.	A_POS adjective, A_NEG noun	1041
2.	A_NEG adjective, A_POS noun	691
3.	J_POS adjective, A_NEG noun	349
4.	A_POS adjective, J_NEG noun	242
5.	A_NEG adjective, J_POS noun	234
6.	J_NEG adjective, A_POS noun	155

Table 16.: Most Frequent tag combinations in noun phrases with conflicting polarities.

We chose to display and analyse the 6 most frequent combination types because, with more than 100 occurrences, we considered them to be most representative. The combination types which feature the `A_POS` and `A_NEG` polarities are by far the most frequent combination types. We would however like to relativize the high number of occurrences: a manual analysis of the lexical combinations quickly revealed that a number of ambiguous or erroneous entries persist in the lexicon. In the case of *vol régulier* for example, it is not only questionable whether the adjective *régulier* "regular, steady" is a polar word, but isolated from context, the ambiguity of the word *vol*, which can mean either "theft" or "flight", becomes apparent. Further cases contain adjectives with questionable or ambiguous polarities, e.g. *chers amis* ("expensive friends" vs. "dear friends"). Others should be modified in the lexicon as diminishers, as the case *légère toux* "mild cough" suggests.

Another reason why the number of occurrences of the combinations with `A_NEG`

and A_POS should be considered with care, is because a number of lexicalised expressions also increase the numbers: for example *succès fou* (lit. "mad success") is a lexicalised expression for "huge success". Similarly, the noun phrase *péché mignon*, literally "sweet sin", means having a weakness for something, but without a pejorative connotation, because usually it is a weakness for something positive and/or without negative consequences. These cases could be included in the lexicon via a multi-word list (as mentioned in section 3.2.2.5). We counted the number of lexicalised combinations we found, and concluded that they represent only a small fraction of the cases. Of the A_POS adjective and A_NEG noun combinations, 3.65% are lexicalised expressions. In the case of the A_NEG adjective and A_POS noun combinations, the percentage is even lower, namely 0.87%. Assuming that we did not detect all the lexicalised cases, we estimate the percentages to approximate 4% and 1% respectively. For the other types of tag combinations, we did not spot a conspicuous amount of wrong or ambiguous tags or lexicalised expressions.

In order to verify the hypotheses formulated in section 6.1.1, we manually selected a sample of 20 cases for each of the 6 most frequent combination types from our result file `outfileNPs.txt`: 10 ADJ+NC cases and 10 NC+ADJ cases. The samples are listed in appendix C.2, together with their evaluation. The evaluation consisted of two steps: first, we evaluated whether the noun phrases are overall negative, positive or ambiguous (cf. hypothesis 1 in section 6.1.1). We attributed the value "yes", "no" or "ambiguous" to each noun phrase. The results are listed in Table 17 below.

Tag Combination	**yes (negative)**	**no (positive)**	**ambiguous**
A_POS adjective, A_NEG noun	14/20 (70%)	4/20 (20%)	2/20 (10%)
A_NEG adjective, A_POS noun	17/20 (85%)	2/20 (10%)	1/20 (5%)
J_POS adjective, A_NEG noun	12/20 (60%)	4/20 (20%)	4/20 (20%)
A_POS adjective, J_NEG noun	18/20 (90%)	0/20 (0%)	2/20 (10%)
A_NEG adjective, J_POS noun	17/20 (85%)	2/20 (10%)	1/20 (5%)
J_NEG adjective, A_POS noun	19/20 (95%)	0/20 (0%)	1/20 (5%)
Total	**97/120 (81%)**	**12/120 (10%)**	**11/20 (9%)**

Table 17.: Verification of hypothesis 1 on NP's with conflicting polarities.

According to our manual evaluation, 97 out of 120 selected conflict cases (which corresponds to 81%) validate hypothesis 1. Indeed, if we consider the cases in (6.1) to (6.3), we can easily identify positive adjectives as intensifiers of their negative head nouns:

(6.1) *célèbre catastrophe* "famous catastrophe" (A_POS adjective + A_NEG noun)

(6.2) *glorieuse incertitude* "glorious uncertainty" (`J_POS adjective + A_NEG noun`)

(6.3) *violation délibérée* "deliberated violation" (`A_POS adjective + J_NEG noun`)

The extracted examples (6.4) to (6.6) illustrate how negative adjectives act as negators or shifters of their positive head nouns:

(6.4) *goût amer* "bitter taste"(`A_NEG adjective + A_POS noun`)

(6.5) *fausses innocences* "false innocence"(`A_NEG adjective + J_POS noun`)

(6.6) *ambition cynique* "cynical ambition"(`J_NEG adjective + A_POS noun`)

We evaluated 19% of the selected cases as either ambiguous or as contradicting our hypothesis. A contradiction of the hypothesis means that the overall polarity of a noun phrase should have been computed as overall positive instead of overall negative, such as in the cases (6.7) to (6.9).

(6.7) *sourire ravageur* "charming smile"(`A_NEG adjective + A_POS noun`)

(6.8) *bouleversante sincérité* "overwhelming sincerity"(`A_NEG adjective + J_POS noun`)

(6.9) *lutte antiterroriste* "antiterrorist fight"(`J_POS adjective + A_NEG noun`)

In (6.7) and (6.8), the computation is inaccurate due to the polysemy of the adjectives: the first meaning of *ravageur* is "destructing"; however, it gathers a positive connotation when combined with a positive noun. Similarly, the first meaning of *bouleversant* is "upsetting"or "distressing", but the meaning is changed into the positive adjective "overwhelming" if combined with a positive noun. While in (6.7) and (6.8) the meaning and consequently the polarity of the adjectives are modified because of the semantics of the noun, in (6.9) we have the opposite phenomenon: the negative noun *lutte* "fight" becomes positive because of the positive adjective.

Ambiguous cases are cases which are considered to be bi-polar, i.e. both negative and positive (with one prevailing polarity), and/or whose surrounding context is considered to be decisive for the overall polarity. Indeed, cases (6.10) to (6.12) are all open for interpretation and discussion:

(6.10) *confusion utile* "useful confusion"(`A_POS adjective + A_NEG noun`)

(6.11) *bon esclave* "good slave"(`A_POS adjective + J_NEG noun`)

(6.12) *réjouissantes excentricités* "uplifting eccentricities"(`A_POS adjective + J_NEG noun`)

These cases' dependency on the context also implies that they remain ambiguous and thus undeterminable as far as their overall polarity is concerned. Moreover, we should bear in mind that we are dealing with evaluative or in any case subjective language: the pragmatic or rhetorical effects of ambiguity, produced by the combination of contradictory meanings, may actually be wanted by the author, and the ambiguity intended to be kept. This is also closely related to the open challenge of irony in sentiment analysis, which for now remains undetectable for state-of-the-art systems. One possibility to deal with ambiguous and more or less lexicalised expressions (which result in a positive overall polarity such as in (6.7) to (6.9)) is, as already mentioned above, to include multiword lists in the lexicon, via a quantitative extraction and annotation of collocations for example. This would however go beyond the scope of our project and of the present analysis.

Overall, hypothesis 1 was proven to be validated, so that we proceeded to the second step of our evaluation, namely to the verification of hypothesis 2, which states that a specific polarity type (`A`, `F` or `J`) can be mapped to the noun phrase level. For this, we again opted for a manual evaluation: for each of the 96 noun phrases that corroborated hypothesis 1, we determined one dominant overall polarity type. The results are displayed in Table 18 below.

Tag Combination	**Overall Dominant Polarity Type**
A_POS adjective, A_NEG noun	A_NEG
A_NEG adjective, A_POS noun	A_NEG
J_POS adjective, A_NEG noun	A_NEG
A_POS adjective, J_NEG noun	J_NEG
A_NEG adjective, J_POS noun	J_NEG
J_NEG adjective, A_POS noun	J_NEG

Table 18.: Verification of hypothesis 2 on NP with conflicting polarities.

In consideration of possible expansions and ameliorations of the present model, we formulate one aggregation rule for each of the 6 analysed combination types:

1. If an **A_POS adjective** modifies an **A_NEG noun**, the A_POS adjective acts as an intensifier. The overall polarity of the **NP is A_NEG**.

2. If an **A_NEG adjective** modifies an **A_POS noun**, the A_NEG adjective shifts the positive polarity of the noun. The overall polarity of the **NP is A_NEG**.

3. If an **J_POS adjective** modifies an **A_NEG noun**, the J_POS adjective adds

a further qualification to the noun. The overall polarity of the **NP is A_NEG**.

4. If a **A_POS adjective** modifies an **J_NEG noun**, the adjective acts as an intensifier of the noun. The overall polarity of the **NP is J_NEG**.

5. If a **A_NEG adjective** modifies an **J_POS noun**, the adjective shifts the noun. The overall polarity of the **NP is reversed to J_NEG**.

6. If an **J_NEG adjective** modifies a **A_POS noun**, the meaning and polarity type of the adjective overrule those of the noun. The overall polarity of the **NP is J_NEG**.

The rules suggest a number of tendencies that need verification as far as the behaviour of the different polarity types is concerned:

- `J_NEG` and `J_POS` polarity words have a stronger polarity weight. They generally overrule the polarity type in the concerned expression, regardless of whether they constitute the noun or the adjective of the noun phrase (cf. rules 4, 5 and 6).

- Unlike `J_NEG` adjectives, `A_NEG` adjectives act as simple shifters (cf. rule 2 and 5). Examples from the combination `A_NEG` adjective + `F_POS` noun corroborate this assumption (e.g. amour tragique "tragic love" → `F_NEG`).

- Whereas `A_POS` adjectives tend to intensify the negative polarity of the `A_NEG` noun that they modify, the `J_POS` adjectives add a qualification to it, which means that they do not necessarily intensify the negative polarity of the noun. (cf. rule 1 and 3).

The last assumption is the one which necessitates the most thorough verification. Indeed, it concerns the most fragile rule, namely the rule about the category where only 60% of the cases validated hypothesis 1: `J_POS` adjective and `A_NEG` noun (cf. Table 17). This category contains the highest percentage of ambiguous cases and of cases which refuted the first hypothesis (20% respectively). The category `A_POS adjective` + `A_NEG noun` also has a 20% portion of cases refuting the hypothesis, but this is due to the high amount of lexicalised expressions in the samples. As illustrated in (6.9) and also in (6.13) below, a number of noun phrases with the combination `J_POS adjective` + `A_NEG noun` result in an overall `A_POS` polarity or even in an overall `J_POS` polarity (6.14).

(6.13) *louable souci* "laudable aim" (lit. "laudable worry")

(6.14) *vertueuse victime* "virtuous victim"

We could argue that since the noun *victime* is coded as `A_NEG passive`, the `J_POS` adjective rightly overrules the negative noun (see section 3.2.1 on the role and characteristics of the `passive` tag). However, if contrasted with the case in (6.15), the validity of this argument becomes questionable.

(6.15) *victimes innocentes* "innocent victims"

Much empirical verification remains open as to the utility of aggregation rules for resolving polarity conflicts at NP level. Larger samples of NP conflicts need to be evaluated and consequently additional rules, at least one for each of the 18 combination types, need to be elaborated and empirically verified. We are in fact aware that the used samples are small in size and that more than one annotator is needed. However, we believe that we could prove the potential utility for aggregation rules for fine-grained and conflicting polarity tags on the NP level. This opens up possibilities to improve the present model (also for the other languages) on the noun phrase level and, once implemented and applied to the analyses in the following section, an improvement is also possible on the verb-phrase level and consequently on the sentence-level.

6.2. Polarity Conflicts in Verb-polarity Frames

6.2.1. Hypotheses

As outlined in previous sections, the principle aim of the project is to model and implement the behaviour of polar verbs in sentiment analysis. This includes not only their standard or default behaviour, but also a non-standard behaviour with conflicting polarities. We believe that including potential polarity conflicts in verb-polarity frames (henceforth also "verb-frame") is going to distinctly increase the performance and accuracy of the system. Polarity conflicts in verb-frames occur when

a) the verb-frame expects an object of positive polarity, but a negative polarity chunk is encountered instead, and vice-versa.

b) a specific polarity effect is produced on the object or subject, but it is incompatible or contradictory with regard to the polarity of the concerned object or subject. The verb *encourager* "encourage" for example, has a positive effect on the object (or on the subject in a passive construction). Its verb-polarity frame further carries the implicit assumption that the argument is positive or,

if not polar, desirable and/or not negatively connotated (cf. 3.3.3 and 4.3.3.5).

As further described in section 4.3.3.5, a number of non-default and potential conflict cases are modelled in the sentiment analysis for verbs. However, apart from being incomplete, they are subject to broad generalisations. As far as the frame-class `fclass_subj_pos_objdir_eff_verb_pos` is concerned for example, if a negative object is encountered, the verb-frame is not marked. This implies that polarity conflicts occur regardless of the semantics of the verb and of the polarity type of a negative object. The examples (6.16) and (6.17) below show how problematic this assumption can be:

(6.16) *L'organisation aide/soutient les victimes du tremblement de terre.*
"The organisation helps/supports the victims of the earthquake."

(6.17) *L'organisation aide/soutient le crime organisé.*
"The organisation helps/supports organised crime."

The frame of the verb *aider* "help" predicts a positive effect on the object. We argue, however, that from a pragmatic and semantic point of view, the sentence in (6.16) results in an overall positive polarity, whereas in (6.17) it results in a negative overall polarity. The same is true if we replace the verb *aider* with *soutenir* "support". The verb-frame of the latter does not set a positive polarity effect but a positive polarity expectation on the object, which is not fulfilled in both cases, each with different consequences.

We argue that the different overall polarity of the sentences (6.16) and (6.17) is a consequence of the two different polarity types of the objects: the noun *victime* in (6.16) is annotated with `A_NEG passive`, whereas *crime* in (6.17) is a `J_NEG` noun (cf. section 3.2.1 for the meaning of the polarity labels). In other words, the hypothesis that we are going to verify is the following, divided into two sub-parts:

1. Polarity conflicts in verb-frames occur when the frame posits a certain polarity effect or polarity expectation on one of its arguments, and the opposite (unexpected) polarity occurs;

2. Whether we are actually dealing with a polarity conflict or not is eventually determined by the fine-grained polarity type of the concerned argument: we assume that `A` and `F` polarity types do not produce polarity conflicts, whereas `J` polarity types always do.

6.2.2. Method

In the context of verbs, it is important to take into account the dependency structure of the sentence, since we want to determine a specific verb argument and its polarity. Hence, we first proceeded to an analysis and tagging process of potential polarity conflicts in VISL. The resulting analyses in VISL are then automatically synthesized and evaluated. Our input material consisted again of our data sets (3.1), which were analysed for sentiment analysis (cf. chapter 4), with the difference that we excluded two rules dealing with non-default-cases (see non-default cases 1. and 2. in section 4.3.3.5), in order to assess as many potential polarity conflicts as possible.

The analysis and tagging process in VISL can be divided into two steps. The first step [A33] consist in the detection and marking of general polarity conflicts in verb-frames. Eight `ADD`-rules are divided into **five sets**. **Set 1** contains two general rules that deal with explicit expectations of verb-frames: if a tagged verb has a child-node that contains a polarity expectation tag and, in the same cohort, a polarity or analysis tag of the opposite polarity, the verb is marked with the tag `POL_CONFLICT`. If in one cohort both the tags `EXP_POS-POLCHUNK-HEAD` and `NEG-NP-POLCHUNK-HEAD` are found for example, the verb is marked. **Set 2** consists of two rules for the detection of polarity conflicts in passive constructions of positive verbs (cf. 2. in section 4.3.3.5). If the subject is negative, e.g. as in *la criminalité est acceptée* "criminality is accepted", the verb is marked. Similarly, the polarity of the subject is analysed for a negative tag in **set 3**. Here, we deal with potential conflicts in active constructions of verbs from the frame-class `fclass_pos_subj_eff`, as for example *Le plan néfaste a abouti.* "The evil plan succeeded.".

Sets 4 and 5 deal with frame-classes where the direct object receives a certain effect. In **set 4**, the concerned verb-frames posit a positive effect on the object; if the direct object is negative, the verb is marked as containing a polarity conflict (cf. 2., section 4.3.3.5). In **set 5**, an opposite and verb-specific rule is formulated for the verb *empêcher* "impede". The standard case assumes a negative effect on the object, regardless of its polarity. We include this rule in order to verify whether a conflict takes place if the object has an a priori negative polarity. The second step in the VISL analysis [A34] consists in filtering the previously marked polarity conflicts according to their polarity type. `SUBSTITUTE`-rules replace the `POL_CONFLICT` tag according to the encountered polarity tag. If a conflict has been found with the `ADD`-rule of set 1 described above for example, and if the child-node that triggered the mentioned rule contains a `J_NEG` tag, the `POL_CONFLICT` tag is replaced with the tag `POL_CONFLICT_J`. The rules are expanded in a way that also conflict cases where a negative noun phrase is composed of a neutral noun and a polar adjective are found. In case of noun phrases with conflicting polarities, the polarity type of the

negative component is considered, although, as we suggested in section 6.1.3 above, this generalisation cannot always be made.

Once the two data sets have been marked for type-specific polarity conflicts, an automatic synthesis of the findings is made, as well as the extraction of the marked sentences [A35]. The analysed VISL data sets are concatenated into one file, which is first processed into a nested list format, in the same way as the data in the analysis of conflicting NPs ([A30] described in section 6.1.2 above). Then a counting process is launched:

a) the total number of sentences with a "root" dependency tag, i.e. sentences which are likely to contain a verb.[5]

b) the number of verbs which were marked as either negative or positive.

c) the number of conflict cases with regard to the involved polarity type, as well as the total number of conflicts in general.

The counts are printed to the terminal screen (cf. Figure 6 in next section). The sentences containing polarity conflicts are moreover appended to a list. This list's content is analysed and, according to the encountered polarity type conflict, the lemma of the involved verb is extracted and put into a respective list. The concerned sentence is moreover written into an output file, also with regard to the polarity type involved in the conflict. In summary, three lists of verbs and three output files are generated. The verbs involved in A and F polarity types are concatenated, and one additional list is generated: it is the intersection of the `A` and `F` conflict verbs with the `J` conflict verbs. The 4 produced verb lists are also eventually printed to the terminal screen, and can also be found in appendix D.1.

6.2.3. Results and Discussion

The counts that result from the above described analysis are displayed in Figure 6 below. The complete verb-lists can be found in appendix D.1. The counts reveal that although the inclusion of the verb-component increases the system's performance (see section 5.3.2), only a small portion (15.42%) of the sentences with verbs contains verbs that were marked as either positive or negative. Negative verbs were tagged slightly more often than positive verbs (53.79% against 46.21%). Of the 30194 sentences with marked verbs, only 505 sentences (i.e. less than 1%) contain polarity conflicts inside verb-frames. Conflicts that involve *Appreciation* polarity

[5] Sentences without a "root" dependency tag usually do not contain verbs. They are often headings, dates, titles or authors' names.

types are by far the most frequent (81.19%), followed by conflicts involving *Judgment* polarity words (14.26%) and *Affect* polarity words (4.55%).

As mentioned in the previous section, we decided to merge the `A` and `F` con-

```
450468 sentences in total.
423829 sentences with a "root" dependency tag, i.e. sentences that are likely to contain
    a verb.
65340 sentences (15.42% of total with root) contain verbs that were marked as either
    POS or NEG.
30194 (46.21%) of the marked verbs were marked as positive;
35146 (53.79%) of the marked verbs were marked as negative.
505 sentences (0.77% of the sentences with marked verbs) appear to contain conflicting
    polarities.

23 (4.55%) with F_NEG or F_POS conflict,
72 (14.26%) with J_NEG or J_POS conflict,
410 (81.19%) with A_NEG or A_POS conflict.
```

Figure 6.: Results of the verb-conflict analysis conducted on the whole news corpus.

flict cases into one category (henceforth `A&F`) because we assume them to behave in the same way. 30 verbs were found to be involved in both `A&F` and `J` conflicts. They are listed alphabetically in Table 19 below. For a better overview, we further divided them into 3 categories: "effect" for verbs pertaining to frame-classes which predict a certain polarity effect on an argument; "expectation" for verbs belonging to frame-classes predicting a specific polarity expectation on the argument; the category "both" contains verbs pertaining to both types of frame-classes.

	"Effect"	**"Expectation"**	**Both**
Verbs	accepter, accueillir, adorer, aider, aimer, alimenter, apprécier, conforter, corriger, cultiver, encourager, favoriser, privilégier, soulager	accorder, défendre, désirer, gagner, mériter, offrir, oublier, perdre, permettre, profiter, promettre, prêner, soutenir, suggérer	assurer, nourrir

Table 19.: Verbs involved in both A&F and J polarity conflicts.

We proceeded to a manual analysis where for each verb we compared the extracted sentences of each conflict type. Some sentences contain parser errors and were marked as conflict sentence although no conflict occurred. Hence, for some verbs, only conflict sentences involving *Appreciation* polarity types were available. Thus, although only little data is available to verify our hypothesis, we gathered a few interesting findings. Our hypothesis formulated in section 6.2.1 above was validated, but only partly. More precisely, a number of aspects and assumptions concerning the concept of polarity conflict in verb frames-turned out to be vague,

but were clarified thanks to the empirical analysis.

As far as part **1** of the hypothesis is concerned, we realised that we made no distinction between the polarities of the verbs. The hypothesis implies that negative as well as positive verbs can be involved in polarity conflicts. However, a number of negative verbs that were marked as having conflicting objects in their frames call for a further precision. The concerned verbs are *oublier* "forget", *perdre* "loose", *manquer* "miss", *dissimuler* and *cacher*, both synonyms for "hide", and *empêcher* "prevent". The latter foresees a negative effect on its object; all the other negative verbs posit a positive polarity expectation on their objects. The analysis revealed that the verb *empêcher* as well as *oublier* and *perdre* should more accurately be modelled as shifter verbs: if the verb modifies a negative object, the verb polarity is positive, and vice versa, as illustrated in the examples below[6]. Only if the object has no polarity, we can assume a default negative verb polarity, as in (6.18).

(6.18) *empêcher la population de se rendre aux urnes* → `verb_NEG`
"prevent the population to go to the polls"

(6.19) *empêcher la guerre civile* → `verb_POS`
"prevent the civil war"

(6.20) *oublier ce côté un peu sévère* → `verb_POS`
"forget its slightly harsh character"

(6.21) *perdre son caractère discriminatoire* → `verb_POS`
"loose its discriminatory character"

For the verb *oublier*, we suggest that conflicts can occur, but only if a `J_NEG`, `A_NEG strong` or `A_NEG passive` object is encountered. In this case, the verb polarity would not be positive but negative, such as the examples in (6.22) and (6.23) suggest:

(6.22) *oublier trop aisément les stupidités* → `verb_NEG`
"forget the stupidities too easily"

(6.23) *oublier les victimes* → `verb_NEG`
"forget the victims"

Of course, empirical verifications are necessary. The same is true for the remaining negative verbs that we listed above: we suggest that they should be reclassified into frame-classes that foresee a negative effect on the object instead of a positive expectation. In the case of verbs with negative effects on their direct objects, we

[6]We shortened the sentences of our data to infinitive verb phrases for better readability. The full sentences, along with a number of additional sentences, can be found in appendix D.2.

claim that no conflicts can occur. Additional data is however needed to validate this claim. Except for the verbs *oublier* and *perdre*, the verbs listed in Table 19 are all of a positive nature. Regarding hypothesis part **1**, we thus claim that polarity conflicts occur only with positive verbs, with exception of the shifter verb *oublier*.

The analysis and comparison of the conflicts involving the positive verbs revealed that part **2** of the hypothesis was validated, but, again, only partly. Indeed, we identified four different conflict classes: the fine-grained polarity tags do determine whether a polarity conflict occurs or not, but they behave differently depending on the verbs. The categories and rules that we identified are the following:

1. `A_NEG`, `F_NEG` and `J_NEG` always cause conflicts, except if they are further labelled as `passive`.
   Verbs: alimenter, nourrir

2. `A_NEG` and `F_NEG` do not produce conflicts, except if they are further labelled as `strong`. `J_NEG` words and expressions always generate conflicts.
   Verbs: accepter, accueillir, accorder, adorer, aider, aimer, apprécier, encourager, défendre, permettre, privilégier, prôner, soutenir, suggérer

3. Nor `A_NEG`, `F_NEG` or `J_NEG` can cause conflicts. The verbs have a positive impact on the negative polarity: they diminish the negative polarity.
   Verbs: corriger, soulager

4. `A_NEG`, `F_NEG` and `J_NEG` always generate conflicts because

   a) the verb intensifies negative expressions.
      Verbs: assurer, conforter, cultiver, favoriser

   b) the proposition establishes a negative polarity effect or relation with regard to somebody or something.
      Verbs: mériter, désirer, offrir, profiter, promettre

We attributed the verbs *alimenter* and *nourrir*, synonyms for "feed" or "nourish" to the first category. We believe that these verbs act as intensifiers on negative expressions, except if the concerned expression is labelled as `passive`, as in the case of *victime* or *blessé* "injured (person)" for instance. The examples below corroborate our claims:

(6.24) *alimenter la polémique* "nourish the controversy"

(6.25) *nourrir le flot des travailleurs pauvres et des SDF* "feed the stream of poor workers and homeless people"

In the case of the second category, we think that the verbs are all semantically linked in that they express a sort of approval (e.g. *accepter* "accept", *encourager* "encourage"), care (*accueillir* "house, accommodate", *aider* "help") or sympathy (e.g. *aimer* "like, love", *apprécier* "appreciate"). Hence, we suggest that conflicts only occur if the negative expression is at the same time controversial, i.e. `A_NEG strong` or `J_NEG` (as in (6.26)). Whether other (weaker) instances of negative expressions, such as in (6.27), create controversy is, in our opinion, too dependent on the context and/or on the adopted perspective.

(6.26) *aimer les gouvernements corrompus et facilement corruptibles*
"like/appreciate corrupted and easily corruptible gouvernments"

(6.27) *accepter la démission* "accept the resignation"

The third category contains two verbs, *soulager* "relieve, alleviate" and *corriger* "correct". They more or less form a category of exception, because we expect no conflicts to occur with these verbs. Indeed, we think that they have a diminishing impact on the negative expressions that they modify (cf. the examples (6.28) and (6.29) below).

(6.28) *soulager la misère* "alleviate misery"

(6.29) *corriger les excès* "correct the excesses"

The fourth and last category deals with verbs that always generate conflicts, regardless of the type of negative polarity they modify. We divided the verbs into two distinct subcategories (a and b). The verbs listed in the subcategory a), *assurer* "assure", *conforter* "strengthen", *cultiver* "cultivate", *favoriser* "favour", have, as far as our data shows, an intensifying effect when they modify negative expressions. The type of negativity is not relevant (cf. (6.30) and (6.31)).

(6.30) a. *favoriser le désespoir* "favour despair"
b. *favoriser le terrorisme* "favour terrorism"

(6.31) a. *cultiver le mensonge et la trahison* "cultivate lies and betrayals"
b. *cultiver la haine* "cultivate hatred"

Unlike the verbs in category 1, the verbs in category 4 are unlikely to encounter A_NEG passive expressions, due to their semantic properties. The verbs in the subcategory b) differ substantially from those in subcategory a) in their semantics, but not in their conflict behaviour. We posit that the concerned verbs, *mériter* "deserve", *désirer* "wish for", *offrir* "offer", *profiter* "profit" and *promettre* "promise", always imply a negative effect on something or somebody and/or information about

a negative relationship between agents in the sentence or text. We provide an example for each verb in order to illustrate our hypothesis.

(6.32) *mériter une peine sévère* "deserve a severe punishment"

(6.33) *désirer la mort* "wish for/desire death"

(6.34) *offrir un score médiocre* "offer a poor score"

(6.35) *profiter des travailleurs illégaux* "take advantage of illegal workers"

(6.36) *promettre la vengeance* "promise revenge"

At this point, we would like to refer to a hypothesis formulated in related literature by the French linguist Yannick-Mathieu, which we presented in section 2.1.2. Yannick-Mathieu (1991) proposes a classification of French sentiment verbs (i.e. verbs expressing feelings or evaluative stances) not according to their syntactic properties, but according to their semantic properties. The author bases her classification on the hypothesis that verbs that do not necessarily share the same syntactic properties can have a semantic link (Yannick-Mathieu, 1991, p. 11). In the context of our empirical analysis of polarity conflicts in verb-frames, we think that this hypothesis is confirmed, also for verbs which are not strictly "sentiment verbs". We think that, despite the scarcity of data, our analyses showed that polarity conflicts in verb-frames are not only determined by the fine-grained polarity types, but also by the semantic properties of verbs.

7. Conclusion

7.1. Expectations and Outcomes

The aim of our project consisted in modelling and implementing a rule-based and verb-centered sentiment analysis for French, based on the principle of compositionality and featuring a new fine-grained polarity tag-set.

We built two expandable resources: on the one hand, the polarity lexicon, which contains around 4400 entries, and on the other hand, the 322 annotated verb-frames, that served as a basis for the modelling of verbs in the sentiment analysis pipeline. We assessed the performance of our system at the document-level, and gathered satisfying results with regard to the positive and negative document rating performance, and to the system's potential for improvement. Indeed, the overall accuracy amounted to 0.64. The impact of the verb-component on the system's performance, however, was less significant than we expected, and negative documents were actually classified more accurately without the verb-component. We believe that this is due to the absence of negation of verbs in our system, and to the fact that the amount of negative entries in the polarity lexicon is almost twice as high as the amount of positive words.

As far as the two more innovative components of our project are concerned, namely the verb-component and the fine-grained polarity labels, we conducted a number of empirical analyses, in order to assess their utility and possible applicability in potential extensions of the current project. We think that despite data sparseness, we were able to formulate interesting and verifiable theories, which can eventually lead to an amelioration of the present model: the system will become able to cope with polarity conflicts, and will potentially also be able to classify the analysed texts according to the prevalent polarity types. For example, the output of the system will consist of a polarity rating on the one hand, and on the other hand, in a classification of the most prominent opinion type, i.e. whether the language is predominantly emotional, factual, or rather moral and judgmental.

7.2. Limitations

There are three main limitations that persist in our project. The first is that there was only one annotator involved in the annotation of the lexicon. Thus, erroneous or ambiguous entries still remain present and are likely to impact on the parser's overall performance. However, as we already mentioned, the lexicon can still be modified and easily updated. The two other limitations are more difficult to rectify. The second limitation concerns the parser that we used for the linguistic pre-processing of the data. The parser not only proved to be slow, but also performed weakly in many aspects. Wrong lemmatisations and phrase attachments are two recurrent and frequent errors that the parser makes, and which have repercussions on the sentiment analysis: polar words are not marked (since the marking relies on the correct lemma form and PoS tag) and verbs do not trigger because the required dependency structure is not found.

The third and last limitation we encountered is concerned with the evaluation tasks rather than with the implementation and performance of the system. We used the CG implementation tool VISL CG-3, because it enabled us to efficiently elaborate and implement our rules based on dependency grammar. The disadvantage was, however, that due to the required parsing, the cohort-format and the tagging-processes, the size of the processed files increases considerably (roughly 10 times). Hence, although we regretted the sparseness of data at several times during the course of our project, we did not add data to our corpus. On the one hand, we did not want to slow down the parser further. On the other hand, we strongly suspect that our evaluation scripts would have had difficulties to cope with the data's size.

A few other limitations remain, that are of a linguistic nature. Our system does not distinguish between the different affirmative grammatical moods. From a pragmatic point of view, however, we think that two grammatical moods, subjunctive and conditional, should be excluded. Indeed, they do not actually affirm a statement, but rather express a desirable or possible state or event. Similarly, verbs in interrogative sentences should also be prevented from triggering. However, this is not possible in the present model. In order to prevent the triggering of specific grammatical forms or moods, additional information needs to be extracted from the CONLL output and to be included in the VISL format and rules. Last but not least, two of the greatest challenges of sentiment analysis, ambiguity and irony, also remain untouched, due to their complexity and unpredictability. Moreover, we think that there are many cases where a clear-cut decision on the polarity is very difficult or even impossible, as for example in the case of "to win a war".

7.3. Further Investigations

Our project shows great potential for further investigations and ameliorations. As far as the lexicon is concerned, the inclusion of multi-word lists (as discussed in section 6.1.3) is likely to improve the system's assessment of noun phrases with conflicting polarities and/or its assessment of lexicalised expressions. For example, the inclusion of a multi-word list could be done via the automatic extraction and manual annotation of collocations. Another missing aspect in the lexicon is the systematic inclusion of an intensity value of the polarities, also with regard to the concerned polarity type. For now, intensity is only included via the presence or absence of the tag "`strong`" and implicitly based on the assumption that the Judgment (J) polarity type is stronger than the Appreciation (A) and Affection (F) polarity types. We believe that a systematic and numeric quantification would help to increase the system's accuracy: a more precise assessment of the polarity weight of the individual words could also help to assess neutral and ambivalent documents more accurately.

The inclusion of polarity intensities is related to the complex and challenging topic of negation. As discussed in section 2.1.3, negation in French can feature more than one negator word, and, moreover, be combined with one or more modifiers: a simple inversion of the polarity is often inaccurate. The inclusion of polarity intensities and of more complex negation rules could help to sensibly improve the assessment of the resulting polarities from the word-level up to the document-level. Of course, the negation of verbs remains a great source for potential improvement. We believe that negative opinions are, at least in our French corpus, often expressed through negated verbs. Needless to say, this is only an hypothesis and not based on quantitative evidence. We however believe that this is also the reason why the verb-component for now has a greater impact on the positive document rating than on the negative document rating.

Last but not least, our system does not make use of the target identification component yet (unlike the German system for example). Targets, such as specific persons or companies, can be marked in the same way as polar words in the prior polarity marking process, i.e. via the inclusion of lists and marking rules in VISL. This would enable a target-specific sentiment analysis, which is certainly advantageous in a commercial or reputation analysis context. Moreover, target identification and rating could, in a further step, allow for an assessment of so-called polarity *relations* among actors (i.e. the targets). Thus, the evaluation of the text could also be adapted to the author's (or other targets') perspective.

References

P. Anand and K. Reschke. "Verb Classes as Evaluativity Functor Classes." Presented at the Interdisciplinary Workshop on Verbs. The Identification and Representation of Verb Features, Pisa, Italy, November 4-5, 2010.

F. Benamara, B. Chardon, Y. Mathieu, V. Popescu, and N. Asher. How do Negation and Modality Impact on Opinions? In *Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics*, ExProM '12, pages 10–18, Jeju, Republic of Korea, 2012.

S. Clematide and M. Klenner. Evaluation and Extension of a Polarity Lexicon for German. In 1st *Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA), ECAI 2010*, pages 7–13, Lisbon, Portugal, August 2010.

D. Dowty, R. Wall, and S. Peters. *Introduction to Montague Semantics*. D. Reidel, Dordrecht, 1981.

A. Esuli and F. Sebastiani. SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining. In *Proceedings of Language Resources and Evaluation (LREC)*, pages 417–422, Genova, Italy, 2006.

A. Fahrni and M. Klenner. Old Wine or Warm Beer: Target-specific Sentiment Anlaysis of Adjectives. In *Symposium on Affective Language in Human and Machine, AISB Convention*, pages 60–63, Aberdeen, Scotland, April 2008.

N. Godbole, M. Srinivasaiah, and S. Skiena. Large-Scale Sentiment Analysis for News and Blogs. In *Proceedings of the First International AAAI Conference on Weblogs and Social Media (ICWSM)*, Boulder, Colorado, USA, 2007.

M. Hu and B. Liu. Mining and Summarizing Customer Reviews. In *Proceedings of the* 10th *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2004)*, pages 168–177, Seattle, Washington, USA, 2004.

M. Klenner and S. Petrakis. Polarity Preference of Verbs: What could Verbs Reveal about the Polarity of their Objects? In G. Bouma, A. Ittoo, E. Métais,

and H. Wortmann, editors, *Proceedings of the* 17th *International Conference on Applications of Natural Language Processing to Information Systems (NLDB 2012)*, Lecture Notes in Computer Science, pages 35–46. Springer, 2012.

M. Klenner, A. Fahrni, and S. Petrakis. PolArt: A Robust Tool for Sentiment Analysis. In 17th *Nordic Conference on Computational Linguistics (NODALIDA 2009)*, pages 235–238, Odense, Denmark, May 2009a.

M. Klenner, S. Petrakis, and A. Fahrni. Robust Compositional Polarity Classification. In *Recent Advances in Natural Language Processing (RANLP)*, Borovets, Bulgaria, 2009b.

A. Küppers and L.-M. Ho-Dac. Private State in Public Media: Subjectivity in French Traditional and Online News. In 1st *Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA), ECAI 2010*, pages 22–29, Lisbon, Portugal, August 2010.

S. Läubli, M. Schranz, U. Christen, and M. Klenner. Sentiment Analysis for Media Reputation Research. In *Proceedings of KONVENS 2012 (PATHOS 2012 Workshop)*, pages 274–281, Vienna, Austria, 2012.

B. Liu. *Sentiment Analysis and Opinion Mining.* Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2012.

B. Liu, M. Hu, and J. Cheng. Opinion Observer: Analyzing and Comparing Opinions on the Web. In *Proceedings of the* 14th *international World Wide Web conference WWW-2005*, pages 10–14, Chiba, Japan, 2005.

J. R. Martin and P. R. R. White. *The Language of Evaluation: Appraisal in English.* Palgarve, London, 2005.

K. Moilanen and S. Pulman. Sentiment Composition. In *Proceedings of Recent Advances in Natural Language Processing (RANLP)*, pages 378–382, Borovets, Bulgaria, 2007.

A. Neviarouskaya, H. Prendinger, and M. Ishizuka. Compositionality Principle in Recognition of Fine-Grained Emotions from Text. In *Proceedings of the International Conference on Weblogs and Social Media, AAAI*, pages 278–281, San Jose, USA, May 2009a.

A. Neviarouskaya, H. Prendinger, and M. Ishizuka. SentiFul: Generating a Reliable Lexicon for Sentiment Analysis. In *Proceedings of International Conference on Affective Computing and Intelligent Interaction*, pages 363–368, Amsterdam, The Netherlands, September 2009b.

A. Neviarouskaya, H. Prendinger, and M. Ishizuka. Semantically Distinct Verb Classes Involved in Sentiment Analysis. In *Proceedings of the IADIS International Conference on Applied Computing 2009*, pages 27–34, Rome, Italy, November 2009c.

A. Neviarouskaya, H. Prendinger, and M. Ishizuka. @AM: Textual Attitude Analysis Model. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 80–88, Los Angeles, California, June 2010a.

A. Neviarouskaya, H. Prendinger, and M. Ishizuka. Affect Analysis Model: Novel Rule-based Approach to Affect Sensing from Text. *Natural Language Engineering*, Vol. 17(1):95–135, 2010b.

J. Nivre. Dependency Grammar and Dependency Parsing. Technical report, Växjö University: School of Mathematics and Systems Engineering, 2005.

P. Paroubek, A. Pak, and D. Mostefa. Annotations for Opinion Mining Evaluation in the Industrial Context of the DOXA project". In *LREC 2010. Seventh International Conference on Language Resources and Evaluation*, pages 1105–1112, Valetta, Malta, June 2010.

S. Petrakis, M. Klenner, É. Ailloud, and A. Fahrni. Composition multilingue de sentiments. In *Traitement Automatique des Langues Naturelles (TALN 2009)*, Senlis, France, June 2009.

L. Polanyi and A. Zaenen. Contextual Valence Shifters. In J. G. Shanahan, Y. Qu, and J. Wiebe, editors, *Computing Attitude and Affect in Text*, volume 20 of *The Information Retrieval Series*, pages 1–10. Springer Netherlands, 2006.

K. Reschke and P. Anand. "Extracting Contextual Evaluativity". Presented at the International Conference on Computational Semantics, Oxford, UK, January 22-24, 2011.

A. Rios and S. Clematide. Constraint Grammar mit VISL. Technical report, Institute of Computational Linguistics, University of Zurich, July 2010.

B. Sagot and D. Fišer. Building a Free French Wordnet from Multilingual Resources. In *Proceedings of Ontolex 2008*, pages 14–19, Marrakech, Maroc, 2008.

L. Tesnière. *Eléments de syntaxe structurale*. Klincksieck, Paris, 1959.

VISL-group. Institute of Language and Communication (ISK), University of Southern Denmark, 2013. http://beta.visl.sdu.dk/constraint grammar.html.

C. Whitelaw, N. Garg, and S. Argamon. Using Appraisal Taxonomies for Sentiment Analysis. In *Proceedings of CIKM-05, the ACM SIGIR Conference on Information and Knowledge Management*, pages 625–631, Bremen, Germany, 2005.

M. Wiegand, A. Balahur, B. Roth, D. Klakow, and A. Montoyo. A Survey on the Role of Negation in Sentiment Analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, pages 60–68, Uppsala, Sweden, July 2010.

T. Wilson, J. Wiebe, and P. Hoffmann. Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP)*, pages 347–354, Vancouver, Canada, 2005.

Y. Yannick-Mathieu. *Les verbes de sentiment: de l'analyse linguistique au traitement automatique.* Sciences du Langage. CNRS Editions, Paris, 1991.

L. Zhang, S. Ferrari, and P. Enjalbert. Opinion Analysis: the Effect of Negation on Polarity and Intensity. In *Proceedings of KONVENS 2012 (PATHOS 2012 Workshop)*, pages 282–290, Vienna, Austria, September 21st 2012.

A. Scripts and Filenames

Scripts and Filenames Referenced in Chapter 3

[A1] `filter_lexicon.py`
Python script used to split the polarity lexicon into individual files, according to the encountered word class.

[A2] `check_file.py`
Python script which verifies that the syntax of the polarity lexicon is accurate. The script outputs erroneous entries, which can then be manually corrected.

[A3] `MostFrequentLemmas_2files.py`
This python script was principally aimed at filtering out the most frequent verbs from our corpus. It was however also used to expand the polarity lexicon in general. The most frequent lemmas of a given word class are extracted from the corpus and compared to the entries of the corresponding lexicon file.

[A4] `add_infinitives.py`
A morphological approach is implemented in python in order to add the base forms of adjectives that derive from past participles to the polarity lexicon. Specific past participles endings are matched and replaced with infinitive endings via regular expressions. The newly lemmatised adjectives are written to the shell for manual analysis.

[A5] `fr_verbs.txt`
This text file originates from the original polarity lexicon. It was generated with the python script `filter_lexicon.py` (cf. [1]), and initially contained 1400 entries. After the expansion (cf. [3]), it lists 1499 entries.

[A6] `MostFreqVerbs.txt`
Contains the 1000 most frequent verb lemmas extracted from the corpus.

[A7] `ResultsComparison.txt`
This text file contains the results of the comparison operated by [3].

Scripts and Filenames Referenced in Chapter 4

[A8] `convert_bonsai2cg.pl`
This perl script is an adaptation of the script `convert_parzu2cg.pl`, which was originally written by Michael Amsler. It was adapted and provided by Nora Hollenstein. It converts the parser's CONLL output into VISL readable CG format via regular expressions.

[A9] `polmarker.bin`
This binary file corresponds to the compilation of two perl scripts. The perl scripts generate one CG file each: the first contains VISL-readable lemma lists (`owl_lists.cgf`, [A12]) generated from the polarity lexicon, and the second performs the actual polarity tagging (`polmarker.cgf`, [A14]). The compilation into the binary file `polmarker.bin` is done via the shell script `build_grammar.sh`: the two perl scripts are executed, as well as a VISL command, which eventually compiles the generated constraint grammar files into the binary file. The polarity marking can thus be executed in one step, via the command

```
$ cat TextInVISLformat.txt | vislcg3 --g polmarker.bin
```

[A10] `owl.txt`
This text file consists of the totality of the polarity lexicon, and serves as input for the perl script `build_owl_cg.pl` [A11].

[A11] `build_owl_cg.pl`
This perl script takes as input the lexicon file `owl.txt` [A10] and produces the output file `owl_lists.cgf` [A12].

[A12] `owl_lists.cgf`
This CG file is the result of the perl script [A11]. It contains all the polarity lexicon entries in different lemma lists, according to their word class and polarity.

[A13] `build_polmarkerfile_cg.pl`
This perl script writes the rules for the polarity tagging in VISL syntax to its output file `polmarker.cgf` [A14]. Via the VISL feature `INCLUDE`, the previously compiled lemma list file [A12] is included in the same output file.

[A14] `polmarker.cgf`
CG output file which results from [A13]. It contains the prior polarity tagging rules and actually carries out the polarity marking.

[A15] `np_shifter_FR.cgf`.
This CG file corresponds to the firs step in the sentiment analysis system, i.e. the modellation of the shifters and their effects on their heads or dependents.

[A16] `cord_np_pp_FR.cgf`
The second step of the sentiment analysis implemented in VISL. It consists of the assessment of NP, PP, and coordination polarities.

[A17] `verb_FR.cgf`
The third and final implementation step of the sentiment analysis system, where the verbs and their effects and/or expectations on their arguments are modelled.

Scripts and Filenames Referenced in Chapter 5

[A18] `evaluate.py`
Python script that launches the whole evaluation process. It is run with three command line arguments (see below).

[A19] `Artikel_F`
Input folder for the evaluation. Its path is the first command line argument for the script `evaluate.py` [A18], and contains the XML files for evaluation.

[A20] `eval_first_ov`
The path for this folder is the second command line argument for the evaluation script [A18]. At the end of the evaluation process, it contains three overview files: cf. [A22], [A25] and [A28].

[A21] `eval_first`
The path for this folder is the third command line argument for the evaluation script [A18]. It is filled with all the output files that are generated during the evaluation process.

[A22] `coded_ratings.txt`
This text file contains an overview of the input data: the input file names, together with their coded ratings.

[A23] `pipe_eval.py`
Python subprocess launched inside `evaluate.py` [A18]. It converts the input files into temporary files, and launches the sentiment analysis system pipeline: the temporary files are parsed and analysed for sentiment as described in

chapter 4. The generated files of each stage of the pipeline are stored in the output folder `eval_first` [A22].

[A24] `rate.py`
The rating of the newly analysed articles (cf. [A23]) is outsourced to this python script.

[A25] `ratings_calculated.txt`
This text file contains an overview of the calculated ratings of the output data (which results from [A23] and [A24]): the output file names, together with their calculated ratings.

[A26] `xmlify_eval.py`
Python subprocess launched inside `pipe_eval.py` [A23]. The analysed output files in VISL format are converted into a complex XML structure. The XML structure is converted (among other) in HTML format.

[A27] `calc_eval.py`
Second and last python subprocess launched inside the script `evaluate.py` [A18]. It carries out the actual evaluation calculations on the basis of the values in [A22] and [A25].

[A28] `ratings_evaluated.txt`
This text file contains the overview of the all the calculated values generated by the script [A27].

[A29] `evaluate_with_parsed_withoutV.py` It corresponds to the script `eval.py`, with the difference that the parsing and the verb-component are omitted when the sentiment analysis pipeline is launched. The parsed format of the documents was stored during a previous evaluation run.

Scripts and Filenames Referenced in Chapter 6

[A30] `extract_NPconflicts.py`
This python script takes as input two VISL files, i.e. the readily analysed versions of our data sets and looks for noun phrases with conflicting polarities.

[A31] `outfileNPs.txt`
First output file generated by [A30]. It contains all the cases which were found, i.e. the part-of-speech and polarity tag combination, together with their respective lexical items.

[A32] `outfileCombinations.txt`
Second output file generated by [A30]. It contains the list of all the types of part-of-speech and polarity tag combinations that were extracted, together with their number of occurrences.

[A33] `VerbAnalysis_step1.cgf`
This VISL script detects and marks general polarity conflicts in verb-frames, regardless of their polarity type.

[A34] `VerbAnalysis_step2.cgf`
VISL script which filters the general polarity conflicts marked by the previous script [A33] according to their polarity type: it replaces the general conflict tags with type-specific tags, according to the encountered polarity tag.

[A35] `extract_Vconflicts.py`
This python script synthetises the findings and extracts the marked sentences out of the VISL files that were previously analysed and marked for type-specific verb-frame conflicts (output of [A34]).

B. Article Format before and after Evaluation Process


```xml
<Artikel>
<Id>1219016</Id>
<Sprache>2</Sprache>
<Titel>Il a négligé les symptômes alarmants</Titel>

<Text>
Justice Un médecin généraliste condamné pour homicide involontaire à Bruxelles.

Une salle comble, coupée en deux. d' un côté les proches des victimes, de l' autre ceux
du médecin poursuivi devant le tribunal correctionnel de Bruxelles pour homicide
involontaire. La salle d' audience se plonge dans un silence froid pour entendre la
lecture du jugement.
La présidente du tribunal, Claire De Gryse, commence par rappeler les faits : Mélanie,
22 ans, consulte un médecin en Alsace. Il diagnostique une varicelle et conseille de
reprendre contact avec un médecin à son retour en Belgique. Quinze jours après, le 4
avril 2004, elle se plaint d' essoufflements et elle consulte un médecin belge, le 5,
qui lui conseille de se reposer.
Le 10 mai, elle prend contact par téléphone avec le prévenu, le Dr Bernard D. Elle fait
de la rétention d' eau, est fatiguée et à bout de souffle. Selon le médecin, c'est la
conséquence de sa varicelle et de sa session d' examens universitaires. Deux jours plus
tard elle le recontacte, avec les mêmes symptômes. Le Dr D. la rassure, lui dit de se
reposer, de surélever ses jambes et lui déconseille de se soumettre à une prise de sang.
Deux semaines après, Mélanie rappelle le médecin qui la reçoit le lendemain. Elle fait un
malaise dans son cabinet. Le médecin reste sur son impression de grosse fatigue et il
déconseille à la patiente de se rendre aux urgences. Mélanie Cailliau décédera dans la
nuit suivante d' une inflammation au coeur ( myocardite aiguë ).
Les parents de Mélanie avaient porté les faits à la connaissance du Conseil de l' ordre
des médecins. En l' absence de toute réaction, rappelle le jugement, ils ont alors déposé
plainte en justice.
Le tribunal s' est interrogé sur ce que savait ou ignorait le médecin au moment de poser
son diagnostic. « Il n' y a aucune raison de penser qu'elle se serait abstenue de parler
d' une prise de poids de 5 kilos, tout comme de son essoufflement, » ponctue le jugement.
Qui estime que le médecin généraliste a négligé des symptômes alarmants et refusé de faire
hospitaliser sa patiente malgré sa demande. « Par l' ensemble de son comportement, basé
sur une idée préconçue, le Dr D. a commis une faute en liaison directe avec la mort de
Mélanie. »
Le tribunal l' a donc condamné à un an de prison avec un sursis total et une amende de
4.125 euros. Il devra en outre dédommager les parents de la jeune victime en leur versant
une somme de 50.000 euros. Le condamné n' a pas souhaité commenter cette décision.
Par contre, le père de Mélanie, Philippe Cailliau, a souligné qu'il menait un combat qui
allait au-delà de la mort de sa fille. Un combat pour la transparence des procédures
( à l' Ordre notamment ), un combat pour éviter que les affaires concernant les médecins
soient enterrées. « C'est un combat pour responsabiliser les médecins, pas un combat
contre leur corporation. Pour éviter aussi qu'il y ait d' autres petites Mélanie. »

</Text>
<Coder>8</Coder>
<Bewertung>1</Bewertung>
</Artikel>
```


Figure A.1.: XML input format of an article with negative coded rating (`1219016_1_f.xml`).

Output der Chain

Output der Chain

alle Sätze mit Markierungen:

Il a négligé les symptômes alarmants . Justice Un médecin généraliste condamné pour homicide involontaire à Bruxelles . Une salle comble , coupée en deux . d' un côté les proches des victimes , de l' autre ceux du médecin poursuivi devant le tribunal correctionnel de Bruxelles pour homicide involontaire . La salle d' audience se plonge dans un silence froid pour entendre la lecture du jugement . La présidente du tribunal , Claire De Gryse , commence par rappeler les faits : Mélanie , 22 ans , consulte un médecin en Alsace . Il diagnostique une varicelle et conseille de reprendre contact avec un médecin à son retour en Belgique . Quinze jours après , le 4 avril 2004 , elle se plaint d' essoufflements et elle consulte un médecin belge , le 5 , qui lui conseille de se reposer . Le 10 mai , elle prend contact par téléphone avec le prévenu , le Dr Bernard D . Elle fait de la rétention d' eau , est fatiguée et à bout de souffle . Selon le médecin , c' est la conséquence de sa varicelle et de sa session d' examens universitaires . Deux jours plus_tard elle le recontacte , avec les mêmes symptômes . Le Dr D . la rassure , lui dit de se reposer , de surélever ses jambes et lui déconseille de se soumettre à une prise de sang . Deux semaines après , Mélanie rappelle le médecin qui la reçoit le lendemain . Elle fait un malaise dans son cabinet . Le médecin reste sur son impression de grosse fatigue et il déconseille à la patiente de se rendre aux urgences . Mélanie Cailliau décédera dans la nuit suivante d' une inflammation au coeur (myocardite aiguë) . Les parents de Mélanie avaient porté les faits à la connaissance du Conseil de_l'_ordre_des médecins . En l' absence de toute réaction , rappelle le jugement , ils ont alors déposé plainte en justice . Le tribunal s' est interrogé sur ce que savait ou ignorait le médecin au moment de poser son diagnostic . « Il n' y a aucune raison de penser qu' elle se serait abstenue de parler d' une prise de poids de 5 kilos , tout comme de son essoufflement , » ponctue le jugement . Qui estime que le médecin généraliste a négligé des symptômes alarmants et refusé de faire hospitaliser sa patiente malgré sa demande . « Par l' ensemble de son comportement , basé sur une idée préconçue , le Dr D . a commis une faute en liaison directe avec la mort de Mélanie . » Le tribunal l' a donc condamné à un an de prison avec un sursis total et une amende de 4.125 euros . Il devra en_outre dédommager les parents de la jeune victime en leur versant une somme de 50.000 euros . Le condamné n' a pas souhaité commenter cette décision . Par contre , le père de Mélanie , Philippe Cailliau , a souligné qu' il menait un combat qui allait au-delà_de la mort de sa fille . Un combat pour la transparence des procédures (à l' Ordre notamment) , un combat pour éviter que les affaires concernant les médecins soient enterrées . « C' est un combat pour responsabiliser les médecins , pas un combat contre leur corporation . Pour éviter aussi qu' il y ait d'_autres petites Mélanie . »

Legende:	
unterstrichen, kursiv:	*Polchunk*
grüner Hintergrund, kursiv:	*POS apriori*
roter Hintergrund, kursiv:	*NEG apriori*
hellblauer Hintergrund, kursiv:	*NEU apriori*
oranger Hintergrund, kursiv:	*INT apriori*
hellgrauer Hintergrund, kursiv:	*DIM apriori*
Fuchsia Hintergrund, kursiv:	*SHI apriori*
kursiv	*Children of Polchunk; no self-polarity*
Graue Schrift:	keine Polarität

Statistiken:	Anzahl	Prozent der Token
apriori POS:	2	0.35
apriori NEG:	28	4.86
POS NP Heads:	2	0.35
NEG NP Heads:	27	4.69
POS PP Heads:	2	0.35
NEG PP Heads:	13	2.26
verb_POS:	1	0.17
verb_NEG:	7	1.22
POS EFFECT Heads:	1	0.17
NEG EFFECT Heads:	14	2.43
POS EXPECTATION Heads:	1	0.17
NEG EXPECTATION Heads:	3	0.52
Total Tokens:	576	
Total positive Heads/verbs	7	1.22
Total negative Heads/verbs	64	11.11

Bewertung ist: 1, also negativ

Figure A.2.: HTML output for an article with negative calculated rating (`1219016_HTMLoutput.html`; implementation by Michael Amsler).

```
<Artikel>
<Id>1219042</Id>
<Sprache>2</Sprache>
<Titel>Jeunes, brillants, passionnés</Titel>

<Text>
Social. Ils ont entre 18 et 30 ans, des idées à foison, un joli budget.
Le Fonds belge de la vocation récompense quinze jeunes en devenir. Plus qu'une bourse, un
projet de vie.

Epatant ! » Jean Van Hamme, scénariste de BD ( XIII, Thorgal, Black et Mortimer... ) est
un « parrain » comblé : ce soir, à l' hôtel Métropole, il remettra de ses mains d'
artiste le prix 2006 de la Fondation belge pour la vocation.
Quinze jeunes de 18 à 30 ans seront à l' honneur. Ils sont cinéastes, infirmières de rue,
photojournaliste ou restaurateur d' objets anciens. Ils ont été sélectionnés par un jury
prestigieux ( scientifiques de renom, experts des arts et des lettres... ) parmi 320
candidats. Ils bénéficieront d' une bourse de 10.000 euros. Et ils succéderont à d'
illustres lauréats ( la danseuse Anne Teresa de Keersmaeker, le volcanologue Gerald
Ersnt, le cinéaste Frédéric Fonteyne... ). « Cette bourse est un tremplin pour la vie ;
un plus pour leur CV, une reconnaissance de leur travail et de leur passion », se
félicite le président de la Fondation, Alain Philippson.
Depuis 1963, cet organisme privé fondé par Emile Bernheim, propriétaire de l' Innovation,
a soutenu 582 jeunes « sans distinction de sexe, d' origine sociale ou de formation
intellectuelle ». Des vocations dans des domaines aussi divers que les sciences exactes,
l' architecture ou les métiers d' arts. En 43 ans, près de 3 millions d' euros ont ainsi
été distribués par la Fondation pour « parfaire ou mener à bien des vocations ». Le Soir
a choisi de vous présenter, de façon succincte, les huit lauréats francophones 2006 qui,
souligne le jury, « sont dotés d' une forte personnalité et font preuve d' une volonté et
d' un enthousiasme persévérants et attachants ».Une manière de montrer que la jeunesse
est aussi porteuse d' idées créatrices, de projets pertinents : « Le monde actuel n' est
pas forcément rigolo pour ces jeunes, poursuit Van Hamme. Ce type de projet est très
positif. Il faut susciter ces vocations qu'on a tous plus ou moins enfouies. Moi, par
exemple, j'ai toujours écrit des nouvelles, des petites histoires. Mais c'est un
excellent prof de français qui a allumé la petite flamme. »
Les lauréats seront suivis de près. « On va les aider sous forme de conseils ou de
collaboration », explique-t-on à la Fondation. l' un pour monter son école de cinéma. l'
autre pour surveiller un volcan au Mexique. Le suivant pour réaliser des contes avec les
enfants psychotiques. Des jeunes en devenir. Qui ont pris leur destin en main. Avec
force, créativité et intelligence.
</Text>
<Coder>8</Coder>
<Bewertung>3</Bewertung>
</Artikel>
```


Figure A.3.: XML input format of an article with positive coded rating (`1219042_3_f.xml`).

Output der Chain

Output der Chain

alle Sätze mit Markierungen:

Jeunes , *brillants* , *passionnés* . Social . Ils ont *entre* *18* et 30 ans , des idées à foison , un *joli* *budget* . Le Fonds belge de la vocation récompense quinze jeunes en devenir . Plus qu' une bourse , un projet de vie . Epatant ! » *Jean* Van Hamme , *scénariste* de BD (*XIII* , *Thorgal* , Black et Mortimer ...) est un « parrain » comblé : ce soir , à l' hôtel Métropole , il remettra de ses mains d' artiste le prix 2006 de la Fondation belge pour la vocation . *Quinze* *jeunes* *de* 18 *à* 30 ans *seront* *à* l' *honneur* . Ils sont *cinéastes* , *infirmières* de rue , photojournaliste ou restaurateur d' objets anciens . Ils ont été sélectionnés *par* un *jury* *prestigieux* (*scientifiques* *de* renom , *experts* des arts et des lettres ...) parmi 320 candidats . *Ils* *bénéficieront* *d'* une *bourse* de 10.000 euros . Et ils succéderont à d' illustres lauréats (la danseuse Anne Teresa de Keersmaeker , le volcanologue Gerald Ersnt , le cinéaste Frédéric Fonteyne ... « Cette bourse est un tremplin pour la vie ; un plus pour leur CV , une *reconnaissance* *de* leur travail et *de* leur *passion* » , se félicite le président de la Fondation , Alain Philippson . Depuis 1963 , cet organisme privé fondé par Emile Bernheim , propriétaire *de* l' *Innovation* , a *soutenu* 582 jeunes « *sans* *distinction* *de* sexe , d' origine sociale ou *de* *formation* *intellectuelle* » . Des vocations dans des domaines aussi divers que les *sciences* *exactes* , l' architecture ou les métiers d' arts . En 43 ans , près_de 3 millions d' euros ont ainsi été distribués par la Fondation pour « parfaire ou mener à *bien* des vocations » . Le Soir a choisi de vous présenter , de façon succincte , les huit lauréats francophones 2006 qui , souligne le jury , « sont dotés d' une forte personnalité et font *preuve* *d'* une *volonté* *et* *d'* un *enthousiasme* *persévérants* et *attachants* » . Une manière de montrer que la *jeunesse* est aussi porteuse d' idées créatrices , de *projets* *pertinents* : « *Le* *monde* *actuel* *n'* *est* *pas* forcément *rigolo* pour ces jeunes , poursuit Van Hamme . *Ce* *type* *de* projet *est* *très* *positif* . Il faut susciter ces vocations qu' on a tous plus_ou_moins enfouies . Moi , par_exemple , j' ai toujours écrit des nouvelles , des petites histoires . Mais c' est un *excellent* *prof* *de* français qui a allumé la petite flamme . » Les lauréats seront suivis de près . « On va les aider sous_forme_de conseils ou de collaboration » , explique -t-on à la Fondation . l' un pour monter son école de cinéma . l' autre pour surveiller un volcan au Mexique . Le suivant pour réaliser des contes *avec* les *enfants* *psychotiques* . Des jeunes en devenir . Qui ont pris leur destin en main . *Avec* *force* , créativité et *intelligence* .

Legende:	
unterstrichen, kursiv:	*Polchunk*
grüner Hintergrund, kursiv:	*POS apriori*
roter Hintergrund, kursiv:	*NEG apriori*
hellblauer Hintergrund, kursiv:	*NEU apriori*
oranger Hintergrund, kursiv:	*INT apriori*
hellgrauer Hintergrund, kursiv:	*DIM apriori*
Fuchsia Hintergrund, kursiv:	*SHI apriori*
kursiv	*Children of Polchunk; no self-polarity*
Graue Schrift:	keine Polarität

Statistiken:	Anzahl	Prozent der Token
apriori POS:	18	3.52
apriori NEG:	1	0.2
POS NP Heads:	19	3.72
NEG NP Heads:	1	0.2
POS PP Heads:	9	1.76
NEG PP Heads:	1	0.2
verb_POS:	5	0.98
verb_NEG:	0	0.0
POS EFFECT Heads:	4	0.78
NEG EFFECT Heads:	0	0.0
POS EXPECTATION Heads:	1	0.2
NEG EXPECTATION Heads:	0	0.0
Total Tokens:	511	
Total positive Heads/verbs	38	7.44
Total negative Heads/verbs	2	0.39

Bewertung ist: 3, also positiv

Figure A.4.: HTML output for an article with positive calculated rating (`1219042_HTMLoutput.html`; implementation by Michael Amsler).

C. NPs with Conflicting Polarities

C.1. Most Frequent Polarity Tag Combinations

	NP Tag Combination	Number of Occurrences
1.	A_NEG noun, A_POS adjective	723
2.	A_POS noun, A_NEG adjective	465
3.	A_NEG noun, J_POS adjective	322
4.	A_POS adjective, A_NEG noun	318
5.	A_NEG adjective, A_POS noun	226
6.	J_NEG noun, A_POS adjective	181
7.	A_NEG adjective, J_POS noun	133
8.	A_POS noun, J_NEG adjective	125
9.	J_POS noun, A_NEG adjective	101
10.	A_POS adjective, J_NEG noun	61
11.	F_POS noun, A_NEG adjective	48
12.	F_NEG noun, A_POS adjective	42
13.	J_NEG noun, J_POS adjective	39
14.	J_NEG adjective, A_POS noun	30
15.	J_POS adjective, A_NEG noun	27
16.	A_POS adjective, F_NEG noun	25
17.	A_NEG noun, F_POS adjective	24
18.	A_NEG adjective, F_POS noun	21
19.	J_POS noun, J_NEG adjective	18
20.	F_NEG noun, J_POS adjective	18
21.	A_POS noun, F_NEG adjective	10
22.	F_NEG adjective, A_POS noun	10
23.	J_NEG adjective, J_POS noun	10
24.	F_POS noun, J_NEG adjective	10
25.	F_POS noun, F_NEG adjective	7
26.	J_POS adjective, F_NEG noun	7

27.	J_POS adjective, J_NEG noun	6
28.	J_NEG noun, F_POS adjective	5
29.	J_POS noun, F_NEG adjective	3
30.	F_POS adjective, A_NEG noun	3
31.	F_NEG adjective, J_POS noun	3
32.	F_NEG noun, F_POS adjective	3
33.	NEG-NP-POLCHUNK-HEAD, A_NEG adj.	2
34.	F_POS adjective, F_NEG noun	1
35.	F_POS adjective, J_NEG noun	1
36.	A_NEG adj., NEG-NP-POLCHUNK-HEAD	1
37.	J_NEG adjective, F_POS noun	1
38.	COORD_HEAD_POS, J_POS adjective	1
39.	POS-NP-POLCHUNK-HEAD, A_POS adj.	1
40.	A_POS adj., POS-NP-POLCHUNK-HEAD	1

Table B.1.: Most frequent tag combinations in conflicting NPs (*all*).

	NP Tag Combination	**Number of Occurrences**
1.	A_POS adjective, A_NEG noun	1041
2.	A_NEG adjective, A_POS noun	691
3.	J_POS adjective, A_NEG noun	349
4.	A_POS adjective, J_NEG noun	242
5.	A_NEG adjective, J_POS noun	234
6.	J_NEG adjective, A_POS noun	155
7.	A_NEG adjective, F_POS noun	69
8.	A_POS adjective, F_NEG noun	67
9.	J_POS adjective, J_NEG noun	45
10.	J_NEG adjective, J_POS noun	28
11.	F_POS adjective, A_NEG noun	27
12.	J_POS adjective, F_NEG noun	25
13.	F_NEG adjective, A_POS noun	20
14.	J_NEG adjective, F_POS noun	11
15.	F_NEG adjective, F_POS noun	7
16.	F_POS adjective, J_NEG noun	6
17.	F_NEG adjective, J_POS noun	6
18.	F_POS adjective, F_NEG noun	4

Table B.2.: Most frequent tag combinations in conflicting NPs (*contracted*).

C.2. Samples of Conflicting NPs with Annotation

NP Tag Combination	Noun Phrase	Annotation	Overall Polarity
1. *A_POS adjective, A_NEG noun*	obscurité spirituelle	negative	A_NEG
	hostilité pure	negative	A_NEG
	confusion utile	ambiguous	-
	menaces précises	negative	A_NEG
	asservissement spirituel	negative	A_NEG
	combats victorieux	positive	A_POS
	handicapés bénéficiaires	positive	A_POS
	dépenses coquettes	negative	A_NEG
	victimes consentantes	negative	A_NEG
	guerres préventives	negative	A_NEG
A_NEG noun, A_POS adjective	profond malaise	negative	A_NEG
	délicieux frissons	positive	A_POS
	célèbre manque	negative	A_NEG
	bonne averse	negative	A_NEG
	jolie brèche	negative	A_NEG
	énergiques pressions	negative	A_NEG
	fameux délai	ambiguous	-
	joyeuse solitude	positive	A_POS
	célèbre catastrophe	negative	A_NEG
	franche agressivité	negative	A_NEG
2. *A_POS adjective, A_NEG noun*	libération explosive	negative	A_NEG
	aide tardive	negative	A_NEG
	amis absents	negative	A_NEG
	sourire ravageur	positive	A_POS
	goût amer	negative	A_NEG
	allocations limitées	negative	A_NEG
	concessions douloureuses	negative	A_NEG
	humour dévastateur	negative	A_NEG
	honoraires disproportionnés	negative	A_NEG

	plaisirs solitaires	negative	A_NEG
A_NEG adjective,	périlleux équilibre	negative	A_NEG
A_POS noun	impardonnable triomphe	negative	A_NEG
	mauvais goût	negative	A_NEG
	fragile bonheur	negative	A_NEG
	faux ami	negative	A_NEG
	encombrant cadeau	negative	A_NEG
	féroce volonté	ambiguous	-
	redoutable précision	positive	A_POS
	écrasante victoire	negative	A_NEG
	piètre consolation	negative	A_NEG
3. *A_NEG noun,*	blessures volontaires	negative	A_NEG
J_POS adjective	victimes innocentes	negative	A_NEG
	lutte antiterroriste	positive	A_POS
	déficit intellectuel	negative	A_NEG
	bataille digne	ambiguous	-
	ennui distingué	negative	A_NEG
	rupture légale	ambiguous	-
	délai raisonnable	ambiguous	-
	crises humanitaires	negative	A_NEG
	sanction morale	negative	A_NEG
J_POS adjective,	digne catastrophe	negative	A_NEG
A_NEG noun	aimable dilettante	negative	A_NEG
	glorieux combats	ambiguous	-
	juste violence	negative	A_NEG
	bienveillante opacité	negative	A_NEG
	judicieux arrêt	positive	A_POS
	louable souci	positive	A_POS
	aimable menace	negative	A_NEG
	glorieuse incertitude	negative	A_NEG
	vertueuse victime	positive	J_POS
4. *J_NEG noun,*	discrimination positive	negative	J_NEG
A_POS adjective	torture systématique	negative	J_NEG
	abus effectif	negative	J_NEG
	pédophilie fructueuse	negative	J_NEG
	assassinat spectaculaire	negative	J_NEG

	violation délibérée	negative	J_NEG
	débauche créative	negative	J_NEG
	délinquants raffinés	negative	J_NEG
	manipulation lucrative	negative	J_NEG
	propagande facile	negative	J_NEG
A_POS adjective, J_NEG noun	joli caprice	negative	J_NEG
	franche injustice	negative	J_NEG
	savoureuse débauche	negative	J_NEG
	tranquille insolence	negative	J_NEG
	délicieuse maladresse	negative	J_NEG
	doux abruti	negative	J_NEG
	bon esclave	ambiguous	-
	réjouissantes excentricités	ambiguous	-
	gentil assassin	negative	J_NEG
	extraordinaires inégalités	negative	J_NEG
5. *A_NEG adjective, J_POS noun*	mauvaise foi	negative	J_NEG
	mauvaise conscience	negative	J_NEG
	bouleversante sincérité	positive	J_POS
	maigres libertés	negative	A_NEG
	lourdes responsabilités	negative	A_NEG
	redoutable honneur	ambiguous	-
	hideuses vérités	negative	J_NEG
	fausses innocences	negative	J_NEG
	fausse bonhomie	negative	J_NEG
	difficile réconciliation	negative	J_NEG
J_POS noun, A_NEG adjective	pacifiques mous	negative	J_NEG
	dignité sombre	negative	J_NEG
	sensibilité tordue	negative	J_NEG
	sensibilité sourcilleuse	negative	J_NEG
	héros sanglant	negative	J_NEG
	intelligence inférieure	negative	J_NEG
	intelligence redoutable	negative	J_NEG
	honneur tardif	positive	J_POS
	compassion insupportable	negative	J_NEG
	pitié dangereuse	negative	J_NEG

6. *A_POS noun, J_NEG adjective*	enthousiasme abject	negative	J_NEG
	alliance passive	negative	J_NEG
	libération sentimentale	ambiguous	-
	coalition lâche	negative	J_NEG
	avantage illicite	negative	J_NEG
	ambition cynique	negative	J_NEG
	plaisir coupable	negative	J_NEG
	rire hystérique	negative	J_NEG
	lucidité sarcastique	negative	J_NEG
	élite dédaigneuse	negative	J_NEG
J_NEG adjective, A_POS noun	prétendus amis	negative	J_NEG
	fallacieuses protéctions	negative	J_NEG
	douteuse blague	negative	J_NEG
	égoïste insouciance	negative	J_NEG
	insolente admission	negative	J_NEG
	vulgaires ambitions	negative	J_NEG
	excessive longévité	negative	J_NEG
	honteux bénéfices	negative	J_NEG
	sanguinaires amis	negative	J_NEG
	ridicule désir	negative	J_NEG

Table B.3.: Samples of conflicting NPs (20 samples for each of the 6 most frequent combinations, cf. Table B.2).

D. Polarity Conflicts in Verb-frames

D.1. Lists of Verbs involved in Polarity Conflicts

Output of the script extract_Vconflicts.py.

Verbs in F-conflicts:
adopter aimer alimenter apprécier calmer conforter dominer favoriser inspirer manquer nourrir offrir perdre vouer épargner

Verbs in J-conflicts:
accepter accorder accueillir adorer aider aimer alimenter améliorer apprécier assurer bénéficier conforter corriger cultiver défendre désirer encourager favoriser gagner garantir mériter nourrir offrir oublier perdre permettre privilégier profiter promettre protéger prôner satisfaire soulager soutenir suggérer épargner

Verbs in A-conflicts:
accepter accomplir accorder accueillir accélérer acquérir admirer adopter adorer aider aimer alimenter animer applaudir approuver apprécier assurer autoriser cacher confier conforter conseiller convenir corriger cultiver célébrer dissimuler défendre désirer empêcher encourager espérer faciliter favoriser financer fêter gagner inaugurer inspirer inviter libérer manquer mériter nourrir offrir oublier pardonner perdre permettre privilégier profiter promettre promouvoir préconiser prôner rater ravir recommander remercier remporter renouveler respecter récupérer rétablir saluer sauver soigner souhaiter soulager soutenir suggérer supporter épargner

Intersection of A-conflicts + F-conflicts verbs with J-conflict verbs:
accepter accorder accueillir adorer aider aimer alimenter apprécier assurer conforter corriger cultiver défendre désirer encourager favoriser gagner mériter nourrir offrir oublier perdre permettre privilégier profiter promettre prôner soulager soutenir suggérer

D.2. Extracted Sentences with Potential Polarity Conflicts in Verbframes

D.2.1. Sentences with Shifter Verbs where no Conflict Occurred

- La violence, les menaces physiques maintes fois mises à exécution par des terroristes soucieux d'établir la charria, **ont empêché** *la population des provinces du Nord de se rendre massivement aux urnes.* (ex. 6.18)
- Il a raison, il **a empêché** *la guerre civile.* (ex. 6.19)
- À l'intérieur, on **oublie** instantanément *ce côté un peu sévère* pour être happé par la beauté et l'harmonie. (ex. 6.20)
- Il est temps d'accepter et de nous adapter à ce "nouveau monde" où l'accès à la culture **perd** *son caractère discriminatoire* et cesser de vouloir en faire une société virtuelle où tout un chacun se sent traqué. (ex. 6.21)

D.2.2. Sentences with Shifter Verbs where Polarity Conflicts Occurred

- On **oublie** trop aisément *les stupidités* de Présidents élus comme George W. (ex. 6.22)
- On **oublie** *les terroristes* et chasse les fraudeurs fiscaux et sociaux!

D.2.3. Sentences with Verbs of Different Semantic Categories

The following extracted sentences (both with and without polarity conflicts) are listed according to the categories defined in section 6.2.3.

Verb-Category 1

`A_NEG`, `F_NEG` and `J_NEG` always trigger conflicts, except if they are further labelled as `passive`.

- Ces clips torrides autant que ses flirts lesbiens dans les boîtes de strip-tease **alimentent** *les polémiques (A_NEG)* comme les phantasmes. (ex. 6.24)

- Il vient surtout faire aveu d'impuissance et de collusion avec le gouvernement en place dont la politique **nourrit** chaque jour *le flot des travailleurs pauvres (A_NEG passive) et des SDF.* (ex. 6.25; *no conflict*)
- Par ailleurs, Clinton a proclamée [sic] la volonté d'investir beaucoup plus que par le passé dans "l'humanité commune" pour aider les pays du tiers-monde à sortir de la misère et de la violence qui **nourrissent** *l'instabilité (A_NEG).*
- Les combats de rue, exécutions sommaires et rapts, qui, hier encore, n'épargnaient personne, **alimentent** toujours *les haines (F_NEG strong).*

Verb-Category 2

`A_NEG` and `F_NEG` do not trigger conflicts, except if they are further labelled as `strong`. `J_NEG` words and expressions always trigger conflicts.

- Contrairement à ce qu'il énonce, les entreprises investisseurs **aiment** bel et bien *les gouvernements corrompus (J_NEG strong) et facilement corruptibles (J_NEG strong)* qui peuvent livrer des marchés sans passer par des voies légales d'appels d'offres et de contrôle parlementaire. (ex. 6.26)
- Yves Leterme, le 22 décembre, le jour où le roi **a accepté** *sa démission (A_NEG).* (ex. 6.27; *no conflict*)
- L'incitation au racisme, la participation à un groupe qui **prône** *la discrimination (J_NEG)* resteront sanctionnés par le tribunal correctionnel.
- Ces visites font croire aux responsables gouvernementaux qu'ils font partie de la communauté internationale, alors qu'en fait, leur État **soutient** *le terrorisme (J_NEG strong).*

Verb-Category 3

Verbs are never involved in conflicts.

- Elle **soulage** *la misère (A_NEG strong) des vieux*, des types névrosés, ou trop gros, ou complexés, ou que leurs femmes refusent. (ex. 6.28)
- Il en **corrige** *les excès (J_NEG).* (ex. 6.29)
- Ce jardin secret qui, heureusement, les anime, c'est "l'aspirine" qui **soulage** *des maux (A_NEG) de la vie.*

- Vous qui êtes une personne réfléchie, vous vous dites que l'exécutif a pris la mesure des difficultés, et qu'il **corrige** *son erreur (A_NEG)* en passant du 15 décembre au 15 février 2010.

Verb-Category 4

`A_NEG`, `F_NEG` and `J_NEG` always trigger conflicts because the verb either intensifies the negative expression, or establishes a negative polarity effect or relation with regard to somebody or something.

- Mais elle veut lutter contre ce qui **favorise** *le désespoir (F_NEG strong).* (ex. 6.30 a.)
- Ils ont joué avec les Etats-Unis, ont été retirés de la liste des Etats qui **favorisent** *le terrorisme (J_NEG strong)* et ont obtenu de garder une partie de leurs actifs nucléaires. (ex. 6.30 b.)
- A demander que des preuves tangibles soient produites, autres que des noms et des dates sur un papier émanant des services de police d'un pays qui **cultive** *le mensonge (A_NEG)* et *la trahison (J_NEG strong)* à son plus haut niveau. (ex. 6.31 a.)
- "Ce film irresponsable **cultive** *la haine (F_NEG strong)* et *la défiance (A_NEG)* à l'égard de l'Occident", déclare président de la branche bavaroise de la CSU. (ex. 6.31 b.)
- Soit le prévenu **mérite** *une peine (A_NEG strong) sévère (A_NEG)*, comme l'a d'ailleurs demandé le parquet, soit il doit être acquitté, ce qu'a plaidé la défense. (ex.6.32)
- Les sauterelles les tourmentent, les hommes chercheront la mort sans la trouver, ils **désireront** *la mort (A_NEG)* et la mort les fuira. (ex.6.33)
- Deux des pick-up, le Nissan Navara (1 étoile sur 5) et le Isuzu D-Max Rodeo (2 étoiles) **offrent** *un score très médiocre (A_NEG)* au niveau de la protection des occupants adultes. (ex.6.34)
- Par contre, les patrons qui **profitent** *des travailleurs illégaux (J_NEG)* en prétendant ne pas connaître leur situation ne sont pas touchés par cette circulaire. (ex.6.35)
- Mais il continue (aussi) à soutenir de manière émotionnelle les organisations qui **promettent** *la vengeance (J_NEG)* en raison des souffrances et de l'humiliation.