Analysing Emotions in Social Media Coverage on Paris Terror Attacks: a Pilot Study

Cynthia Van Hee, Celine Verleye and Els Lefever
LT[3], Language and Translation Technology Team
Ghent University, Belgium
Groot-Brittanniëlaan 45, 9000 Ghent
Email: firstname.lastname@ugent.be

***Abstract*—Social media provide an increasingly used platform for crisis communication. Governments need to understand how publics consume and react to crisis information via social media. One option to do this is by applying emotion analysis. In this pilot study, we target the November 2015 terrorist attacks in Paris as a case study for emotion analysis and detection. We constructed a Dutch Facebook corpus manually annotated with i) Ekman's basic emotions and ii) irony use. The annotations reveal that *anger* is the most recurrent emotion, however the basic emotions do not cover all emotions in the dataset. The corpus also exhibits a fair number of ironic utterances, mostly expressing emotions like *disgust* and *anger*. The experimental results show that the detection of some emotions (e.g., *fear*) is challenging compared to others and that the classifier suffers from data sparseness.**

Keywords–Emotion detection; Social media; Natural language processing; Terrorism

I. Introduction

Social media have become primary communication tools for everyday conversations. More and more, they are also an important means of communication during crises [1], [2], allowing organisations and governments to inform the public, calm down anxiety and understand people's behaviour in such situations [3]. A recent example of this are the November 2015 Paris attacks, a series of coordinated terrorist attacks on 13 November 2015 in Paris by which 130 people lost their lives and many people were injured [4]. During the attacks, social media were extensively used by people looking for –or offering– shelter, and as a medium for spreading photos and information about missing people in the region [5]. Facebook activated the Paris Safety Check application allowing users to inform relatives about their safety and news channels provided up-to-date information and safety instructions via the platform. After the attacks, Facebook was also used by people to show their support for France and to react to the events.

As a result of their popularity, social networking sites constitute a rich source of information about the public opinion. Over the past decade, user-generated content has been investigated extensively in the field of sentiment and emotion analysis. Sentiment analysis involves machine learning techniques for determining the polarity of a text (i.e., positive or negative) [6], without taking into account specific emotions. The latter belongs to the field of emotion classification, which is a more fine-grained form of sentiment analysis that focuses on extracting emotions from text like joy, anger, and fear [7].

This paper describes a pilot study in which we apply machine learning techniques to unravel the emotions expressed on Facebook after the Paris attacks. To this end, we collected 483 Dutch Facebook reactions to news announcements covering the events. The data are retrieved from the Facebook pages of two Flemish news channels. The corpus is manually labeled for emotion-related categories including Ekman's basic six emotions [8]. Based on the annotations, we explore the feasibility of automatic emotion recognition and report our findings.

The remainder of the paper is structured as follows: in Section II, we give a brief overview of related work in the field of emotion detection. Section III describes the corpus and presents the annotation framework with some examples. Section IV elaborates on the emotion classification experiments. Finally, in Section V, we draw some conclusions and present prospects for future research.

II. Related Research

The past decade has seen an increased research interest in the field of sentiment and emotion analysis. In the framework of SemEval, the International Workshop on Semantic Evaluation[9], benchmark datasets have been made publicly available and several sentiment and emotion classification systems have been developed recently. Automatic emotion detection has been applied to different text genres including weblogs [10], emails [11], [12], news headlines [13], suicide notes [14], and tweets [3], [15]. Many systems for automatic emotion classification focus on the six basic emotions distinguished by Ekman [8], being *joy, fear, anger, disgust, sorrow* and *surprise*. Some studies, however, revealed more complicated emotions in text. For instance Plutchik [16] suggested eight bipolar primary emotions: *joy* versus *sadness*; *anger* versus *fear*; *trust* versus *disgust*; and *surprise* versus *anticipation*. Pestian et al. [17] distinguished sixteen emotion categories relevant to the domain of suicide notes. Finally, Yan & Turtle [7] composed a list of 28 emotions based on manual Twitter annotations.

Table I presents an overview of the state of the art in automatic emotion detection. Most of the work that is listed focuses on Twitter data and all but one (Yan & Turtle [7] describe a multiclass-based approach) conduct binary classification experiments per emotion category. In short, state-of-the-art emotion classifiers rely on machine learning algorithms such as LIBLINEAR, Naïve Bayes, Support Vector Machines, and k-Nearest Neighbors (k-NN). Often exploited features, i.e., information about text properties that may be relevant for emotion classification, include n-grams (i.e., sequences of n following words or characters), punctuation, Part-of-Speech

 ISBN: 978-1-61208-519-7

TABLE I. STATE-OF-THE-ART APPROACHES TO EMOTION DETECTION.

Reference	Corpus	# Emotion categories	Features	Results
Strapparava & Mihalcea [13]	1,25K news headlines	6	n-grams, sentiment lexicons, PMI, syntactic features	F= 0% – 32.78%
Wang et al. [18]	2.5M tweets	7	n-grams, sentiment/emotion lexicons, PoS tags	F= 13.90% – 72.10%
Roberts et al. [19]	7K tweets	7	n-grams, sentiment/emotion lexicons, PMI, punctuation, LDA	F= 60.80% – 74.00%
Mohammad et al. [20]	20K tweets	6	n-grams, sentiment/emotion lexicons	F= 18.70% – 62.40%
Yan & Turtle [7]	5.5K tweets	28	n-grams	F= 51.00% – 57.00%

tags, information from lexical resources such as WordNet-Affect [21], and topic information. The classification results vary among the emotion categories and often reveal that emotions like *joy* and *sadness* are more likely to be recognised than others [13], [18], [19].

III. CORPUS

To train and test the emotion detection classifiers, we collected a series of Facebook posts on the subject of the November 2015 Paris attacks. The corpus comprises 483 Dutch Facebook reactions to news announcements covering the attacks. The announcements date from 14 to 26 November 2015 and were posted on the Facebook page of two Flemish news channels being *Vlaamse Televisie Maatschappij (VTM)*, the main channel of commercial TV in Flanders and Brussels, and *Vlaamse Radio- en Televisieomroeporganisatie (VRT)*, the main channel of the Flemish public broadcaster. Table II presents some corpus examples covering direct reactions to the attacks (examples 1 and 2), as well as topics including house searches and safety measures implemented in Brussels (examples 3 and 4), the raid in which the alleged brain of the attacks was killed (examples 5 and 7), and communications about the threat level in the capital (example 6). After collecting the corpus, all posts were annotated for emotion and irony, the details of which are presented in the next paragraph.

A. Corpus Annotation

As mentioned earlier, the Facebook corpus was annotated for emotions and irony by trained linguists. The emotion annotation was based on Ekman's basic emotions [8]: *joy, fear, anger, disgust, sorrow* and *surprise*. We also included the label *Other* for ambiguous posts and posts expressing another emotion than one of the basic six, and *None* for posts exhibiting no emotions at all. The resulting set of manually labeled posts serves as the gold standard for the experiments. Table II presents an example for each emotion class with its corpus frequency. It should be noted that some posts received multiple labels. The sum of the second column values in the table thus reflects the total of emotion labels that were assigned for the entire corpus. Furthermore, all posts were annotated for the presence of (verbal) irony, the motivation for which is twofold: firstly we hypothesise that the subject will cause people to venture criticism, which is often 'softened' by using irony [22]. Indeed, tweets have proven rich in figurative language like irony [23], hence it wil be interesting to see if the same applies to the current dataset. Secondly, we want to investigate to what extent the presence of irony impacts the performance of the automatic emotion classifier. The next paragraph provides more details on this annotation with some ironic examples.

B. Annotation Analysis

Table II presents the different emotion classes that were annotated and provides a corpus example for each class. As described earlier, in addition to the basic emotions, we included *Other* as an annotation category. Interestingly, 278 instances were assigned this label, which means that in approximately 60% of the corpus the expressed emotion could not be matched to any of Ekman's basic six [8]. A closer inspection of the *Other* category reveals that many of these instances have a mocking or criticising tone and often express emotions like indignation and indifference (e.g., 'Yeah bla, bla, bla...', 'Guess I'm going to sleep. We'll see how it ends tomorrow *(...)*'). An analysis of the emotion distribution by gender reveals that women express more fear (14%) and sorrow (4%) as opposed to men (8% and 2%, respectively). Anger on the other hand, is the most frequent emotion expressed by men (19% vs. 16% by women). The observations seem to support the gender stereotyping of emotions [24], although further research on a larger dataset is needed.

With regards to the use of irony, we observe that approximately 20% of the corpus is labeled as ironic, which supports the findings of Ghosh et al. [23]. Here, we present two examples of ironic instances:

(1) Spijtig da fie *(sic.)* terrorist geen 60 ree waar je 50 mag! DAN zouden ze em wel hebben. **EN:** Too bad that the terrorist wasn't driving 60 where the speed limit is 50! THEN they would have caught him.

(2) Och al een geluk dat diene mens zoveel betaald *(sic.)* wordt om ons dit mee te delen! Had dat nooit zelf kunnen bedenken. **EN:** Good thing hat man is paid so much to communicate this to the public! Never could have come up with this myself.

Also, more ironic utterances are posted by men than by women (70% vs. 20%) –no author information was found for the remaining 10%. A closer look at the emotions expressed in ironic utterances reveals that the irony in the corpus often co-occurs with anger, disgust and *other* (Fig. 1).

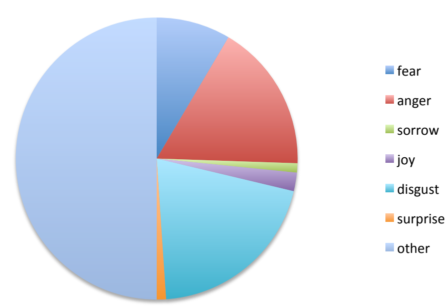

Figure 1. The distribution of ironic utterances according to the different emotion categories.

Generally, we see that the irony in these instances is mainly used for two purposes: i) expressing criticism towards the

TABLE II. CORPUS EXAMPLES.

Emotion category	# Instances	Relative freq.	Corpus example	Translation
anger	117	25.60%	1) Onnozelaars jullie maken ons van alles wijs !	1) You stupid people make us believe anything!
disgust	61	13.35%	2) Zoiets doet een beest nog niet	2) Not even an animal would do this ....
joy	62	13.57%	3) Knap van jou! jij hebt mijn verkiezingstem!	3) Excellent! You can count on my vote!
fear	52	11.38%	4) Pffff kids die bang zijn, wij zijn zenuwachtig, ... Niks om te lachen!	4) Pffff kids that are afraid, we parents that are nervous, ... Nothing to laugh about!
sorrow	10	2.19%	5) Ik treur voor zijn ouders...	5) I feel sorry for his parents...
surprise	7	1.53%	6) Snap er niks van..... Eerst zochten ze 1 terrorist en het was niveau 4, nu zoeken ze 2 terroristen en nu is het niveau 3 ?????	6) Do not get it ..... First, they were looking for one terrorist and the level was four, now they are looking for two terrorists and now the level is 3 ?????
other	278	60.83%	7) Woorden maar weinig initiatief...	7) Words, but little initiative...

Belgian government and police, and ii) lightening the subject by using irony as a form of humour, for instance by mocking with the alleged brain of the attacks. Examples of the latter tend to be more ludic than the former. However, both uses of irony share the purpose of expressing criticism towards some entity, which supports the hypothesis that irony is often used to express criticism in a less face-threatening way [22].

IV. EXPERIMENTS

We evaluated the feasibility of emotion classification in Facebook data by means of a series of binary classification experiments. For the experiments, we only considered posts in which at least one emotion category was identified by the annotators, which resulted in an experimental corpus containing 457 instances. For each emotion category –including *Other*– a binary experiment was run to predict whether the emotion is present (classification label “1”) in an instance or not (label “0”). This resulted in seven binary experiments with one emotion category as the positive class, whereas the remaining emotion categories represent the negative class. Instances that were annotated with more than one emotion category (e.g., expressing both anger and fear), are subject to detection by the different corresponding classifiers.

As the classification algorithm we used LIBSVM [25] with linear kernel. As evaluation measures, we report (tenfold cross-validated) (1) precision, (2) recall and (3) F_1-score for the positive class, calculated as follows:

$$Precision = \frac{Number\ of\ correctly\ predicted\ labels}{Total\ number\ of\ predicted\ labels} \quad (1)$$

$$Recall = \frac{Number\ of\ correctly\ predicted\ labels}{Total\ number\ of\ gold\ standard\ labels} \quad (2)$$

$$F-score = \frac{2(Precision\ *\ Recall)}{Precision\ +\ Recall} \quad (3)$$

In addition, we report accuracy figures, which simply divide the number of true predictions (both positive and negative class) by the total number of instances.

As a preprocessing step, all posts were tokenised using the LeTs Preprocess Toolkit [26]. For each classifier, the following features were exploited:

- **Bags-of-words:** token unigrams, bigrams and trigrams.
- **Sentiment features** based on two existing sentiment lexicons for Dutch [27], [28]:
    - the number of positive, negative and neutral tokens in the instance;
    - the overall polarity, i.e., the sum of the values of the identified polarity words in the instance.

Table III presents the experimental results for all binary classifiers by means of accuracy, precision, recall and F_1-score. As we approach the automatic emotion classification task as a detection taks, we only considered the positive class labels (i.e., the instances containing the emotion in question) for calculating precision and recall. In contrast, the accuracy results are measured on the complete data set (i.e., all positive and negative instances).

TABLE III. EXPERIMENTAL RESULTS PER EMOTION CLASSIFIER.

Emotion category	Accuracy	Precision	Recall	F_1-score
Anger	72.21%	42.86%	25.64%	32.09%
Joy	89.28%	76.00%	30.65%	43.68%
Fear	86.00%	25.00%	11.54%	15.79%
Disgust	89.06%	66.67%	36.07%	46.81%
Surprise	98.47%	-	-	-
Sorrow	97.81%	-	-	-
Other	71.55%	75.87%	78.06%	**76.95%**

Not considering *Other*, we see that the system performance is highest for the category *Disgust* (F_1= 46.81%), followed by *Joy* (F_1= 43.68%). The category *Other* scoring best would suggest that, albeit ambiguous, the category encompasses instances that share a number of characteristics. Another explanation for the good result would be the high relative frequency of the emotion class in the corpus compared to the other categories. The *Surprise* and *Sorrow* classifiers consistently predict the negative class, resulting in an F_1-score of zero and an accuracy equal to the proportion of negative class instances. Presumably, there are insufficient training examples in the corpus for both categories, which causes the system to fail to build a good model for recognising new instances of these classes.

A qualitative analysis of the systems' output reveals that many misclassifications could be the result of the systems exploiting only lexical information. For the *Joy* category for instance, we see a fair number of false negatives that contain negative sentiment words while expressing a positive sentiment overall (e.g., ‘It's a shame I can only press the like button once!’). Inversely, false positives often include sentences with positive words while expressing an overall negative emotion (e.g., ‘The government should guarantee a good policy *(...)*’). With respect to the category *Anger*, we see that many false positives contain flooded punctuation (e.g., ‘Good job guys!!!!’), which would indicate that the system considers heavy punctuation as an indication of anger. An explanation for the poor performance of the category *Fear* would be that such emotion expressions (e.g., ‘What will happen now?’,

 ISBN: 978-1-61208-519-7

'Should we keep the kids at home tomorrow?') are much less lexicalised than expressions of anger, for instance.

A more general conclusion that can be drawn from the analysis is that many instances are ambiguous, i.e., they exhibit more than one emotion category. We see that instances containing only one emotion category are more often correctly classified than instances expressing multiple emotions. An analysis of the annotated categories shows that *Joy, Disgust* and *Other* are often the only emotion category that was identified (in 65% of the cases), whereas *Fear* and *Anger* were more often used in combination with other emotion categories (only in 37% of the cases it was the only expressed emotion). This is also reflected in Table III. We also see a fair number of ironic utterances among the wrongly classified instances, which would suggest that irony indeed affects the classification performance (cf. Section III-A).

When comparing the results to the state of the art, we see that generally, the classifiers perform less well than other systems that are trained on much larger corpora (Table I). Nevertheless, this pilot study provides valuable insights into the emotions expressed in the aftermath of a series of terrorist attacks. The main conclusions are the following:

1) Ekman's basic emotions [8] are insufficient to describe all emotions in the corpus. Expanding the list would reduce the number of ambiguous annotations and scale down the *Other* class.
2) The emotion classifiers mainly rely on lexical clues, which are often insufficient to determine the correct emotion class.
3) Many instances contained multiple emotion categories. Since a binary classification task forces the system to choose one label, it would be interesting to see whether a multiclass approach works better.
4) The results for sparse emotion categories (e.g., *Surprise*) are very low, indicating that a strong correlation exists between the occurrence of a class and the system's performance.

V. Conclusions and Future Work

In this paper, we analysed the emotions expressed online in the aftermath of the November 2015 Paris attacks. The analysis reveals that anger is one of the most salient emotions. Governments should bear this in mind when communicating with the public. Since *Other* remains the largest emotion category in the corpus, we suggest to expand the list of basic emotions. The results of the binary classification experiments show that emotion classification is not a trivial task and that the system's performance clearly suffers from data sparseness. If we discard the category *Other*, the best results are achieved for the emotion categories *Disgust* and *Joy*. This would suggest that these categories are more explicit or highly lexicalised when compared to the others. We see an inverse correlation between classification performance and the proportion of ambiguous instances (i.e., instances expressing multiple emotions) in the corpus. For instance, the proportion of ambiguous instances for the *Joy* category is 32% whereas this is 62% for *Anger*. F_1-scores for the corresponding classifiers are 43.68% and 32.09%, respectively. Another interesting observation is the good performance for the category *Other*, which was assigned to tweets that are ambiguous or that express another emotion than one of the basic six. When looking at the use of irony, we see that many ironic utterances in the corpus co-occur with the emotions anger, disgust and *other*. A closer look into the latter revealed that many of these instances contain emotions like indignation, and indifference (cf. Section III-B).

This paper presents a pilot study to emotion detection in Dutch crisis communication. To be able to generalise our findings, more experiments are needed on a larger dataset, which will be the main focus in future work. Additionally, we aim to enhance the performance of our classifiers by adding more complex features including topic models, Linguistic Inquiry and Word Count (LIWC) features and syntactic information. Another interesting direction for future work is automatic irony recognition. Since the classifier exploits sentiment lexicon features, its performance is affected by ironic utterances that contain positive sentiment words while actually conveying a negative sentiment.

References

[1] A. Schwarz, "How publics use social media to respond to blame games in crisis communication: The Love Parade tragedy in Duisburg 2010," Public Relations Review, vol. 38, no. 3, 2012, pp. 430–437, ISSN: 0363-8111.

[2] Y. Jin, A. Pang, and G. T. Cameron, "Integrated crisis mapping: Towards a publics-based, emotion-driven conceptualization in crisis communication," Sphera Publica, vol. 7, no. 1, 2007, pp. 81–96, ISSN: 1180-9210.

[3] B.-K. H. Vo and N. Collier, "Twitter emotion analysis in earthquake situations," International Journal of Computational Linguistics and Applications, vol. 4, no. 1, 2013, pp. 159–173.

[4] "November 2015 Paris attacks," 2015, URL: https://en.wikipedia.org/wiki/November_2015_Paris_attacks/ [accessed: 2016-10-03].

[5] "Quel est le rôle des réseaux sociaux dans des événements comme les attentats de Paris?" 2015, URL: http://www.la-croix.com/Actualite/France/Quel-est-le-role-des-reseaux-sociaux-dans-des-evenements-comme-les-attentats-de-Paris-2015-11-15-1380592/ [accessed: 2016-10-03].

[6] B. Pang and L. Lee, "Opinion Mining and Sentiment Analysis," Foundations and Trends in Information Retrieval, vol. 2, no. 1-2, 2008, pp. 1–135, ISSN: 1554-0669.

[7] J. S. Y. Liew and H. R. Turtle, "Exploring Fine-Grained Emotion Detection in Tweets," in Proceedings of the NAACL Student Research Workshop, June 13–15, 2016, San Diego, California, USA. Association for Computational Linguistics, Jun. 2016, pp. 73–80.

[8] P. Ekman, "An Argument for Basic Emotions," Cognition and Emotion, vol. 6, no. 3, 1992, pp. 169–200.

[9] "SemEval-2016 : Semantic Evaluation Exercises," 2016, URL: http://alt.qcri.org/semeval2016/ [accessed: 2016-10-03].

[10] R. Mihalcea and H. Liu, "A corpus-based approach to finding happiness," in Proceedings of AAAI-CAAW-06, the Spring Symposia on Computational Approaches to Analyzing Weblogs, 2006.

[11] H. Liu, H. Lieberman, and T. Selker, "A Model of Textual Affect Sensing Using Real-world Knowledge," in Proceedings of the 8^{th} International Conference on Intelligent User Interfaces (IUI), January 12–15, 2003, Miami, Florida, USA. ACM, Jan. 2003, pp. 125–132, ISBN: 1-58113-586-6.

[12] S. M. Mohammad and T. Yang, "Tracking Sentiment in Mail: How Genders Differ on Emotional Axes," in Proceedings of the 2^{nd} Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA), June 24, 2011, Portland, Oregon, USA. Association for Computational Linguistics, Jun. 2011, pp. 70–79.

[13] C. Strapparava and R. Mihalcea, "Learning to Identify Emotions in Text," in Proceedings of the 2008 ACM Symposium on Applied Computing (SAC), March 16–20, Fortaleza, Ceará, Brazil. ACM, Mar. 2008, pp. 1556–1560, ISBN: 978-1-59593-753-7.

[14] B. Desmet and V. Hoste, "Emotion detection in suicide notes," Expert Systems with Applications, vol. 40, no. 16, 2013, pp. 6351–6358, ISSN: 0957-4174.

[15] S. M. Mohammad, “#Emotional Tweets,” in *SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012), June 7–8, Montreal, Canada. Association for Computational Linguistics, Jun. 2012, pp. 246–255.

[16] R. Plutchik, The Emotions: Facts, theories, and a new model. Random House, 1962.

[17] J. P. Pestian, P. Matykiewicz, M. Linn-Gust, B. South, O. Uzuner, J. Wiebe, K. B. Cohen, J. Hurdle, and C. Brew, “Sentiment analysis of suicide notes: A shared task,” Biomedical informatics insights, vol. 5, no. Suppl. 1, 2012, p. 3.

[18] W. Wang, L. Chen, K. Thirunarayan, and A. P. Sheth, “Harnessing Twitter “Big Data” for Automatic Emotion Identification,” in 2012 International Conference on Privacy, Security, Risk and Trust (PASSAT 2012), and 2012 International Confernece on Social Computing (SocialCom 2012), September 3–5, 2012, Amsterdam, Netherlands. IEEE Computer Society, Sep. 2012, pp. 587–592, ISBN: 978-0-7695-4848-7.

[19] K. Roberts, M. A. Roach, J. Johnson, J. Guthrie, and S. M. Harabagiu, “EmpaTweet: Annotating and Detecting Emotions on Twitter,” in Proceedings of the 8^{th} International Conference on Language Resources and Evaluation (LREC), May 21–27, 2012, Istanbul, Turkey. European Language Resources Association (ELRA), May 2012, pp. 3806–3813, ISBN: 978-2-9517408-7-7.

[20] S. M. Mohammad and S. Kiritchenko, “Using hashtags to capture fine emotion categories from tweets,” Computational Intelligence, vol. 31, no. 2, 2015, pp. 301–326, ISSN: 0824-7935.

[21] C. Strapparava and A. Valitutti, “WordNet-Affect: An affective extension of WordNet,” in Proceedings of the 4^{th} International Conference on Language Resources and Evaluation (LREC), May 26–28, Lisbon, Portugal. ELRA, May 2004, pp. 1083–1086.

[22] P. Brown and S. C. Levinson, Politeness: Some Universals in Language Usage. Cambridge University Press, Feb. 1987, ISBN: 9780521313551.

[23] A. Ghosh and T. Veale, “Fracking Sarcasm using Neural Network,” in Proceedings of the 7^{th} Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA), June 16, 2016, San Diego, California, USA. Association for Computational Linguistics, Jun. 2016, pp. 161–169.

[24] E. A. Plant, J. S. Hyde, D. Keltner, and P. G. Devine, “The Gender Stereotyping of Emotions,” Psychology of Women Quarterly, vol. 24, no. 1, 2000, pp. 81–92, ISSN: 1471-6402.

[25] C.-C. Chang and C.-J. Lin, “LIBSVM: A Library for Support Vector Machines,” ACM Transactions on Intelligent Systems and Technology (TIST), vol. 2, no. 3, 2011, pp. 27:1–27:27, ISSN: 2157-6904.

[26] M. Van de Kauter, G. Coorman, E. Lefever, B. Desmet, L. Macken, and V. Hoste, “LeTs preprocess: the multilingual LT3 linguistic preprocessing toolkit,” Computational Linguistics in the Netherlands Journal, vol. 3, 2013, pp. 103–120, ISSN: 2211-4009.

[27] T. De Smedt and W. Daelemans, ““Vreselijk mooi!” (Terribly Beautiful!): A Subjectivity Lexicon for Dutch Adjectives,” in Proceedings of the 8^{th} International Conference on Language Resources and Evaluation (LREC), May 21–27, 2012, Istanbul, Turkey, 2012, pp. 3568–3572.

[28] V. Jijkoun and K. Hofmann, “Generating a Non-English Subjectivity Lexicon: Relations That Matter,” in Proceedings of the 12^{th} Conference of the European Chapter of the Association for Computational Linguistics (EACL), March 30–April 3, Athens, Greece, Mar. 2009, pp. 398–405.