# *AAL-2009-2-049, ALIAS*
# *D4.5*
# *Module for retrieval of opinionated content*

| | |
|---|---|
| **Due Date of Deliverable** | 2012-12-31 |
| **Actual Submission Date** | 2013-03-05 |
| **Workpackage:** | 4 |
| **Dissemination Level:** | Public |
| **Nature:** | Report |
| **Approval Status:** | Final |
| **Version:** | v1.0 |
| **Total Number of Pages:** | 18 |
| **Filename:** | D4.5-EURECOM-OpinionatedContent.pdf |
| **Keyword list:** | Sentiment analysis, opinion mining, events, news |

**Abstract**
Detecting what people feel or think about particular topics has been the subject of several applications and research studies. In particular, people's opinions are an important key to enhance event experience and direct the decision upon attending future events. There is a need to explore and aggregate in real-time the online sentiment that could be retrieved from many sources such as news articles and microblogs. However, automated opinion discovery is a complex task that requires a high-level text analysis, particularly in micro-posts characterized by short informal text. In this report, we provide an overview about relevant opinion sources for events and we discuss a variety of issues related to sentiment classification.

## History

| Version | Date | Reason | RevisedBy |
|---|---|---|---|
| 0.1 | 2012-12-15 | created [EURECOM] | Houda Khrouf |
| 0.2 | 2013-02-22 | updated [EURECOM] | Raphael Troncy |
| 0.3 | 2013-02-25 | first review [TUM-GSING] | Wolfram Schneider |
| 0.4 | 2013-03-04 | second review [COGNESYS] | Karsten Lenkitsch |
| 0.5 | 2013-03-04 | addressing comments [EURECOM] | Houda Khrouf |
| 1.0 | 2013-03-05 | finalized [EURECOM] | Raphael Troncy |

## Authors


| Partner | Name | Phone / Fax / Email |
|---|---|---|
| EURECOM | Raphaël Troncy | Tel: ++33 4 9300 8242<br>Fax: ++33 4 9300 8200<br>Email: Raphael.Troncy@eurecom.fr |
| EURECOM | Houda Khrouf | Email: Houda.Khrouf@eurecom.fr |

## Table of Contents

## 1 Summary

Detecting what people feel or think about particular topics has been the subject of several applications and research studies. In particular, people's opinions are an important key to enhance event experience and direct the decision upon attending future events. There is a need to explore and aggregate in real-time the online sentiment that could be retrieved from many sources such as news articles and microblogs. However, automated opinion discovery is a complex task that requires a high-level text analysis, particularly in microposts characterized by short informal text. In this report, we provide an overview about relevant opinion sources for events and we discuss a variety of issues related to sentiment classification.

## 2 Abbreviations and Acronyms

| | |
|---|---|
| API | Application Programming Interface |
| CRF | Conditional Random Fields |
| LODE | An ontology for Linking Open Descriptions of Events |
| OWL | The Web Ontology Language is a knowledge representation language based on description logics. It has an RDF syntax and in its dialect OWL-Full (one the three flavors of OWL) includes the RDF/S semantics |
| NER | Named Entity Recognition |
| NLP | Natural Language Processing |
| POS | Part-Of-Speech |
| RDF | The Resource Description Framework is a knowledge representation language based on a triple model, and serves as foundation for other semantic web languages such as RDFS or OWL |
| SIOC | Semantically-Interlinked Online Communities |
| sLDA | supervised Latent Dirichlet Allocation |
| SVM | Support Vector Machine |

## 3 Introduction

The ALIAS robot has to provide a range of functionalities that promote the social inclusion of elderly and engage them in different type of events. Retrieving informational and opinionated content that stimulate the user's knowledge of events, e.g by reading articles from online newspapers and blog post, is therefore of particular interest in this deliverable. The opinionated content is actually often considered as an important cue that describes events, and it can also be exploited for recommendation purposes. However, tracking people's sentiments about events as expressed in online forums, newspapers or social networks is a challenging task. It is difficult to retrieve relevant content, and to extract the opinions expressed in the sentences in order to summarize the overall sentiment of people. These tasks rely on high-level text analysis involving subjectivity detection and polarity classification to determine whether a document, a sentence or an entity is associated with a positive, a negative or a neutral opinion. Despite the intuition that events would be associated with sentiment fluctuations, there has been very little attention to deeply study the nature of this event-sentiment relationship.

In this document, we provide an overview of some definitions important in the study of sentiment analysis (Chapter 4). Then, we describe relevant opinion sources for events and the techniques used to retrieve relevant content (Chapter 5). Once data collected, to determine the overall sentiment about events, we detail some existing techniques that could be applied based on machine learning and opinion words dictionaries (Chapter 6). Finally, we give our conclusion and outline the future work (Chapter 7).

# 4 Background

Mining and retrieval of opinionated content has recently gained attention in research studies. Discovering what people think about entities, events and their properties is an important task that promotes the decision making. Opinions are usually subjective expressions that describe people's sentiments, appraisals or feelings towards a variety of topics [1]. A plethora of online services such as blogs, Internet forums and social networks host an ever-increasing amount of users' feedbacks and views on products and many other entities. However, it is difficult for a human reader to find relevant sources, extract related sentences with opinions, read them, summarize them, and organize them into usable forms [1]. Thus, there is a need for an automatic opinion discovery from single or heterogeneous sources, which is the focus of sentiment analysis, also called opinion mining.

The sentiment analysis is based on Natural Language Processing (NLP) techniques trying to infer people's sentiments using their language expressions. Sentiment analysis concentrates on attitudes, whereas traditional text mining focuses on the analysis of facts [2]. It is in general cast as text classification problem where two sub-topics have been mostly addressed in research studies. The first topic aims to classify an opinionated document as expressing positive or negative sentiment, known as document-level sentiment classification. The second topic classifies individual sentences as subjective or objective that means whether a sentence expresses an opinion or not, known as subjectivity classification. Then, the subjective sentence is classified as expressing positive, negative or neutral opinion, known as sentence-level sentiment analysis. On the other hand, some research studies [3, 4] have also investigated another problem called the feature-based sentiment analysis. The goal of this analysis is to detect the targets on which opinions have been expressed in a sentence. This is important when we want to produce detailed analysis about an entity and its related features. For instance, in a product review sentence, a set of product features and their related opinions can be identified and classified as positive or negative sentiment. For example, an event is associated to a set of features such as artists, location and attendees, and one sentence may contain opinions related to one or more of these features.

In order to provide a clear description of opinionated document, we report the models of object and opinionated document defined in [1].

- **Model of an object:** An object o is represented with a finite set of features, $F = f_1, f_2, ..., f_n$, which includes the object itself as a special feature. Each feature $f_i \in F$ can be expressed with any one of a finite set of words or phrases $W_i = w_{i1}, w_{i2}, ..., w_{im}$, which are synonyms of the feature, or indicated by any one of a finite set of feature indicators $I_i = i_{i1}, i_{i2}, ..., i_{iq}$ of the feature.
- **Model of an opinionated document:** A general opinionated document d contains

opinions on a set of objects $o_1, o_2, ..., o_q$ from a set of opinion holders $h_1, h_2, ..., h_p$. The opinions on each object $o_j$ are expressed on a subset $F_j$ of features of $o_j$.

# 5 Data Sources for Opinion Retrieval

The sentiment analysis involves in building a system to collect and examine opinions about the product made in blog posts, comments, reviews or tweets [2]. Nowadays, opinions are in general shared over a large variety of web services. For example, the opinions about events are described in different newspaper texts, micro-posts and other important documents where cues such as occurrence time, location and agents are often specified. However, retrieving relevant informations from different sources and detecting emotional states from texts are challenging tasks. In this chapter, we focus on the first task and we describe the sources and techniques used to build relevant corpus, from which emotional states can be detected. Examples of opinions sources are news sites and micro-bogging services, for which we detail how to retrieve data and link it with relevant topics.

## 5.1 News Sites

There exist a variety of websites that report the latest news about social events or other related entities. Examples of online news sites are Guardian[1], BBC[2] and New York Times[3] providing open data available for re-use on the Internet. The news sites provide interesting facts that describe an event environment and could potentially hold subjective insights. For instance, the *Radiohead Concert*[4] given on the *6th of November 2012* in *New Zealand* has been reported in *UnderTheRadar.Co.Nz*, a music news website. The article[5] contains the following sentences: *"Radiohead played an epic show last night at Auckland's Vector Arena. A show that attacked all the senses, face-meltingly loud and visually stunning"*. A strong positive opinion is clearly expressed and highlighted through a set of high intensity positive adjectives. Broadly speaking, the news websites are popular source of opinions, but they typically contain much less clearly stated opinions compared with other sources, which makes harder the opinion mining [5]. This is because the news articles are in general narrative-oriented reports detailing facts about what has been happened. A few attention is paid to express explicit feedbacks, which are in turn reported by few authors (e.g the author of the article) leading to a low opinion diversity.

We use the web APIs provided by the sites to retrieve relevant articles on a particular event or feature. We exploit the explicit relationships that exist between EventMedia and the news web sites. For example, Guardian has recently used Linked Data in its open platform, so that the identifier supplied by a third party could be utilized to locate Guardian

[1] http://www.guardian.co.uk/
[2] http://www.bbc.co.uk/
[3] www.nytimes.com/
[4] http://www.last.fm/event/3213752
[5] http://www.undertheradar.co.nz/news/5977/Live-Photos-Radiohead---Vector-Arena-Auckland.utr

content. For instance, MusicBrainz IDs has been added to Guardian articles, and consequently the lookup by these identifiers via Guardian API is made possible. It is worth to note that MusicBrainz is an open music encyclopedia that collects, and makes available to the public, music metadata. There exist another example called Echo Nest, a project devoted to aggregating, indexing, using, and sharing vast troves of music data. Its API enables the lookup of highly relevant news articles also using MusicBrainz IDs. Thanks to our reconciliation framework explained in the Deliverable D4.3 [6], we have been able to link nearly 40 000 agents in EventMedia with MusicBrainz dataset. In consequence, we simply use the identifiers when it is possible to retrieve news articles about agents involved in particular event. Similarly, we also exploit the connections between EventMedia and DBpedia established during the reconciliation to retrieve the articles of New York Times about agents linked to DBpedia. By this approach, we highlight the importance of Linked Data to greatly simplify information retrieval. In addition, we mainly focus on building a corpus of news articles using agents instead of other entities such as events or locations. In fact, most of articles with reference to events contain also a reference to involved agents. Then, to discover opinions about each entity (i.e: event), a feature-based analysis can be performed (see section 6).

### 5.2 Micro-blogging Services

Microblogging services have recently been a very popular communication tool generating millions of daily messages, some of which concerns events of different types. Because of a free format of messages and an easy accessibility of microblogging platforms, Internet users tend to shift from traditional communication tools (such as traditional blogs or mailing lists) to microblogging services [7]. Social networks such as Twitter, Facebook or Google Plus provide now services for microblogging, where authors discuss a variety of topics about their life and share opinions. In addition, it is found that the number of elderly who spend time on these sites has recently increased, so they provide interesting sources in the context of ALIAS project. For instance, according to recently conducted study [8] by Pew Internet (an internet research company), the number of social networking users aged 65 or more increased from 13% in 2009 to 33% in 2011. Given the ease with which one could create and maintain online relationships, elderly are more motivated to share information and to create a community of friends on social networks.

Among these services, Twitter is the most popular microblog website enabling users to post 140 characters status update messages (tweets). In particular, participants are more and more engaged in real-time discussion on Twitter during real-world events. Twitter users share information about upcoming events, the events the users are attending and events being broadcasted. In [9], a study shows that an important event can be expected to trigger more informational tweeting, in contrast to personal communication. In Table 5.1, we report examples of typical opinionated posts from Twitter about some events.

The microblogging services have been a subject of many research studies in sentiment

| Event | Tweet | Polarity |
|---|---|---|
| Coldplay and Jay-Z Concert on 31th of December 2012 | What a better way to welcome the new year then attend a Jay-Z and Coldplay concert at the Barclays Center... | Positive |
| Winter Festival in Newcastle | It's looking fantastic at the rehearsal for the Winter Carnival | Positive |
| 29th Chaos Communication Congress in Hamburg | A really enjoyable #29c3 talk by Natalie Silvanovich on her reversing Tamagotchis | Positive |

Table 5.1: Examples of Twitter posts with expressed opinions about particular events

analysis. In particular, detecting opinions about particular events comprises two tasks. More precisely, as data shared in microblogging services are often noisy and unstructured, detecting sentiments should first cope with retrieving relevant content about particular event, and secondly mine and summarize opinions. In this section, we focus on the first task that has to connect structured data with noisy content.

Many microblogging services provide an API to allow developers to collect data programmatically. For example, the Twitter API allows making HTTP calls to their servers with parameters to get filtered data. Twitter API allows filtering by location, keywords and author. Retrieving relevant content about known events can be ensured using a set of terms such as event title, agents' names and location. Then, an alignment can be performed to filter relevant tweets and associate them with particular event. In [10], the authors present an automated approach to align tweets with the events they illustrate. They use two machine learning techniques: proximity-based clustering and Naive Bayes classification. In [11], our alignment approach is based on named entity extraction (NEE) to tackle the noisy nature of tweets. The evaluation shows a precision of about 60% and recall of about 47%.

There also exists a plethora of tools to perform sentiment analysis of Twitter messages. First, the user enters a search term to retrieve related tweets, and then he gets back all the positive, negative and neutral classification of those tweets, along with some graphics such as pie charts or graphs. Typical basic online tools are Twitter Sentiment[6], TweetFeel[7] and Twitrratr[8].

[6] http://www.sentiment140.com/
[7] http://www.tweetfeel.com/
[8] http://www.twitrratr.com/

# 6 Sentiment Classification

Sentiment classification is the task to determine the polarity of a document as expressing a negative, positive or neutral opinion. There are a number of challenging aspects of this task. Opinions in natural language are very often expressed in subtle and complex ways, presenting challenges which may not be easily addressed by simple text categorization approaches such as n-gram or keyword identification approaches. Recognizing the semantic impact of words or phrases is a challenging task in itself, but in many cases the overarching sentiment of a text is not the same as that of decontextualized snippets. Negative reviews may contain many apparently positive phrases even while maintaining a strongly negative tone, and the opposite is also common. Most of existing techniques are based on supervised learning, and also on unsupervised methods and nature language processing (NLP) techniques. In particular, a much attention has been paid recently to the sentiment analysis from micro-posts because of their noisy nature. The micro-posts tend to be less grammatical than longer posts and typically contain a frequent use of abbreviations, emoticons and sarcasm words, which are difficult for a machine to detect. In this section, we describe some related work in the field of sentiment classification and more generally, in the field of text mining.

## 6.1 Supervised Learning

The machine (ML) approach applicable to sentiment analysis mostly belongs to supervised classification in general and text classification techniques in particular. A number of supervised learning techniques have been applied to sentiment classification such as Naive Bayes, maximum entropy (ME), and support vector machines (SVM) [12, 13]. For standard machine learning, a set of texts annotated for polarity by human coders are used to train an algorithm to detect features that associate with positive, negative, and neutral categories. The trained algorithm can then look for the same features in new texts to predict their polarity. In general, the machine learning approaches make use of syntactic and/or linguistic features. Examples of features are: (i) the n-grams which means the frequency of occurrence of all n consecutive words; (ii) the part of speech tags (POS), particularly the adjectives which are considered as important indicators of subjectivity and opinion; (iii) the opinion bearing words which are commonly used to express an opinion such as *good*, *beautiful*, *wonderful* used to express positive sentiment; (iv) the negation words which often change the sentiment orientation, for example, considering the difference between the phrases ”good” and ”not good”. For example, the work in [13] aimed to classify the tweets on the basis of positive or negative sentiment using n-gram and POS features, and trained on instances which had been annotated according to positive and negative emoticons. The evaluation underlines the high accuracy achieved with machine learning algorithms (Naive Bayes, Maximum Entropy, and SVM) when trained with emoticon data. In [2], the authors provide an insight about the difference between some ML methods employed to classify sentiments in movie and product reviews. It has been shown that the discriminative models such as support vector machines (SVM) outperform generative models. Nevertheless, despite their performance, the supervised methods are highly sensitive to the domain from which the training data are extracted. The reason behind is that words and even language constructs used in different domains for expressing opinions can be substantially different [2].

With regards to our interest, few are the research studies that address sentiment analysis about events. In [14], an approach has been proposed to mine future events from news articles, and identify related sentiment for recommendation purpose. The assumption behind is that positive feelings are related to entertainment event such as festival, concert and sport event, while the negative feelings are induced by events such as accident, crime, slow traffic, and poor weather. The approach is based on document level sentiment classification where the identified events from a news article are assigned the news article sentiment. Two classification methods have been applied, namely the supervised Latent Dirichlet Allocation (sLDA as a generative graphical model) [15] and the Support Vector Machine (SVM as a discriminative classifier). The results show that the Linear SVM and the Gaussian SVM perform better than sLDA and the human-based classifiers. One limitation of this approach is the assumption that the sentiment polarity is reflected by the event category, and there is no deep analysis of the article content. Other studies have also analyzed event sentiment from Twitter in order to detect whether there is a correlation between events and increases in sentiment strength. Most of them are based on unsupervised methods (see section 6.2), since it is challenging to provide sufficient training data from minute texts.

### 6.2 Unsupervised Techniques

There exists many unsupervised techniques for sentiment classification. Most of them are lexicon-based approaches that use dictionaries of opinion bearing words that can be created manually or automatically using a set of seed words and an online dictionary. Examples of dictionaries are the Harvard General Inquirer[1], the WordNet-Affect[2] based on [16] and SentiWordNet[3]. For a given text, all words are extracted and annotated with their polarity value using the dictionary scores. The polarity scores are in turn aggregated into a single score for the text. A typical lexicon-based example is the Twitrratr[4] website that uses a manually-made list of positive and negative keywords to classify the tweets as positive, negative or neutral. Figure 6.1 shows an example of opinionated tweets in Twitrratr about the music band "Coldplay". In conjunction with lexicon-based approach,

[1] http://www.wjh.harvard.edu/~inquirer/

[2] http://wndomains.fbk.eu/wnaffect.html

[3] http://sentiwordnet.isti.cnr.it/

[4] http://www.twitrratr.com/

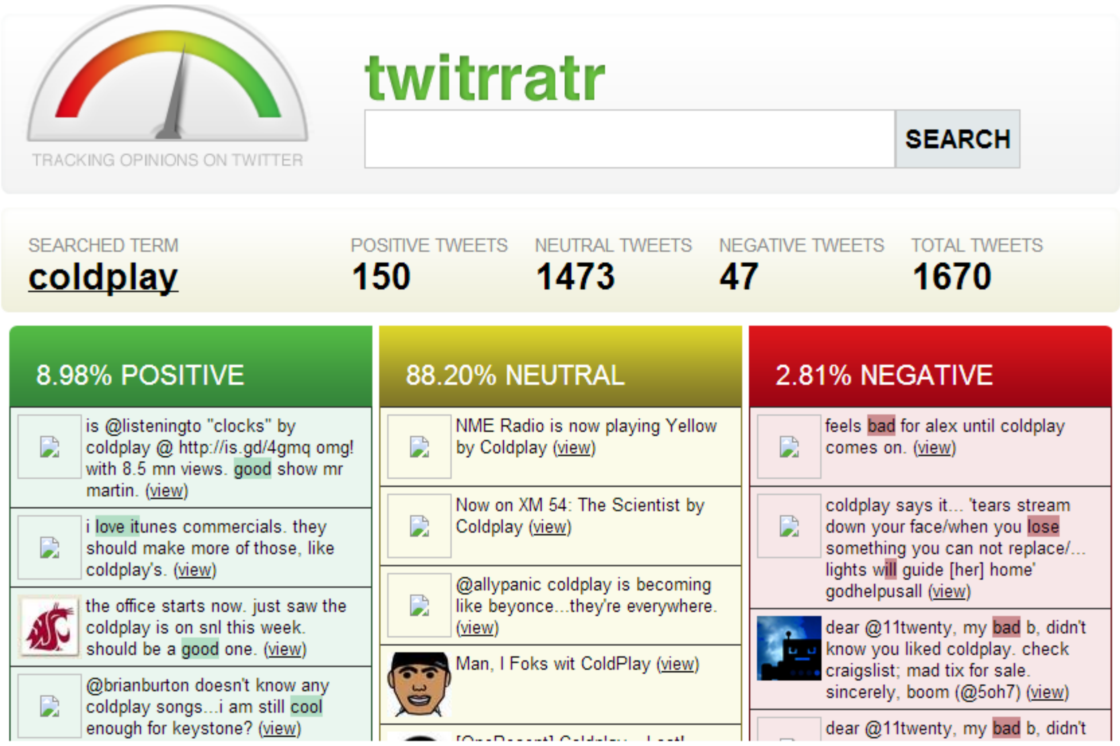


Figure 6.1: Opinionated tweets in Twitrratr about the music band "Coldplay"

a linguistic analysis of the grammatical structure is also exploited to predict the polarity of a given text. For instance, linguistic algorithms may attempt to identify context, negations, superlatives, and idioms as part of the polarity prediction process [17].

In [18], Thelwall et al. use the SentiStrength algorithm to analyze sentiment from Twitter posts about popular events. SentiStrength [19] is designed for short informal English text with abbreviations and slang. It exploits a list of sentiment words with associated strength ranging from 2 to 5. The algorithm integrates many functionalities such as a machine learning optimization of strength values, a spelling correction, a negation words list and an emoticon list, the repeated letters and the exclamation mark which reflect sentiment intensification. The results highlight that popular events are normally associated with increases in negative sentiment strength, and that peaks of interest in events have stronger positive sentiment than the time before the peak. Similarly, another approach [20] analyze Twitter posts to explore the relations between popular events and fluctuations of public moods. It measures the sentiment of each tweet using an extended version of the Profile of Mood States (POMS), a lexicon-based approach assessing six individual dimensions of mood, namely *Tension*, *Depression*, *Anger*, *Vigour*, *Fatigue*, and *Confusion*. The results highlight that social, political, cultural and economic events are correlated with significant, even if delayed fluctuations of public mood levels along a range of different mood dimensions.

### 6.3 Feature-based Analysis

The feature-based analysis is a fine-grained analysis that mines opinions in context and determines the features which have been opinionated. In fact, most of current systems do not consider the target when classifying the sentiment, and they only operate at the document level or at the sentence level. For example, the tweets in Twittrar are classified as positive or negative, but there is no correlation between the search keyword and the sentiment detected. For instance, in the tweet: "*I feel bad for alex until coldplay comes on*", there is no negative opinion about the search keyword "Coldplay", although classified as negative tweet according to Twittrar website. To overcome this problem, some of existing techniques have been used such as the Named Entity Extraction (NER), the Conditional Random Fields (CRF), and the rule-based methods which consider the grammatical dependency relations detected when parsing the text. For example, Maynard et al. [3] first identify the relevant entity and then look for linguistic relations that associate opinions to this entity (for instance, an entity may be linked to a verb expressing like or dislike as its direct object). In their system, they propose to integrate an event recognition module that identify events to be used as possible targets for the opinions. They adopt a pattern-based method to identify domain-specific events using Entity Recognition and linguistic relations. For example, events can be expressed by text elements such as event-referring nouns ("crisis", "accident"), noun phrases ("economic growth"), etc. Although their approach is interesting, there is still much work in progress to be evaluated. Moreover, to the best of our knowledge, there is no work that mine opinions about the elements that can be part of an event environment, and thus of the global sentiment. A research area that can, of course, be more investigated in order to provide an accurate and fine-grained analysis of sentiments related to events.

# 7 Conclusion

This report describes the different tasks required to mine and classify opinionated content with particular focus on events. We give an overview about the existing techniques employed in information retrieval, natural language processing and computational linguistics. From the technical point of view, to enhance the performance of sentiment classification, it is found that different types of features, classification algorithms and linguistic methods are combined in an efficient way in order to overcome their individual drawbacks and benefit from each other's assets. In particular, an attention has been paid to event sentiment classification in social media which highlight the correlation between events and increases in sentiment strength. In the future, we plan to experiment with existing techniques to deeply analyze the associative relationship between events and sentiments in EventMedia.

## Bibliography


[1] B. Liu, "Sentiment analysis and subjectivity," in *Handbook of Natural Language Processing, Second Edition. Taylor and Francis Group, Boca*, 2010.

[2] G.Vinodhini and RM.Chandrasekaran, "Sentiment analysis and opinion mining: A survey," *International Journal of advanced Research in Computer Science and Software Engineering*, vol. 2, pp. 282–292, June 2012.

[3] D. Maynard, K. Bontcheva, and D. Rout, "Challenges in developing opinion mining tools for social media," in *@NLP can u tag #usergeneratedcontent?! Workshop at LREC 2012*, (Istanbul, Turkey), 2012.

[4] M. Hu and B. Liu, "Mining and summarizing customer reviews," in *10$^{th}$ ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, (Seattle, WA, USA), 2004.

[5] A. Balahur and R. Steinberger, "Rethinking sentiment analysis in the news: from theory to practice and back," in *1$^{st}$ Workshop On Opinion Mining And Sentiment Analysis*, (Sevilla, Spain), 2009.

[6] R. Troncy, H. Khrouf, G. Atemezing, A. Fialho, and L. Hardman, "Module for knowledge enrichment of event descriptions." ALIAS, Deliverable 4.3, 2011.

[7] A. Pak and P. Paroubek, "Twitter as a corpus for sentiment analysis and opinion mining," in *7$^{th}$ International Conference on Language Resources and Evaluation (LREC'10)*, (Valletta, Malta), 2010.

[8] K. Zickuhr and M. Madden, "Older adults and internet use." http://www.pewinternet.org/Reports/2012/Older-adults-and-internet-use.aspx, June 2012.

[9] A. L. Hughes and L. Palen, "Twitter adoption and use in mass convergence and emergency events," vol. 6, no. 3, pp. 248–260, 2009.

[10] M. Rowe and M. Stankovic, "Aligning Tweets with Events: Automation via Semantics," *Semantic Web Journal*, 2011.

[11] H. Khrouf, G. Atemezing, G. Rizzo, R. Troncy, and T. Steiner, "Aggregating social media for enhancing conference experiences," in *1$^{st}$ International Workshop on Real-Time Analysis and Mining of Social Streams*, (Dublin, IRLANDE), 2012.

[12] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up? sentiment classification using machine learning techniques," in *Conference on Empirical Methods in Natural Language Processing*, pp. 79–86, 2002.

[13] A. Go, R. Bhayani, and L. Huang, “Twitter sentiment classfication using distant supervision,” tech. rep., Stanford, 2009.

[14] S.-S. Ho, M. Lieberman, P. Wang, and H. Samet, “Mining future spatiotemporal events and their sentiment from online news articles for location-aware recommendation system,” in *1$^{st}$ ACM SIGSPATIAL International Workshop on Mobile Geographic Information Systems*, (California, USA), 2012.

[15] D. M. Blei and J. D. McAuliffe, “Supervised topic models,” in *Advances in Neural Information Processing Systems (NIPS)*, MIT Press, 2007.

[16] M. Taboada, M. Tofiloski, J. Brooke, K. Voll, and M. Stede, “Lexicon-based methods for sentiment analysis,” *Computational Linguistics*, vol. 37, no. 2, pp. 267–307, 2011.

[17] T. Wilson, J. Wiebe, and P. Hoffmann, “Recognizing contextual polarity: An exploration of features for phrase-level sentiment analysis,” *Computational Linguistics*, pp. 399–433, 2009.

[18] M. Thelwall, K. Buckley, and G. Paltoglou, “Sentiment in twitter events,” *Journal of the American Society for Information Science and Technology*, vol. 62, no. 2, pp. 406–418, 2010.

[19] M. Thelwall, K. Buckley, G. Paltoglou, D. Cai, and A. Kappas, “Sentiment in short strength detection informal text,” *Journal of the American Society for Information Science and Technology*, vol. 61, dec 2010.

[20] J. Bollen, H. Mao, and A. Pepe, “Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena,” in *5$^{th}$ International AAAI Conference on Weblogs and Social Media*, (Barcelona, Spain), 2011.

[21] J. Blitzer, M. Dredze, and F. Pereira, “Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification,” in *Annual Meeting Association for Computational Linguistics*, (Prague, Czech Republic), 2007.