

SentiProfiler: Creating Comparable Visual Profiles of Sentimental Content in Texts

Tuomo Kakkonen
School of Computing
University of Eastern Finland
tuomo.kakkonen@uef.fi

Gordana Galić Kakkonen
Department of Croatian Language and
Literature
University of Split, Croatia
ggalic@ffst.hr

Abstract

We introduce the concept of sentiment profiles, representations of emotional content in texts and the SentiProfiler system for creating and visualizing such profiles. We also demonstrate the practical applicability of the system in literary research by describing its use in analyzing novels in the Gothic fiction genre. Our results indicate that the system is able to support literary research by providing valuable insights into the emotional content of Gothic novels.

1 Introduction

In recent years, non-topical text analysis has become an active area of research. In contrast to “traditional” text analytics in which the aim is to extract and process the facts and concepts, non-topical text analysis aims at characterizing the attitudes, opinions and feelings present in a text.

Automatic detection and classification of sentiments has several potential areas of application. Hence, it has become an established field of *natural language processing* (NLP) research and a rising number of papers and articles dedicated to the analysis of the affective content of texts is being published each year. Feelings and sentiments are often represented within texts in complex and subtle ways, which makes their automatic detection an interesting research challenge.

Much of the work in *sentiment analysis* (SA) is concentrated on business applications, in particular the analysis of user-generated online content, such as customer reviews and tweets, in order to tap into what is sometimes referred to as the “Wisdom of the Crowds.” This work departs

from the general trends in SA research in more ways than one.

First, we apply SA to a field in which it has not, to the best of our knowledge, been applied before, namely the analysis of novels for literary research. Second, rather than concentrating on determining the polarity (negative vs. positive) of the analyzed document, we aim at providing more detailed classification affective content of the target texts. Third, although visualization of SA results has attracted some interest from the opinion mining and SA community, research contributions that particularly address the visual comparison of SA results are scarce.

We introduce a system called SentiProfiler that generates visual representations of affective content of texts and uses techniques to outline the differences and similarities between the pairs of texts under scrutiny. The design of the tool is centered on the idea of *sentiment profiles* (SPs), hierarchical representations of affective contents of the input documents. The SPs and the visual graphs representing them that the system generates are in effect summaries of the sentiment content of the input documents. The system makes use of various NLP and visualization resources and tools in tandem with Semantic Web technologies.

In addition calculating frequencies of sentiment-bearing words, a scoring measure is used to determine the prevalence of sentiments in the target texts. The hierarchical organization of the sentiment classes enables the system to aggregate these scores in order to gauge the prevalence of specific groups of sentiments. Our visualization technique uses vertex colors to denote the differences between the SPs that are being compared. This allows for quick browsing and detection of differences between the emotional content of the target documents.

We demonstrate the practical applicability of the SentiProfiler system by considering the ways in which it can support literary research by allowing visual comparisons of pairs of SPs created from works of a literary genre rich in dark and gloomy topics (and hence negative emotions), namely Gothic literature. The fact that Gothic novels are divided into two distinct sub-genres of terror and horror allows us to experiment with the comparison functions of the system.

The paper is organized as follows. In Section 2, we outline the background of this research by shortly summarizing works on visualization of SA results and describing the tools and technologies on top of which SentiProfiler has been built. Section 3 introduces the architecture of the SentiProfiler system and explains the functioning of each of its main components. The practical applicability of the system is demonstrated and discussed in Section 4. Section 5 concludes by summarizing the paper and considering directions for future work.

2 Introduction

2.1 Related work

Most work on SA visualization has dealt with summarizing analysis results for collections of documents. Often, basic visualization methods, such as bar charts (Liu et al., 2005) and temporal graphs (Fukuhara et al., 2007) are utilized. The Pulse system by Gamon et al. (2005) generates treemap visualizations that display topic clusters and their associated sentiments. The size of the boxes indicates the number of sentences in the topic cluster, and the color denotes the average sentiment of the sentences belonging to that topic. Chen et al. (2006) generated multiple visualizations (such as decision trees and term variation networks) in order to enable the analysis of conflicting opinions.

One of the rare works that discusses visual comparison of SA results (Gregory et al., 2006) introduced a system that combines lexicon lookup-based SA with a visualization engine. The paper described an experiment with the well-known Hu and Liu (2004) customer review dataset.

The most relevant of the recent work on visualization of SA results is presented in Wu et al. (2010). They introduced OpinionSeer, a system that visualizes hotel customer feedbacks that is based on a visualization-centric analysis technique that considers uncertainty for modeling

and analyzing customer opinions. They also suggested a type of visual representation that conveys customer opinions by augmenting scatterplots and radial visualization.

The only research work we are aware of that has applied SA to literary research is reported in (Taboda et al., 2006). In contrast to our work, rather than analyzing novels, Taboda et al. used SA techniques to extract information on the reputation of six early 20th century authors based on writings concerning them.

2.2 Tools and technologies applied

The SentiProfiler system uses WordNet-Affect (Strapparava and Valitutti, 2004) as the source for emotion-bearing words. WordNet-Affect is a linguistic resource for the lexical representation of affective knowledge. It was developed on the basis of WordNet (Miller, 1995) through the selection and labeling of the *synsets* (the WordNet technical term for semantically equivalent words) representing affective concepts. WordNet-Affect defines a hierarchy of emotions, in which the items are referred to as emotional categories. Each emotional category is linked with a set of WordNet synsets that contain the words that are connected with the emotional category. The WordNet-Affect hierarchy contains four main categories of emotions: negative, positive, ambiguous and neutral.

The WordNet-Affect hierarchy and the corresponding synsets from WordNet are represented in SentiProfiler as an ontology that is automatically generated from the WordNet-Affect hierarchy and WordNet synset definitions. The ontology is created and accessed by using the Jena framework (<http://jena.sourceforge.net/>). Jena is a well-known and stable Java framework for building Semantic Web applications that provides, among other things, an API for manipulating Semantic Web resources in the Resource Description Framework (RDF) format.

SentiProfiler makes use of several other well-known freely available Java tools and libraries. The MIT Java WordNet Interface (JWI) (<http://projects.csail.mit.edu/jwi/>) is an API for interfacing with WordNet. JWI is used by SentiProfiler for retrieving the relevant synsets from a local copy of the WordNet dictionaries.

GATE (Cunningham et al., 2002) is a widely used and flexible framework for developing text analysis systems. It is commonly applied in the NLP research community. SentiProfiler utilizes GATE in ontology-based tagging of sentiment-bearing words.

JUNG (the Java Universal Network/Graph Framework) (<http://jung.sourceforge.net/>) is a library for the modeling, analysis, and visualization of graphs. It provides an extensible set of graph operations, visualization methods and layouts. SentiProfiler uses JUNG for the visualization of sentiment profiles.

3 SentiProfiler

3.1 Introduction

The SentiProfiler system consists of three main components: ontology and ontology factory, sentiment analyzer and SP visualizer. Figure 1 outlines the system architecture.

As shown in Figure 1, the ontology that describes the hierarchy of the emotions to be detected is generated automatically from WN-Affect and WordNet data (see Section 3.2). The sentiment analyzer component (Section 3.3) consists of a GATE pipeline and a set of Java classes that analyze the tags assigned by GATE and creates the SPs of the input documents. The visualizer component (Section 3.3) is responsible for displaying SPs in a format that supports easy visual comparison.

3.2 Ontology of sentiments

The automatic ontology creation process from WordNet-Affect and WordNet data works as follows. First, the XML file containing the relevant part of the WordNet-Affect hierarchy is converted and stored in a graph data structure. Next, an ontology representing the sentiment hierarchy is created that contains the WordNet-Affect sentiment categories as classes. We refer to these as

sentiment classes. The words from the relevant WordNet synsets are used as the individuals that instantiate each sentiment class. Finally, the ontology is written in an RDF/XML file that can be browsed and edited with any RDF-aware ontology editor. This allows automatically generated ontologies to be manually inspected and modified to the specific needs particular analysis tasks.

The JitterOnto ontology of negative sentiments that was generated for the experiments on Gothic literature described in Section 4 consisted of the 147 classes under the *negative-emotion* branch of the WordNet-Affect hierarchy. The maximum depth from the root class *negative-emotion* to a leaf sentiment class was five. This was the case, for example, with the sentiment class *negative-emotion/general-dislike/negative-fear/negative-unconcern/heartlessness/cruelty*.

There are 823 individuals in the ontology, which means that the sentiment classes are described by a total of 823 nouns, verbs and adjectives. Hence, each class has on average of 5.6 instantiations. For instance, the *cruelty* class mentioned above is instantiated by the words “cruel,” “cruelly,” “cruelty,” “mercilessness,” “pitilessness,” “ruthlessness” and “unkind.”

Figure 2 shows an extract from the branch of the ontology that has to do with the set of emotions that are classified as *negative-fear*. The other seven classes that immediately follow the root class *negative-emotion* are as follows: *anxiety*, *daze*, *despair*, *general-dislike*, *humility*, *ingratitude* and *sadness*.

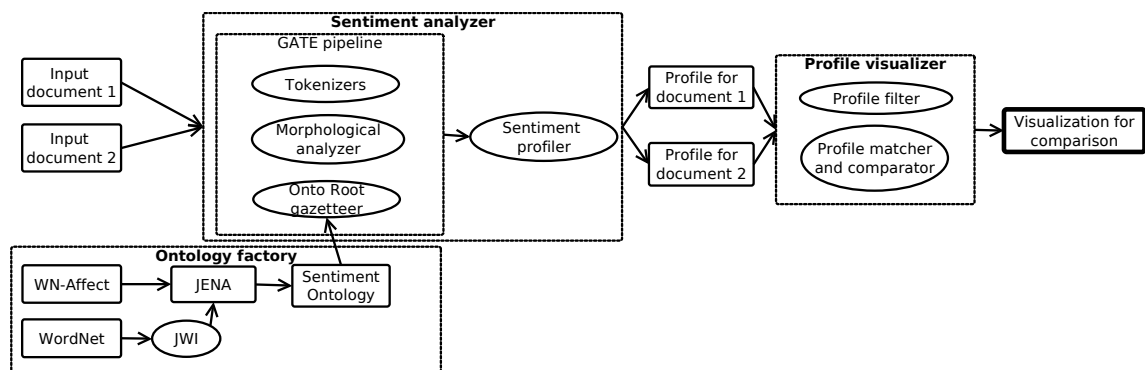


Figure 1. Architecture of the SentiProfiler system.




Figure 2. Extract from JitterOnto showing the two levels of more specific sentiment classes under the parent class negative-fear. The figure was created with the Protégé editor (Stanford Center for Biomedical Informatics Research, 2011).

3.3 Analysis of affective content

Sentiment analysis in SentiProfiler consists of creating a SP of an input document. A SP is a hierarchy of sentiment classes that contains all the classes of the source ontology that occurred in the document, and the classes that are part of a path from such a class to the root (*negative-emotion* in case of the ontology used for the experiments in this paper). That is, those sentiment classes that did not appear in the input document or have any children that appeared in the document are not included in the SP.

For instance, let us have a document in which the only negative sentiment classes that appear are *chill* and *timidity*. Following the paths from *negative-emotion* in Figure 2 to *chill* and *timidity* gives us the SP of the document.

The creation of SPs consist of three phases: detection of sentiment-bearing words, relating each such word with the relevant sentiment class and, finally, constructing the hierarchy that describes the SP of a document.

SA is performed in SentiProfiler with a GATE pipeline that consists of three basic ANNIE components (sentence splitter, word tokenizer and POS tagger), GATE morphological analyzer and an ontology-based tagging tool. Onto Root Gazetteer (Damljanovic et al., 2008) is a GATE processing resource that dynamically constructs a gazetteer from an ontology and creates, in combination with other GATE components, on-

tology-based annotations on the given content. In the SentiProfiler GATE pipeline Onto Root Gazetteer marks up in the input document the words that match with an individual found in the ontology. The relevant sentiment class names are used as the tags in the output.

Next, a graph presentation of the sentiment class hierarchy (i.e. the sentiment profile) is created for the input document in which each graph vertex is associated with the number of times a word relating to the relevant sentiment class appears in the document. In addition to the frequency counts we define a score that measures the prevalence of each sentiment class. The *sentiment class scores* (SCMs) measure the relative frequency with which a specific sentiment s_i appears in a document. SCM score is defined as follows:

$$SCM1_i = \frac{|S_i|}{|words|}$$

where $|S_i|$ is the number of times a word instantiating the sentiment class s_i appears in the document and $|S_i|$ is the total number of word tokens in the document. For instance, let us have a document with 1000 word tokens. Three of the sentiment-bearing words belong to the class s_1 . The SCM score for the sentiment class s_1 is 0.15.

An *aggregate SCM score* (ASCM) is defined for all the non-leaf sentiment classes. It is calculated as the sum of the SCM scores of all the sentiment classes that succeed the current sentiment classes in the hierarchy plus the SCM score of

the current sentiment class. This score provides a way of comparing whole branches of the SP rather than one single sentiment class at a time. Figure 3 illustrates the concept.




Figure 3. An example of SCM and ASCM scores. The scores that are not inside parentheses are SCMs (multiplied by 10 for easier interpretation). The figures inside the parentheses are aggregate SCM values.

3.4 Visualization

In the visualization of an SP, each sentiment class is represented as a vertex that indicates, in addition to the name of the sentiment class, any combination (depending on the system configuration) of frequency, SCM and ASCM values. User can also observe the occurrences of the sentiment-bearing words pertaining to specific sentiment category along with the context in which the word appeared.

The matching algorithm compares the profiles in order to find the sentiment classes that are present in only one of the profiles. We refer to these as *additional sentiment classes*. It also calculates the differences between the scores of those classes that occur in both of the profiles. The sentiment classes that receive a higher score are referred to as *higher score sentiment classes*. As a result of the matching process, the vertices representing additional and higher score sentiment classes are denoted by specific user-modifiable colors in the visualization.

4 Experiments with Gothic Novels

We evaluated the practical applicability of the SentiProfiler system by analyzing the SPs of Gothic novels. Such novels consist of stories of “terror and suspense, usually set in a gloomy old castle or monastery” (Baldick, 2004). The Gothic literary genre is further divided into works of terror and horror. Many explanations of the distinction between the two subgenres have been put forward in the literary research community (for example, Botting, 1996). In essence, the distinction between terror and horror can be summarized as the intensity and the type of emotions

they depict and evoke in the fictional character as well as in the reader him/herself. Although the works of both subgenres of the Gothic novels contain emotions such as anxiety, fear and gloom, in terror these emotions are more connected to a threat, real or perceived, rather than actual events of cruelty and violence.

What makes this genre of novels so suitable for testing the profile comparison capabilities of SentiProfiler is, first, that they can be expected to contain a relatively high amount of (negative) emotional content. Second, the fact that there are two subgenres of Gothic novels provides a way of testing the practical use of our system in comparing SPs. If SentiProfiler is able, in addition to creating SP visualizations of Gothic novels, to distinguish differences in the types of emotions present in the SPs generated for samples of the two Gothic novel subgenres it provides evidence that the system can be used as a practical tool for supporting literature research.

Due to the nature of the target texts, we concentrated our analysis on negative emotions. What we were interested in observing, in particular, were differences in the SPs that support the distinction between the two subgenres of Gothic literature as it is understood in the theory of literature. What we were expecting to be able to recognize is that horror novels contain emotions that can be described as a sort of an “aftershock”, a display of disgust that appears after a horrendous event has occurred. Terror, in contrast, raises anxiety and timidity caused by the fear of something terrible happening in the near future. In a sense, terror is of a more “psychological” nature than horror. Varma (1966) puts it succinctly: the difference between the two subgenres “is the difference between awful apprehension and sickening realization: between the smell of death and stumbling against a corpse.”

Figure 4 illustrates an extract from the comparison of a horror and terror novel. The two novels used in the comparison were Matthew Lewis’s (1796) “The Monk: A romance”. It is considered as one of the prime examples of novels in Gothic horror. Ann Radcliffe’s (1794) “The Mysteries of Udolpho” enjoys a similar status in the Gothic terror genre¹.

¹ All the novels for the research reported in this paper were obtained from the Project Gutenberg web site at <http://www.gutenberg.org>.

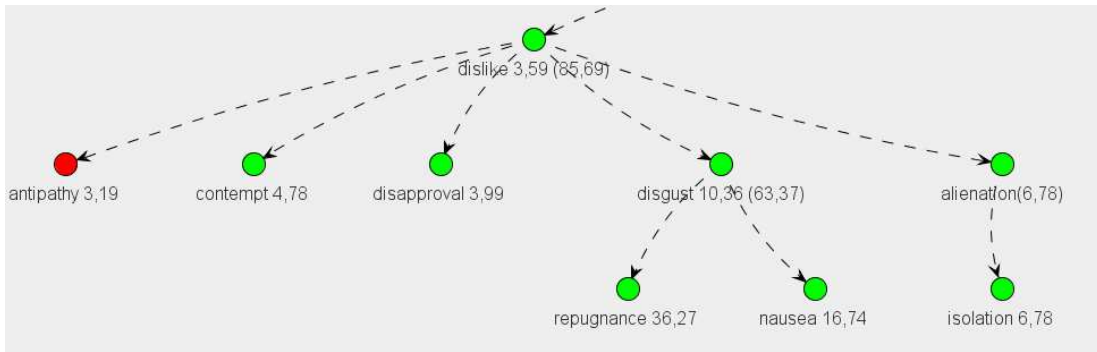


Figure 4. An extract from a sentiment profile comparison of the horror novel (*The Monk*) with the terror novel (*Udolpho*) from the perspective of the horror novel. Green vertices denote higher score sentiment classes. The red vertex (*antipathy*) indicates an additional sentiment class.

The SP extract in Figure 4 indicates, as expected under the definitions of the two subgenres, that the novel representative of the horror subgenre receives higher SCM scores for the sentiment classes pertaining to, for example, disgust and nausea. The same was observed, for example, for the sentiment class cruelty (not shown in the figure).

Figure 5 shows an extract from the same pair-wise comparison that was depicted in Figure 4, but in this case from the perspective of the terror novel.

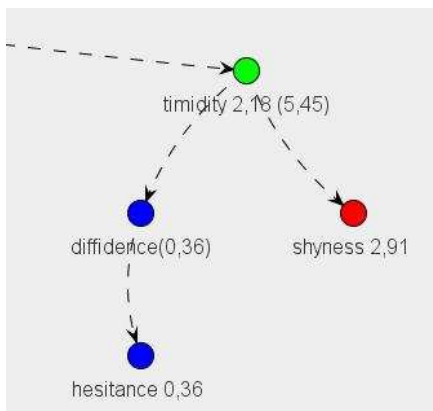


Figure 5. An extract from a comparison of the terror novel with the horror novel showing additional sentiment classes *diffidence* and *hesitance* and higher score class *timidity*.

As visualized in Figure 5, as expected, the terror novel more frequently contained senti-

ments that have to do with emotions that can be described as “less intense” and less connected with cruelty and the shock caused by acts of violence. In addition to the sentiment classes illustrated in Figure 5, this was the case, for instance, with the sentiment classes *impatience* and *depression*.

The color coding of vertices and the ability to zoom in and out (see Figure 6 for an example) enable the literary scholar to easily locate of the branches in the SP that show a significant number of differences between the novels explored. Moreover, the dialog that shows the contexts in which the words linked with a specific sentiment class occurred helps the comparative literature researcher to make more detailed analyses of, for instance, style and discourse.

Table 1 summarizes some of the observed differences between the SPs created based on the terror and horror novel discussed above, and the following two canonical works in the horror and terror subgenres: *Frankenstein*; or, *The Modern Prometheus* (Shelley, 1818) and *the Castle of Otranto* (Walpole, 1764).

The table does not give an exhaustive list of all the differences, but rather concentrates on those classes that support the notion of dividing the Gothic subgenres based on the “severity” of emotions.




Figure 6. Zoomed-out view of a sentiment profile comparison that demonstrates how the system can be used for locating interesting branches in a sentiment profile. In this sample case, attention is drawn to the large branch in the middle, as it contains several blue vertices (denoting additional sentiment classes) and a mixture of red and green vertices. The user can zoom in to the interesting section of the visualization by turning the mouse wheel.

Novel pair		Sentiment classes			
		Higher score		Additional	
T	H	T	H	T	H
Udolpho	Frankenstein	timidity, shyness, anxiousness	horror, disgust, hate, lividity	hesitance, solicitude, insecurity	-
Otranto	Monk	sorrow, depression, anxiety	horror, disgust, repugnance	-	nausea, lividity
Udolpho	Monk	timidity, anxiety, impatience	horror, disgust, repugnance, cruelty	hesitance, diffidence	nausea, lividity
Otranto	Frankenstein	hopelessness, impatience	horror, disgust, repugnance	-	nausea, trepidation

Table 1: Pair-wise comparison of the two horror and terror novels. The two columns under the heading “higher score” lists examples of the sentiment classes that received higher score in the terror (T) and the horror (H) novel, respectively. The columns under the heading “additional” contain examples of sentiment classes that were present only either in the terror or the horror novel.

The results reported in Table 1 indicate that differences can indeed be observed in the relative frequency and presence of certain sentiment classes between representatives of the two subgenres of Gothic literature. Sentiment classes such as timidity, anxiety and shyness were more frequent in the terror novels than they were in the representatives of the horror subgenre. Horror and disgust were more frequent sentiment classes in horror novels in all the four pair-wise comparisons. Nausea was among sentiments that were present only in the horror novels included in the comparison.

It is also interesting to note that the terror novel *Castle of Otranto* had additional categories aggravation and wrath (not shown in Table 1) that were not present in either of the horror novels. This observation did not seem to support the expected distinction between terror and horror novels. However, further analysis of relevant research literature (for instance, Hume (1969)) revealed that, while *Otranto* is often considered as a terror novel, it is somewhat of a borderline case between the two subgenres. Being able to capture such an ambiguity gives additional support to the practical applicability of SentiProfiler.

5 Conclusion

We introduced the concept of sentiment profiles (SPs) and the SentiProfiler system for creating easily comparable SPs created from pairs of documents. The system is built on the basis of various well-known NLP and Semantic Web technologies and tools. We demonstrated the use of the system by describing how it can support research in comparative literature. The visual comparisons allow the literary scholar to gain

insights into the target texts that would be difficult, if not impossible, to obtain with the traditional “pen and paper” research methods that are typically used in the field.

The preliminary experiments we reported in Section 4 indicated that the tool can provide interesting insights for literary researchers. SentiProfiler was able to detect differences between the sentiments present in example novels of Gothic terror and horror subgenres. Moreover, many of the differences observed by the tool supported the literature theoretical distinction between these two subgenres.

Our planned future work includes applying the SentiProfiler tool to conduct a comprehensive study of a larger set of Gothic novels in order to verify whether the commonly accepted definition of the difference between the emotional content present in horror and terror subgenres holds true.

There are various ways in which SentiProfiler could be improved and extended. First, we plan to apply the system to a study of Gothic novels. Since we are dealing with books from the eighteenth and nineteenth centuries, extending the JitterOnto ontology of negative emotions with words that are not in use in modern English would presumably increase the accuracy of the system in that particular area of application.

The system has many potential uses beyond literary research. Negative emotions play a role in anti-social behavior. One of the future applications of SentiProfiler and the JittersOnto ontology used in the experiment reported in this paper includes automatic detection and prediction of potential acts of extreme anti-social behavior (such as school shootings) based on messages posted online. It is important to note that, despite the focus of this paper, SentiProfiler is designed

and implemented in a way that allows any of the WordNet-Affect sentiment hierarchy branches, and in fact any ontologically represented class hierarchy, to be used. Hence, there is no reason why the tool and the method could not be applied in more general case of SA, rather than focusing on negative emotional content.

References

- Chris Baldick. 2004. *The Concise Dictionary of Literary Terms (2nd edition)*. Oxford University Press, Oxford, UK and New York, USA.
- Fred Botting. 1996. *Gothic (The New Critical Idiom)*. Routledge, New York, USA.
- Chaomei Chen, Fidelia Ibekwe-SanJuan, Eric SanJuan and Chris Weaver. 2006. Visual Analysis of Conflicting Opinions. *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*, Baltimore, Maryland, USA.
- Hamish Cunningham, Diana Maynard, Kalina Bontcheva and Valentin Tablan. 2002. GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications. *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistic*, Philadelphia, Pennsylvania, USA.
- Danica Damljanovic, Valentin Tablan and Kalina Bontcheva. 2008 A Text-based Query Interface to OWL Ontologies. *Proceedings of the 6th Language Resources and Evaluation Conference*, Marrakech, Morocco.
- Tomohiro Fukuhara, Hiroshi Nakagawa and Toyooki Nishida. 2007. Understanding Sentiment of People from News Articles: Temporal Sentiment Analysis of Social Events. *Proceedings of the International Conference on Weblogs and Social Media*, Boulder, Colorado, USA.
- Michael Gamon, Anthony Aue, Simon Corston-Oliver and Eric Ringger. 2005. Pulse: Mining Customer Opinions from Free Text. *Proceedings of the 6th International Symposium on Intelligent Data Analysis*, Madrid, Spain.
- Michelle L. Gregory, Nancy A. Chinchor, Paul Whitney, Richard Carter, Elizabeth Hetzler and Alan Turner. 2006. User-Directed Sentiment Analysis: Visualizing the Affective Content of Documents. *Proceedings of the Workshop on Sentiment and Subjectivity in Text*, Sydney, Australia.
- Minqing Hu and Bing Liu. 2004. Mining Opinion Features in Customer Reviews. *Proceedings of 19th National Conference on Artificial Intelligence*, San Jose, California, USA.

Acknowledgments

This work was supported by the project entitled “Detecting and Visualizing Changes in Emotions in Texts” funded by the Academy of Finland.

- Robert D. Hume. 1969. Gothic Versus Romantic: A Reevaluation of the Gothic Novel. *PMLA*, 84(2): 282-90.
- Matthew G. Lewis. 1796. *The Monk: A Romance*. Printed for J. Saunders, Waterford, Ireland.
- Bing Liu, Minqing Hu and Junsheng Cheng. 2005. Opinion Observer: Analyzing and Comparing Opinions on the Web. *Proc. of the 14th International Conference on World Wide Web*, Chiba, Japan.
- George A. Miller. 1995. WordNet: A Lexical Database for English. *Communications of the ACM*, 38(11):39-41.
- Ann Radcliffe. 1794. *The Mysteries of Udolpho*. Printed for G.G. and J. Robinson, London, UK.
- Mary Shelley. 1818. *Frankenstein; or, The Modern Prometheus*. Lackington, Hughes, Harding, Mavor & Jones, London, UK.
- Stanford Center for Biomedical Informatics Research. 2011. *Protégé*. <http://protege.stanford.edu/> (Accessed March 23rd, 2011).
- Carlo Strapparava and Alessandro Valitutti. 2004. WordNet-Affect: an Affective Extension of WordNet. *Proceedings of the 4th International Conference on Language Resources and Evaluation*, Lisbon, Portugal.
- Maitte Taboada, Mary Ann Gillies and Paul McFetridge 2006. Sentiment Classification Techniques for Tracking Literary Reputation. *Proceedings of LREC Workshop Towards Computational Models of Literary Analysis*, Genoa, Italy.
- Horace Walpole. 1764. *The Castle of Otranto*. Thomas Lownds, London, UK.
- Devendra P. Varma. 1966. *The Gothic Flame: Being a History of the Gothic Novel in England - Its Origins, Efflorescence, Disintegration, and Residuary Influences*. Scarecrow Press, New York, USA.
- Yingcai Wu, Furu Wei, Shixia Liu, Norman Au, Weiwei Cui, Hong Zhou and Huamin Qu. 2010. OpinionSeer: Interactive Visualization of Hotel Customer Feedback. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1109-1118.