



From once upon a time to happily ever after: Tracking emotions in mail and books

Saif M. Mohammad

Emerging Technologies, National Research Council Canada, Ottawa, Ontario, Canada, K1A 0R6

ARTICLE INFO

Available online 23 May 2012

Keywords:

Emotion analysis
Sentiment analysis
Lexicon
Email
Fairy tales
Enron Corpus
Google Books Corpus

ABSTRACT

In this paper, we show how sentiment analysis can be used in tandem with effective visualizations to quantify and track emotions in mail and books. We study a number of specific datasets and show, among other things, how collections of texts can be organized for affect-based search and how books portray different entities through co-occurring emotion words. Analysis of the Enron Email Corpus reveals that there are marked differences across genders in how they use emotion words in work-place email. Finally, we show that fairy tales have more extreme emotion densities than novels.

Crown Copyright © 2012 Published by Elsevier B.V. All rights reserved.

1. Introduction

Emotions are an integral part of how humans perceive and communicate with the outside world. We convey emotion through our facial expressions, our speech, and through our writing. A given sentence may be pertinent to many different entities and determining the emotions evoked by one entity in another is fairly challenging—often requiring information not present in the sentence itself. For example, consider this headline in a newspaper from 2009:

When your cartoon can get you killed.

The article is about the controversy surrounding a particular episode of the television series *South Park*. The entities involved here are the creators of the show and the extremists issuing them death threats. The sentence has a writer and a large number of readers. All of these people may be expressing or feeling certain emotions. However, identifying the emotions associated with different entities requires not just the analysis of the target sentence, but often also of the context, entity behavior, and world knowledge. Thus, it is not surprising that current methods that attempt this task have relatively low accuracies [48].

However, for the first time in our history, we now have access to hundreds of thousands of digitized mail, books, and social-media communications. Even though making accurate predictions of individual instances may be error prone, simple methods can be used to draw reliable conclusions from many occurrences of a target entity. In this paper, we show how we created a large word–emotion association lexicon by crowdsourcing (Section 3), and use it to analyze the

use of emotion words in large collections of text. Specifically, we show how sentiment analysis can be used in tandem with effective visualizations to quantify and track emotions in mail (Sections 4–6) and in books (Sections 7–11). Many of these techniques can also apply to data from other forms of communication, such as Twitter feeds.

The lexicon we created has manual annotations of a word's associations with positive polarity negative polarity, and eight emotions—joy, sadness, anger, fear, trust, disgust, surprise, anticipation. These emotions have been argued to be the eight basic and prototypical emotions [43].

1.1. Emotion analysis of mail

Letters have long been a channel to convey emotions, explicitly and implicitly, and now with the widespread usage of email, we have access to unprecedented amounts of text that we ourselves have written. Automatic analysis and tracking of emotions in mail has a number of benefits including:

1. *Decision Support Tool*: Helping physicians identify patients who have a higher likelihood of attempting suicide [29,39] The 2011 Informatics for Integrating Biology and the Bedside (i2b2) challenge by the National Center for Biomedical Computing is on detecting emotions in suicide notes.
2. *Social Analysis*: Understanding how genders communicate through work-place and personal email [7].
3. *Productivity and Self-Assessment Tool*: Tracking emotions towards people and entities, over time. For example, did a certain managerial course bring about a measurable change in one's interpersonal communication?

E-mail address: saif.mohammad@nrc-cnrc.gc.ca.

4. *Health Applications*: Determining if there is a correlation between the emotional content of letters and changes in a person's social, economic, or physiological state. Sudden and persistent changes in the amount of emotion words in mail may be a sign of psychological disorder.
5. *Search*: Enabling affect-based search. For example, efforts to improve customer satisfaction can benefit by searching the received mail for snippets expressing anger [13,15].
6. *Writing Aids*: Assisting in writing emails that convey only the desired emotion, and avoiding misinterpretation [26].

In Section 4, we show comparative analyses of emotion words in love letters, hate mail, and suicide notes. This is done: (a) To determine the distribution of emotion words in these types of mail, as a first step towards more sophisticated emotion analysis (for example, in developing a depression–happiness scale for Application 1), and (b) To use these corpora as a testbed to establish that the emotion lexicon and the visualizations we propose help interpret the emotions in text. In Section 5, we analyze how men and women differ in the kinds of emotion words they use in work-place email (Application 2). Finally, in Section 6, we show how emotion analysis can be integrated with email services such as Gmail to help people track emotions in the emails they send and receive (Application 3).

1.2. Emotion analysis of books

Literary texts, such as novels, fairy tales, fables, romances, and epics tend to be rich in emotions. With widespread digitization of text, we now have easy access to large amounts of such literary texts. Project Gutenberg provides access to 34,000 books [24].¹ Google is providing n-gram sequences, and their frequencies, from more than 5.2 million digitized books, as part of the *Google Books Corpus (GBC)* [30].² Emotion analysis of books has many applications, including:

1. *Search*: Allowing search based on emotions. For example, retrieving the darkest of the Brothers Grimm fairy tales, or finding snippets from the Sherlock Holmes series that build the highest sense of anticipation and suspense.
2. *Social Analysis*: Identifying how books have portrayed different people and entities over time. For example, the distribution of emotion words used in proximity to mentions of women, race, and homosexuals.
3. *Comparative analysis of literary works, genres, and writing styles*: For example, is the distribution of emotion words in fairy tales significantly different from that in novels? Do women authors use a different distribution of emotion words than their male counterparts? Did Hans C. Andersen use emotion words differently than Beatrix Potter?
4. *Summarization*: For example, automatically generating summaries that capture the different emotional states of the characters in a novel.
5. *Analyzing Persuasion Tactics*: Analyzing emotion words and their role in persuasion [4,28].

We present a number of visualizations that help track and analyze the use of emotion words in individual texts and across very large collections, which is especially useful in Applications 1, 2, and 3 described above (Sections 7 and 8). Using the Google Books Corpus we show how to determine emotion associations portrayed in books towards different entities (Section 9). We introduce the concept of emotion word density, and using the Brothers Grimm fairy tales as

an example, we show how collections of text can be organized for better search (Section 10). Finally, for the first time, we compare a collection of novels and a collection of fairy tales using an emotion lexicon to show that fairy tales have a much wider distribution of emotion word densities than novels (Section 11).

This work is part of a broader project to provide an affect-based interface to Project Gutenberg. Given a search query, the goal is to provide users with relevant visualizations presented in this paper, and the ability to search for text snippets that have high emotion word densities.

2. Related work

Over the last decade, there has been considerable work in sentiment analysis, especially in determining whether a term has a positive or negative polarity [25,34,54]. There is also work in more sophisticated aspects of sentiment, for example, in detecting emotions such as anger, joy, sadness, fear, surprise, and disgust [2,5,33]. The technology is still developing and it can be unpredictable when dealing with short sentences, but it has been shown to be reliable when drawing conclusions from large amounts of text [14,37,41].

Automatically analyzing affect in emails has primarily been done for automatic gender identification [9,10], but it has relied mostly on surface features such as exclamations and very small emotion lexicons. The WordNet Affect Lexicon (WAL) [49] has a few hundred words annotated with associations to a number of affect categories including the six Ekman emotions (joy, sadness, anger, fear, disgust, and surprise).³ General Inquirer (GI) [47] has 11,788 words labeled with 182 categories of word tags, including positive and negative polarity.⁴ Affective Norms for English Words (ANEW) has pleasure (happy–unhappy), arousal (excited–calm), and dominance (controlled–in control) ratings for 1034 words.⁵ Mohammad and Turney [33] compiled emotion annotations for about 4000 words with eight emotions (six of Ekman, trust, and anticipation).

Empirical assessment of emotions in literary texts has sometimes relied on human annotation of the texts, but this has restricted the number of texts analyzed. For example, Alm and Sproat [1] annotated twenty two Brothers Grimm fairy tales to show that fairy tales often began with a neutral sentence and ended with a happy sentence. Here we use out-of-context word–emotion associations and analyze individual texts to very large collections.

Automatic systems for analyzing emotional content of text follow many different approaches: a number of these systems look for specific emotion denoting words [18], some determine the tendency of terms to co-occur with seed words whose emotions are known [45], some use hand-coded rules [35,36], and some use machine learning and a number of emotion features, including emotion denoting words [2,3].

Much recent work focuses on six emotions studied by Ekman [17]. These emotions—joy, sadness, anger, fear, disgust, and surprise—are a subset of the eight proposed by Plutchik [43]. There is less work on complex emotions, for example, work by Pear et al. [42] which focuses on politeness, rudeness, embarrassment, formality, persuasion, deception, confidence, and disbelief. Francisco and Gervás [19] marked sentences in fairy tales with tags for pleasantness, activation, and dominance, using lexicons of words associated with the three categories.

There has also been some interesting work in visualizing emotions [21,44,50]. Mohammad [31] describes work on identifying colors associated with emotion words.

¹ Project Gutenberg: <http://www.gutenberg.org>.

² GBC: <http://ngrams.googlelabs.com/datasets>.

³ WAL: <http://wdomains.fbk.eu/wnaffect.html>.

⁴ GI: <http://www.wjh.harvard.edu/~inquirer>.

⁵ ANEW: <http://csea.php.ufl.edu/media/anevmessage.html>.

3. Emotion analysis

3.1. Emotion lexicon

We created a large word–emotion association lexicon by crowdsourcing to Amazon’s mechanical Turk.⁶ We follow the method outlined in Mohammad and Turney [33]. Unlike Mohammad and Turney, who used the *Macquarie Thesaurus* [6], we use the *Roget Thesaurus* as the source for target terms.⁷ Since the 1911 US edition of *Roget’s* is available freely in the public domain, it allows us to distribute our emotion lexicon without the burden of restrictive licenses. We annotated only those words that occurred more than 120,000 times in the Google n-gram corpus.⁸

The *Roget’s Thesaurus* groups related words into about a thousand categories, which can be thought of as coarse senses or concepts [55]. If a word is ambiguous, then it is listed in more than one category. Since a word may have different emotion associations when used in different senses, we obtained annotations at word-sense level by first asking an automatically generated word-choice question pertaining to the target:

Q1. Which word is closest in meaning to *shark* (target)?

- *car*
- *tree*
- *fish*
- *olive*

The near-synonym is taken from the thesaurus, and the distractors are randomly chosen words. This question guides the annotator to the desired sense of the target word. It is followed by ten questions asking if the target is associated with positive sentiment, negative sentiment, anger, fear, joy, sadness, disgust, surprise, trust, and anticipation. The questions are phrased exactly as described in Mohammad and Turney [33].

If an annotator answers Q1 incorrectly, then we discard information obtained from the remaining questions. Thus, even though we do not have correct answers to the emotion association questions, likely incorrect annotations are filtered out. About 10% of the annotations were discarded because of an incorrect response to Q1.

Each term is annotated by 5 different people. For 74.4% of the instances, all five annotators agreed on whether a term is associated with a particular emotion or not. For 16.9% of the instances four out of five people agreed with each other. The information from multiple annotators for a particular term is combined by taking the majority vote. The lexicon has entries for about 24,200 word–sense pairs. The information from different senses of a word is combined by taking the union of all emotions associated with the different senses of the word. This resulted in a word-level emotion association lexicon for about 14,200 word types. These files are together referred to as the *NRC Emotion Lexicon version 0.92*.

3.2. Text analysis

Given a target text, the system determines which of the words exist in our emotion lexicon and calculates ratios such as the number of words associated with a particular emotion to the total number of emotion words in the text. This simple approach may not be reliable in determining if a particular sentence is expressing a certain emotion, but it is reliable in determining if a large piece of text has more emotional expressions compared to others in a corpus. Example applications include detecting spikes in anger words in close proximity to mentions of a target product in a Twitter data stream [13,15], and

⁶ Mechanical Turk: www.mturk.com/mturk/welcome.

⁷ Macquarie Thesaurus: www.macquarieonline.com.au. Roget’s Thesaurus: www.gutenberg.org/ebooks/10681.

⁸ The Google n-gram corpus is available through the LDC.

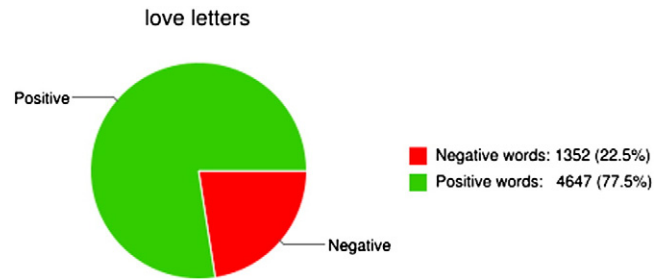


Fig. 1. Percentage of positive and negative words in the love letters corpus.

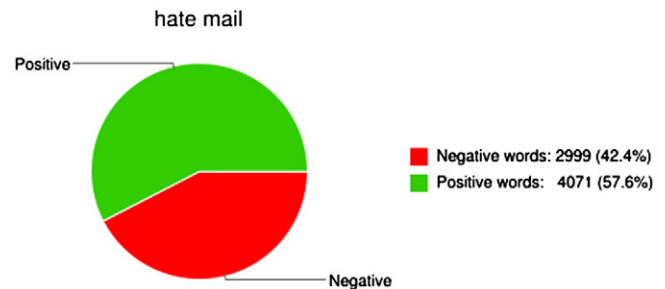


Fig. 2. Percentage of positive and negative words in the hate mail corpus.

literary analyses of text, for example, how novels and fairy tales differ in the use of emotion words [32].

Sections 4–6 describe the use of the NRC emotion lexicon to analyze mail, whereas Sections 7–11 describe its use to analyze books.

4. Emotional mail: love letters, hate mail, suicide notes

In this section, we quantitatively compare the emotion words in love letters, hate mail, and suicide notes. We compiled a *love letters corpus (LLC) v 0.1* by extracting 348 postings from lovingyou.com.⁹ We created a *hate mail corpus (HMC) v 0.1* by collecting 279 pieces of hate mail sent to the *Millenium Project*.¹⁰ The *suicide notes corpus (SNC) v 0.1* has 21 notes taken from Art Kleiner’s website.¹¹ We will continue to add more data to these corpora as we find them, and all three corpora are freely available.

Figs. 1–3 show the percentages of positive and negative words in the love letters corpus, hate mail corpus, and the suicide notes corpus. Figs. 4–6 show the percentages of different emotion words in the three corpora.

Fig. 7 is a bar graph showing the difference of emotion percentages in love letters and hate mail. Observe that as expected, love letters have a higher ratio of joy and trust words, whereas hate mail have a higher ratio of fear, sadness, disgust, and anger words.

The bar graph is effective at conveying the extent to which one emotion is more prominent in one text than another, but it does not convey the source of these emotions. Therefore, we calculate the *RelativeSalience* of an emotion word w across two target texts T_1 and T_2 :

$$\text{RelativeSalience}(w|T_1, T_2) = \frac{f_1}{N_1} - \frac{f_2}{N_2} \quad (1)$$

⁹ LLC: <http://www.lovingyou.com/content/inspiration/loveletters-topic.php?ID=lovingyou>.

¹⁰ HMC: <http://www.ratbags.com>.

¹¹ SNC: <http://www.well.com/art/suicidenotes.html#w>.




Fig. 3. Percentage of positive and negative words in the suicide notes corpus.




Fig. 5. Percentage of emotion words in the hate mail corpus.




Fig. 4. Percentage of emotion words in the love letters corpus.




Fig. 6. Percentage of emotion words in the suicide notes corpus.

Where, f_1 and f_2 are the frequencies of w in T_1 and T_2 , respectively. N_1 and N_2 are the total number of word tokens in T_1 and T_2 . Fig. 9 depicts a RelativeSalience word cloud of joy words in the love letters corpus with respect to the hate mail corpus. As expected, love letters, much more than hate mail, have terms such as *loving, baby, beautiful, feeling, and smile*. This is a nice sanity check of the manually created emotion lexicon. We used Google's freely available software to create the word clouds shown in this paper.¹² The most salient fear words in the suicide notes with respect to love letters, in decreasing order, were: *hell, kill, broke, worship, sorrow, afraid, loneliness, endless, shaking, and devil* (bar graph and word cloud not shown due to space constraints).

Fig. 8 is a difference bar graph for suicide notes and hate mail. Fig. 10 depicts a RelativeSalience word cloud of disgust words in the hate mail corpus with respect to the suicide notes corpus. The cloud shows many words that seem expected, for example *ignorant, quack, fraudulent, illegal, lying, and damage*. Words such as *cancer and disease* are prominent in this hate mail corpus because the *Millenium Project* denigrates various alternative treatment websites for cancer and other diseases, and consequently receives angry emails from some cancer patients and physicians.

5. Men and women: emotional differences in work-place email

There is a large amount of research at the intersection of gender and language (see bibliographies compiled by Schiffman [46] and Sunderland et al. [51]). It is widely believed that men and women use language differently, and this is true even in computer-mediated communications such as email [7]. It is claimed that women tend to foster personal relations [12,16] whereas men communicate for social position [52]. Women tend to share concerns and support others [7] whereas men prefer to talk about activities [8,11].

Otterbacher [40] investigated stylistic differences in how men and women write product reviews. Thelwall et al. [53] examine how men and women communicate over social networks such as MySpace. Here, for the first time, we use an emotion lexicon of more than 14,000 words to investigate gender differences in work-place communications. The analysis shown here does not prove the propositions mentioned above; however, it provides empirical support to the claim that men and women use emotion words to different degrees.

We chose the Enron email corpus [23] as the source of work-place communications because it remains the only large publicly available collection of emails.¹³ It consists of more than 200,000 emails sent between October 1998 and June 2002 by 150 people in senior managerial positions at the Enron Corporation, a former American energy, commodities, and services company. The emails largely pertain to official business but also contain personal communication.

In addition to the body of the email, the corpus has meta-information such as the time stamp and the email addresses of the sender and receiver. Just as in Cheng et al. [9]: (1) we removed emails whose body had fewer than 50 words or more than 200 words, (2) we identified the gender of each of the 150 people solely from their names. If the name was not a clear indicator of gender, then the person was marked as "gender-untagged". This resulted in tagging 41 employees as female and 89 as male; 20 were left gender-untagged. Emails sent from and to gender-untagged employees were removed from all further analysis, leaving 32,045 mails (19,920 mails sent by men and 12,125 mails sent by women). We then determined the number of emotion words in emails written by men, in emails written by women, in emails written by men to women, men to men, women to men, and women to women.

5.1. Analysis

Fig. 11 shows the difference in percentages of emotion words in emails sent by men from the percentage of emotion words in

¹² Google WordCloud: <http://visapi-gadgets.googlecode.com/svn/trunk/wordcloud/doc.html>.

¹³ The Enron email corpus is available at <http://www-2.cs.cmu.edu/enron>.




Fig. 7. Difference in percentages of emotion words in the love letters corpus and the hate mail corpus. The RelativeSalience word cloud for the joy bar is shown in the figure to the right (Fig. 8).




Fig. 8. Difference in percentages of emotion words in the suicide notes corpus and the hate mail corpus.



Fig. 9. Love letters corpus - hate mail corpus: RelativeSalience word cloud for joy.



Fig. 10. Suicide notes - hate mail: RelativeSalience word cloud for disgust.

emails sent by women. Observe the marked difference is in the percentage of trust words. The men used many more trust words than women. Fig. 13 shows the RelativeSalience word cloud of these trust words.

Fig. 12 shows the difference in percentages of emotion words in emails sent to women and the percentage of emotion words in emails sent to men. Observe the marked difference once again in the percentage of trust words and joy words. The men received emails with more trust words, whereas the women received emails with more joy words. Fig. 14 shows the RelativeSalience word cloud of joy.

Fig. 15 shows the difference in emotion words in emails sent by men to women and the emotions in mails sent by men to men. Apart from trust words, there is a marked difference in the percentage of anticipation words. The men used many more anticipation words when writing to women, than when writing to other men. Fig. 17 shows the RelativeSalience word cloud of these anticipation words.

Figs. 16, 18–20 show difference bar graphs and RelativeSalience word clouds analyzing some other possible pairs of correspondences.

5.2. Discussion

Figs. 12, 15, 16, and 19 support the claim that when writing to women, both men and women use more joyous and cheerful words than when writing to men. Figs. 12, 15 and 16 show that both men




Fig. 11. Difference in percentages of emotion words in emails sent by women and emails sent by men.




Fig. 12. Difference in percentages of emotion words in emails sent to women and emails sent to men.



Fig. 13. Emails by women - emails by men: RelativeSaliency word cloud of trust.



Fig. 14. Emails to women - emails to men: RelativeSaliency word cloud of joy.




Fig. 15. Difference in percentages of emotion words in emails sent by men to women and by men to men.




Fig. 16. Difference in percentages of emotion words in emails sent by women to women and by women to men.



Fig. 17. Emails by men to women - emails by men to men: RelativeSaliency word cloud of anticipation.



Fig. 18. Emails by women to women - emails by women to men: RelativeSaliency word cloud of sadness.




Fig. 19. Difference in percentages of emotion words in emails sent by men to men and by women to women.



Fig. 20. Emails by men to men - emails by women to women: RelativeSaliency word cloud of fear.

and women use lots of trust words when writing to men. Figs. 11, 16, and 19 are consistent with the notion that women use more cheerful words in emails than men. The sadness values in these figures are consistent with the claim that women tend to share their worries with other women more often than men with other men, men with women, and women with men. The fear values in the Figs. 15 and 19 suggest that men prefer to use a lot of fear words, especially when communicating with other men. Thus, women communicate relatively more on the joy–sadness axis, whereas men have a preference for the trust–fear axis. It is interesting how there is a markedly higher percentage of anticipation words in cross-gender communication than in same-sex communication (Figs. 15, 16, and 19).

6. Tracking sentiment in personal email

In the previous section, we showed analyses of sets of emails that were sent across a network of individuals. In this section, we show visualizations catered toward individuals—who in most cases have access to only the emails they send and receive. We are using Google Apps API to develop an application that integrates with Gmail (Google's email service), to provide users with the ability to track their emotions towards people they correspond with.¹⁴ Below we show some of these visualizations by selecting John Arnold, a former employee at Enron, as a stand-in for the actual user.

Fig. 21 shows the percentage of positive and negative words in emails sent by John Arnold to his colleagues. John can select any of the bars in the figure to reveal the difference in percentages of emotion words in emails sent to that particular person and all the emails sent out. Fig. 22 shows such a graph pertaining to Andy Zipper. Fig. 23 shows the percentage of positive and negative words in each of the emails sent by John to Andy.

In the future, we will make a public call for volunteers interested in our Gmail emotion application, and we will request access to

frequencies of emotion words in their emails for a large-scale analysis of personal email. The application will protect the privacy of the users by passing emotion word frequencies, gender, and age, but no text, names, or email ids.

7. Distribution of emotion words in books

As mentioned earlier, literary texts, such as novels, fairy tales, fables, romances, and epics are effective channels for conveying human emotions. Further, different genres may be rich in different emotions. In this section, we show the use of the emotion lexicon, and the visualizations proposed earlier in the paper, to analyze books. In the remaining sections, we show visualizations designed specifically for books, although many of them too may be applied to email and social-media data.

Fig. 24 conveys the difference bar graph between Shakespeare's famous tragedy, *Hamlet*, and his comedy, *As you like it*. Figs. 25 and 26 show snippets of the RelativeSaliency word clouds for trust and sadness, respectively. Observe how one can clearly see that *Hamlet* has more fear, sadness, disgust, and anger, and less joy, trust, and anticipation. Figs. 25 and 26 depict snippets of RelativeSaliency word clouds of trust words and sadness words across *Hamlet* and *As You Like it*.

8. Flow of emotions

Literary researchers as well as casual readers may be interested in noting how the use of emotion words has varied through the course of a book. Figs. 27–29 show the flow of joy, trust, and fear in *As You Like it* (comedy), *Hamlet* (tragedy), and *Frankenstein* (horror), respectively. As expected, the visualizations depict the novels to be progressively more dark than the previous ones in the list. Also note that *Frankenstein* is much darker in the final chapters.




Fig. 21. Emails sent by John Arnold.

¹⁴ Google Apps API: <http://code.google.com/googleapps/docs>.




Fig. 22. Difference in percentages of emotion words in emails sent by John Arnold to Andy Zipper and emails sent by John to all.




Fig. 23. Emails sent by John Arnold to Andy Zipper.

9. Determining emotion associations through co-occurring words

Words found in proximity of target entities can be good indicators of emotions associated with the targets. Google recently released n-gram frequency data from all the books they scanned up to July 15, 2009.¹⁵ It is a digitized version of about 5.2 million books, and the English portion has about 361 billion words. The data consists of 5-grams, frequency in a particular year, and the year. We analyzed the 5-gram files (about 800 GB of data) to quantify the emotions associated with different target entities. We ignored data from books published before 1800 as that period is less comprehensively covered by Google books. We chose to group the data into five-year bins, though other groupings are reasonable as well. Given a target entity of interest, the system identifies all 5-grams that contain the target word, identifies all the emotion words in those n-grams (other than the target word itself), and calculates percentages of emotion words.

Fig. 30 shows the percentage of fear words in the n-grams of different countries. Observe, that there is a marked rise of fear words around World War I (1914–1918) for Germany, America, and China. There is a spike for China around 1900, likely due to the unrest leading up to the Boxer Rebellion (1898–1901).¹⁶ The 1810–1814 spike for China is probably correlated with descriptions of piracy in the South China Seas, since the era of the commoner-pirates of mid-Qing dynasty came to an end in 1810.¹⁷ India does not see a spike during World War I, but has a spike in the 1940's probably reflecting heightened vigor in the independence struggle (Quit India Movement of 1942¹⁸) and growing involvement in World War II (1939–1945).¹⁹

Fig. 31 shows two curves for the percentages of joy words in 5-grams that include *woman* and *man*, respectively. Fig. 32 shows similar curves for anger words. We find a larger percentage of joy words associated with *woman* and a larger percentage of anger words associated with

man. The curves for trust words for *man* and *woman* are almost identical (figure not shown).

10. Emotion word density

Apart from showing the flow of emotion words, the use of emotion words in a book can also be quantified by calculating the number of emotion words one is expected to see on reading every X words. We will refer to this metric as *emotion word density*, or simply *emotion density* for short. All emotion densities reported in this paper are for X = 10, 000. The dotted line in Fig. 33 shows the negative word density plot of 192 fairy tales collected by Brothers Grimm. The joy and sadness word densities are also shown—the thin and thick lines, respectively. A person interested in understanding the use of emotion words in the fairy tales collected by Brothers Grimm can further select any two fairy tales from the plot, say *Cinderella* and *Godfather Death*, to reveal a bar graph showing the difference in percentages of emotions in the two texts (figure not shown due to space constraints). The RelativeSalience word cloud of fear (figure not shown) shows that the prominent fear words in *Godfather Death* are *death*, *ill*, *beware*, *poverty*, *devil*, *astray*, *risk*, *illness*, *threatening*, *horrified* and *revenge*.

Arranging documents as per emotion word density has a number of other applications too, such as arranging articles from different newspapers on the same event in order of different emotion densities, arranging different user demographics as per their emotion densities towards target entities in the blogs they write, and so on.

11. Novels and fairy tales: emotional differences

Novels and fairy tales are two popular forms of literary prose. Both forms tell a story, but a fairy tale has certain distinct characteristics such as (a) archetypal characters (peasant, king) (b) clear identification of good and bad characters, (c) happy ending, (d) presence of magic and magical creatures, and (d) a clear moral [20]. Fairy tales are extremely popular and appeal to audiences through emotions—they convey personal concerns, subliminal fears, wishes, and fantasies in an exaggerated manner [20,22,38]. However, there have not been any large-scale empirical studies to compare affect in fairy tales and novels. Here for the first time, we compare the use of emotion-associated words in fairy tales and novels using a large lexicon.

Specifically, we are interested in determining whether: (1) fairy tales on average have a higher emotional density than novels, (2) different fairy tales focus on different emotions such that some fairy tales have high densities for certain emotion, whereas others have low emotional densities for those same emotions.

We used the Corpus of English Novels (CEN) and the Fairy Tale Corpus (FTC) for our experiments.²⁰ The Corpus of English Novels is

¹⁵ Google books data: <http://ngrams.googlelabs.com/datasets>.

¹⁶ http://en.wikipedia.org/wiki/Boxer_Rebellion.

¹⁷ http://www.ias.nl/nl/36/IIAS_NL36_07.pdf.

¹⁸ http://en.wikipedia.org/wiki/Quit_India_Movement.

¹⁹ http://en.wikipedia.org/wiki/India_in_World_War_II.

²⁰ CEN: <https://perswww.kuleuven.be/~u0044428/cen.htm>. FTC: https://www.l2f.inesc-id.pt/wiki/index.php/Fairy_tale_corpus.

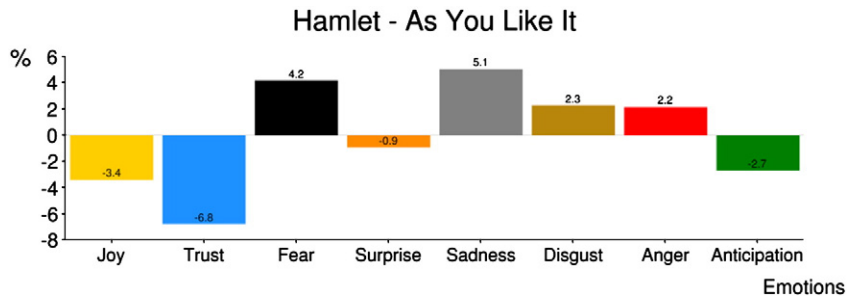


Fig. 24. Difference in percentage scores for each of the eight basic emotions in Hamlet and As you like it.



Fig. 25. Hamlet – As You Like It: RelativeSalience word cloud for trust words.



Fig. 26. Hamlet – As You Like It: RelativeSalience word cloud for sadness words.

a collection of 292 novels written between 1881 and 1922 by 25 British and American novelists. It was compiled from Project Gutenberg at the Catholic University of Leuven by Hendrik de Smet. It consists of about 26 million words. The Fairy Tale Corpus [27] has 453 stories, close to 1 million words, downloaded from Project Gutenberg. Even though many fairy tales have a strong oral tradition, the stories in this collection were compiled, translated, or penned in the 19th century by the Brothers Grimm, Beatrix Potter, and Hans C. Andersen to name a few.

We calculated the polarity and emotion word density of each of the novels in CEN and each of the fairy tales in FTC. Table 1 lists the mean densities as well as standard deviation for each of the eight basic emotions in the two corpora.

We find that the mean densities for anger and sadness across CEN and FTC are not significantly different. However, fairy tales have significantly higher anticipation, disgust, joy, and surprise densities when compared to novels ($p < 0.001$). On the other hand, they have significantly lower trust word density than novels. Further, the standard deviations for all eight emotions are significantly different across the two corpora ($p < 0.001$). The fairy tales, in general, have a much larger standard deviation than the novels. Thus for each of the 8 emotions, there are more fairy tales than novels having high emotion densities and there are more fairy tales than novels having low emotion densities.

Table 2 lists the mean densities as well as standard deviation for negative and positive polarity words in the two corpora. The table states, for example, that for every 10,000 words in the CEN, one can expect to see about 1670 negative words. We find that fairy tales, on average, have a significantly lower number of negative terms, and a significantly higher number of positive words ($p < 0.001$).

In order to obtain a better sense of the distribution of emotion densities, we generated histograms by counting all texts that had emotion densities between 0–99, 100–199, 200–399, and so on. A large standard deviation for fairy tales could be due to one of at least two reasons: (1) the histogram has a bimodal distribution—most of the fairy tales have extreme emotion densities (either much higher than that of the novels, or much smaller). (2) the histogram approaches a normal distribution such that more fairy tales than novels have extreme emotion densities. Figs. 34 through 39 show




Fig. 27. Timeline of the emotions in As You Like It.




Fig. 28. Timeline of the emotions in Hamlet.




Fig. 29. Timeline of the emotions in Frankenstein.

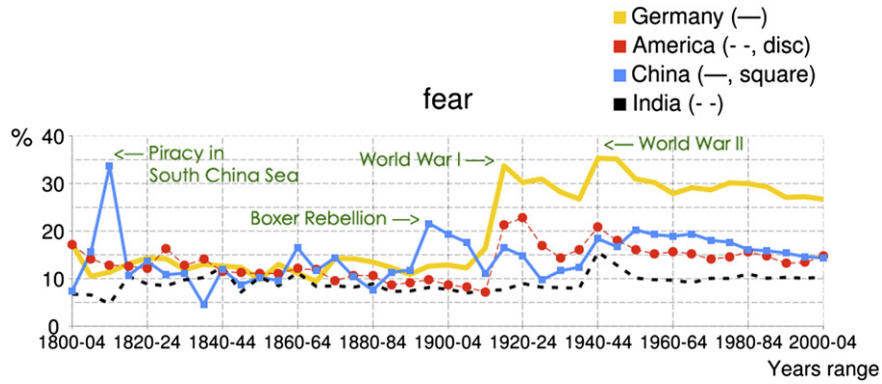


Fig. 30. Percentage of fear words in close proximity to occurrences of America, China, Germany, and India in books from the year 1800 to 2004. Source: 5-gram data released by Google.




Fig. 31. Percentage of joy words in close proximity to occurrences of man and woman in books from the year 1800 to 2004. Source: 5-gram data released by Google.




Fig. 32. Percentage of anger words in close proximity to occurrences of man and woman in books.

Table 1

Density of emotion words in novels and fairy tales: number of emotion words in every 10,000 words.

	Anger		Anticipation		Disgust		Fear	
	Mean	Σ	Mean	Σ	Mean	Σ	Mean	Σ
CEN	746	±162	1230	±126	591	±135	975	±225
FTC	749	±393	1394	±460	682	±460	910	±454

	Joy		Sadness		Surprise		Trust	
	Mean	Σ	Mean	Σ	Mean	Σ	Mean	Σ
CEN	1164	±196	785	±159	628	±93	1473	±190
FTC	1417	±467	814	±443	680	±325	1348	±491

Table 2

Density of polarity words in novels and fairy tales: number of polar words in every 10,000 words.

	Negative		Positive	
	Mean	Σ	Mean	σ
CEN	1670	±243	2602	±278
FTC	1543	±613	2808	±726




Fig. 33. The Brothers Grimm fairy tales arranged in increasing order of negative word density (number of negative words in every 10,000 words). The plot is of 192 stories but the x-axis has labels for only a few due to lack of space. A user may select any two tales, say Cinderella and Godfather Death (see arrows), to generate further visualizations such as the difference bar graph and salience word clouds.




Fig. 34. Histogram of texts with different negative word densities. On the x-axis: 1 refers to density between 0 and 100, 2 refers to 100 to 200, and so on.




Fig. 37. Histogram of texts with different positive word densities. On the x-axis: 1 refers to density between 0 and 100, 2 refers to 100 to 200, and so on.

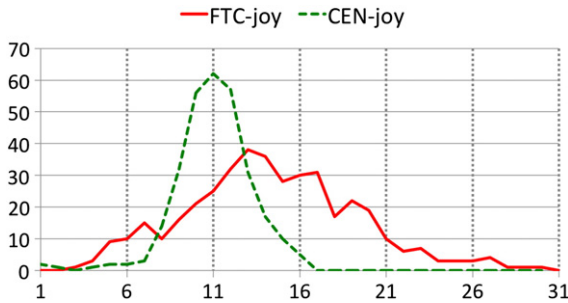


Fig. 35. Histogram of texts with different joy word densities.




Fig. 38. Histogram of texts with different surprise word densities.




Fig. 36. Histogram of texts with different anger word densities.




Fig. 39. Histogram of texts with different anticipation word densities.

histograms comparing novels and fairy tales for positive and negative polarities, as well as for a few emotions. Observe that fairy tales do not have a bimodal distribution, and case (2) holds true.

12. Conclusions and future work

We have created a large word–emotion association lexicon by crowdsourcing, and used it to analyze and track the distribution of emotion words in books and mail.²¹ We showed how different visualizations and word clouds can be used to effectively interpret the results of the emotion analysis.

We compared emotion words in love letters, hate mail, and suicide notes. We analyzed work–place email and showed that women use and receive a relatively larger number of joy and sadness words, whereas men use and receive a relatively larger number of trust and fear words.

We introduced the concept of emotion word density, and using the Brothers Grimm fairy tales as an example, we showed how collections of text can be organized for better search. Using the Google Books Corpus we showed how to determine emotion associations portrayed in books towards different entities. Finally, for the first

time, we compared a collection of novels and a collection of fairy tales using the emotion lexicon to show that fairy tales have a much wider distribution of emotion word densities than novels. This work is part of a broader project to provide an affect-based interface to Project Gutenberg. Given a search query, the goal is to provide users with relevant plots presented in this paper. Further, they will be able to search for snippets from multiple texts that have strong emotion word densities. We are also interested in applying these techniques to understand emotions in social-media communications, such as Twitter feeds.

Acknowledgments

Thanks to Tony (Wenda) Yang for coding an online emotion analyzer. Grateful thanks to Peter Turney and Tara Small for many wonderful ideas.

References

- [1] C.O. Alm, R. Sproat, Emotional sequencing and development in fairy tales, Springer, 2005, pp. 668–674.

²¹ Please send an e-mail to saif.mohammad@nrc-cnrc.gc.ca to obtain the latest version of the NRC Emotion Lexicon.

- [2] C.O. Alm, D. Roth, R. Sproat, Emotions from text: Machine learning for text-based emotion prediction, in: Proceedings of the Joint Conference on HLT–EMNLP, Vancouver, Canada, 2005.
- [3] S. Aman, S. Szpakowicz, Identifying expressions of emotion in text, in: V. Matoušek, P. Mautner (Eds.), Text, Speech and Dialogue, Lecture Notes in Computer Science, vol. 4629, Springer, Berlin / Heidelberg, 2007, pp. 196–205.
- [4] R. Bales, Persuasion in the French personal novel: Studies of Chateaubriand, Constant, Balzac, Nerval, and Fromentin, Summa Publications, Birmingham, Alabama, 1997.
- [5] J. Bellegarda, Emotion analysis using latent affective folding and embedding, in: Proceedings of the NAACL–HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, Los Angeles, California, 2010.
- [6] J. Bernard (Ed.), The Macquarie Thesaurus, Macquarie Library, Sydney, Australia, 1986.
- [7] B. Boneva, R. Kraut, D. Frohlich, Using e-mail for personal relationships, The American Behavioral Scientist 45 (2001) 530–549.
- [8] M.A. Caldwell, L.A. Peplau, Sex differences in same-sex friendships, Sex Roles 8 (1982) 721–732.
- [9] N. Cheng, X. Chen, R. Chandramouli, K. Subbalakshmi, Gender identification from e-mails, Computational Intelligence and Data Mining, CIDM '09. IEEE Symposium on, , 2009, pp. 154–158.
- [10] M. Corney, O. de Vel, A. Anderson, G. Mohay, Gender-preferential text mining of e-mail discourse, Computer Security Applications Conference, Annual, 2002, pp. 282–289.
- [11] L.R. Davidson, L. Duberman, Friendship: communication and interactional patterns in same-sex dyads, Sex Roles 8 (1982) 809–822, <http://dx.doi.org/10.1007/BF00287852>.
- [12] K. Deaux, B. Major, Putting gender into context: an interactive model of gender-related behavior, Psychological Review 94 (1987) 369–389.
- [13] A.B.C. Díaz, F.J.M. Ruiz, The consumer's reaction to delays in service, International Journal of Service Industry Management 13 (2002).
- [14] P. Dodds, C. Danforth, Measuring the happiness of large-scale written expression: songs, blogs, and presidents, Journal of Happiness Studies 11 (2010) 441–456, <http://dx.doi.org/10.1007/s10902-009-9150-9>.
- [15] L. Dubé, M.F. Maute, The antecedents of brand switching, brand loyalty and verbal responses to service failure, Advances in Services Marketing and Management 5 (1996) 127–151.
- [16] A.H. Eagly, V.J. Steffen, Gender stereotypes stem from the distribution of women and men into social roles, Journal of Personality and Social Psychology 46 (1984) 735–754.
- [17] P. Ekman, An argument for basic emotions, Cognition and Emotion 6 (1992) 169–200.
- [18] C. Elliott, The affective reasoner: A process model of emotions in a multi-agent system, Ph.D. thesis, Institute for the Learning Sciences, Northwestern University, 1992.
- [19] V. Francisco, P. Gervás, Automated mark up of affective information in English texts, in: P. Sojka, I. Kopeček, K. Pala (Eds.), Text, Speech and Dialogue, Lecture Notes in Computer Science, vol. 4188, Springer, Berlin / Heidelberg, 2006, pp. 375–382.
- [20] S.S. Jones, The Fairy Tale: The Magic Mirror of the Imagination, Routledge, 2002.
- [21] A. Kalra, K. Karahalios, Texttone: Expressing Emotion Through Text, in: M.F. Costabile, F. Patern (Eds.), Human–Computer Interaction – INTERACT 2005, Lecture Notes in Computer Science, vol. 3585, Springer, Berlin / Heidelberg, 2005, pp. 966–969.
- [22] V. Kast, Through Emotions to Maturity: Psychological Readings of Fairy Tales, Fromm Intl, 1993.
- [23] B. Klimt, Y. Yang, Introducing the enron corpus, in: CEAS, 2004.
- [24] M. Lebert, Project Gutenberg (1971–2009), Benediction Classic, 2009.
- [25] A. Lehrer, Semantic fields and lexical structure, North-Holland, American Elsevier, Amsterdam, NY, 1974.
- [26] H. Liu, H. Lieberman, T. Selker, A model of Textual Affect Sensing Using Real-world Knowledge, Proceedings of the 8th international conference on Intelligent user interfaces, IUI '03, ACM, New York, NY, 2003, pp. 125–132.
- [27] P.V. Lobo, D. Martins de Matos, Fairy tale corpus organization using latent semantic mapping and an item-to-item top-n recommendation algorithm, in: Language Resources and Evaluation Conference - LREC 2010, European Language Resources Association (ELRA), Malta.
- [28] P. Mannix, The rhetoric of antinuclear fiction: persuasive strategies in novels and films, Bucknell University Press, Associated University Presses, London, 1992.
- [29] P. Matykievicz, W. Duch, J.P. Pestian, Clustering Semantic Spaces of Suicide Notes and Newsgroups Articles, Proceedings of the Workshop on Current Trends in Bio-medical Natural Language Processing, BioNLP '09, Association for Computational Linguistics, Stroudsburg, PA, USA, 2009, pp. 179–184.
- [30] J.B. Michel, Y.K. Shen, A.P. Aiden, A. Veres, M.K. Gray, T.G.B. Team, J.P. Pickett, D. Hoiberg, D. Clancy, P. Norvig, J. Orwant, S. Pinker, M.A. Nowak, E.L. Aiden, Quantitative analysis of culture using millions of digitized books, Science 331 (2011) 176–182.
- [31] S.M. Mohammad, Even the abstract have colour: Consensus in word-colour associations, in: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Portland, OR, USA.
- [32] S.M. Mohammad, From once upon a time to happily ever after: Tracking emotions in novels and fairy tales, in: Proceedings of the ACL 2011 Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH), Portland, OR, USA, 2011.
- [33] S.M. Mohammad, P.D. Turney, Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon, in: Proceedings of the NAACL–HLT Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, LA, California, 2010.
- [34] S.M. Mohammad, C. Dunne, B. Dorr, Generating high-coverage semantic orientation lexicons from overtly marked words and a thesaurus, in: Proceedings of Empirical Methods in Natural Language Processing (EMNLP-2009), Singapore, 2009, pp. 599–608.
- [35] A. Neviarouskaya, H. Prendinger, M. Ishizuka, Affect analysis model: novel rule-based approach to affect sensing from text, Natural Language Engineering 17 (2010) 95–135.
- [36] A. Neviarouskaya, H. Prendinger, M. Ishizuka, Compositionality principle in recognition of fine-grained emotions from text, in: Proceedings of the Proceedings of the Third International Conference on Weblogs and Social Media (ICWSM-09), San Jose, California, 2009, pp. 278–281.
- [37] D.E. O'Leary, Blog mining-review and extensions: from each according to his opinion, Decision Support Systems 51 (2011) 821–830 (Recent Advances in Data, Text, and Media Mining and Information Issues in Supply Chain and in Service System Design).
- [38] C. Orenstein, Little Red Riding Hood Unlocked: Sex, Morality, And The Evolution Of A Fairy Tale, Basic Books, 2003.
- [39] C.E. Osgood, E.G. Walker, Motivation and language behavior: a content analysis of suicide notes, Journal of Abnormal and Social Psychology 59 (1959) 58–67.
- [40] J. Otterbacher, Inferring Gender of Movie Reviewers: Exploiting Writing Style, Content and Metadata, Proceedings of the 19th ACM international conference on Information and knowledge management, CIKM '10, ACM, New York, NY, USA, 2010, pp. 369–378.
- [41] B. Pang, L. Lee, Opinion mining and sentiment analysis, Foundations and Trends in Information Retrieval 2 (2008) 1–135.
- [42] L. Pearl, M. Steyvers, Identifying emotions, intentions, and attitudes in text using a game with a purpose, in: Proceedings of the NAACL–HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, Los Angeles, California, 2010.
- [43] R. Plutchik, A general psychoevolutionary theory of emotion, Emotion: Theory, research, and experience, 1, 1980, pp. 3–33.
- [44] R. Rashid, J. Aitken, D. Fels, Expressing Emotions Using Animated Text Captions, Computers Helping People with Special Needs, Lecture Notes in Computer Science, vol. 4061, Springer, Berlin / Heidelberg, 2006, pp. 24–31.
- [45] J. Read, Recognising affect in text using pointwise-mutual information, Ph.D. thesis, Department of Informatics, University of Sussex, 2004.
- [46] H. Schifman, Bibliography of gender and language, <http://www.sas.upenn.edu/~haroldfs/popcult/bibliogs/gender/genbib.htm>.
- [47] P. Stone, D.C. Dunphy, M.S. Smith, D.M. Ogilvie, The General Inquirer: A Computer Approach to Content Analysis associates, The MIT Press, 1966.
- [48] C. Strapparava, R. Mihalcea, Semeval-2007 task 14: Affective text, Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007), Association for Computational Linguistics, Prague, Czech Republic, 2007, pp. 70–74.
- [49] C. Strapparava, A. Valitutti, Wordnet-Affect: An affective extension of WordNet, in: Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC-2004), Lisbon, Portugal, 2004, pp. 1083–1086.
- [50] P. Subasic, A. Huettner, Affect analysis of text using fuzzy semantic typing, IEEE Transaction on Fuzzy Systems 9 (2001) 483–496.
- [51] J. Sunderland, R.F. Duann, P. Baker, Gender and genre bibliography, <http://www.ling.lanccs.ac.uk/pubs/clsl/clsl122.pdf>.
- [52] D. Tannen, You Just Don't Understand: Women and Men in Conversation, Random House, 1991.
- [53] M. Thelwall, D. Wilkinson, S. Uppal, Data mining emotion in social network communication: gender differences in myspace, Journal of the American Society for Information Science and Technology 61 (2010) 190–199.
- [54] P. Turney, M. Littman, Measuring praise and criticism: Inference of semantic orientation from association, ACM Transactions on Information Systems (TOIS) 21 (2003) 315–346.
- [55] D. Yarowsky, C. Dunne, Word-sense disambiguation using statistical models of Roget's categories trained on large corpora, in: Proceedings of the 14th International Conference on Computational Linguistics (COLING-92), Nantes, France, 1992, pp. 454–460.

Dr. Saif Mohammad is a Research Officer at the Institute for Information Technology, National Research Council Canada (NRC). He received his Ph.D. in Computer Science from University of Toronto in January 2008. He was a Research Associate in the Institute of Advanced Computer Studies at the University of Maryland, College Park, before joining NRC in 2009. Saif's research interests are in Natural Language Processing, especially Lexical Semantics. He develops computational models for emotion detection, word-color associations, semantic distance, and lexical contrast.