People on Twitter usually smile @the-users-they-mention

Sentiment Analysis: detecting mentions and emoticons in one tweet

Melvin Roest

Bachelor Student of Information, Multimedia and Management
VU University, Amsterdam, The Netherlands
m.r.roest@vu.nl

Abstract

The internet could be considered to facilitate the virtual side of social interaction. However, it is not clear to what extent psychological principles apply to behavior on the internet. Social behavior is especially observable in social networks, such as Facebook, and microblogging services, such as Twitter. In this paper, we explore how positive and negative emoticons in a tweet could be used as a measure of opinion towards referred users in the same tweet. Furthermore, we discuss how our results compare against findings in (positive) psychology. From a psychological standpoint, our results indicate that Twitter is a flourishing community. On average, above 80 percent of the emoticons are positive in our data set of mentioned tweets. Further analysis showed that there are some exceptional outliers of Twitter users that are mentioned in tweets with mostly negative or positive emoticons. One Twitter user that is almost only referred to with negative emoticons is discussed in a case, another Twitter user that is almost only referred to with positive emoticons is discussed in another case. The possible reason of the two outliers are linked to the field of marketing.

0 Foreword

This paper is a bachelor thesis written by me and supervised by Spyros Voulgaris. I remember when I told my initial idea: "I want something to do with information retrieval and graph theory, I want to cluster people into clusters via the Twitter social graph and then do some information retrieval on it with regards to happiness and sadness. It's a bit foggy but I am working on the details". With that statement he accepted to supervise me.

For four weeks I searched for data sets about Twitter, but I could not find any. Eventually Spyros and I began brainstorming when it became clear that data sets were too hard to find (due to Twitters recent policy). During the brainstorm session Spyros came to a question which became the beginning of this thesis: "What would happen if you search for @mentions in a tweet and emoticons, and relate that to the user who is mentioned?" The answer involved a lot of programming work. Most of that programming work will not be shown in this thesis.

This is an exciting question in my mind, because it combines my two favorite fields in science which I know something about: business informatics and psychology. For me personally, this thesis represents my informal side - reading psychology research in my spare time - and my formal side - studying business informatics. It's also the reason why I enjoyed doing this project, to put it in my words, it was awesome!

I would like to acknowledge all people who helped on this paper. These people are: my supervisor Spyros Voulgaris (thanks for the flexible meeting hours!), my proofreaders Thomas van Drunen and Rebecca van Haastrecht. I would also like to thank Guido van Oorschot for his tips in capturing Twitter data. Thanks to the four of you, this thesis became a reality. Then there are people who helped me generate ideas which allowed me to think broader: Brendan Meeder (of Carnegie Mellon University), Jeffrey Bruijntjes and Maarten van Steen, thanks! And of course, I would like to thank everyone who read my abstract, you are an army of little proofreaders who are able to overthrow any small country! :) A thanks goes out to Willem van Hage, thanks for proofreading my introduction and for being the second reader of this thesis. Finally, I would like to thank my grandpa and grandma for being so supportive during this time.

A note to the one who will be reading this: when I write "we" I mean myself, apparently it is a convention to use the word "we". Furthermore, this paper is in English because Spyros (my supervisor) his best languages are Greek and English. Since my Greek is not that great I chose English.

To the reader: I hope you enjoy this paper and that it makes you smile, just like the positive emoticons.

1 Introduction

Humans have the need to be social among one and another (Myers 2010, 344-355). Microblogging services like Twitter provide a way of being social by following other Twitter users. These users write about almost everything that is conceivable within 140 characters or less, because Twitter applies a 140 character limit. The service became wildly popular in a short time-span[1]. A possible explanation for this is because Twitter is easy to use. At this point in time, hundreds of millions of tweets are tweeted every day.

The data derived from these tweets could be useful for scientists and companies interested in trends and the content of people (and the occasional automated bot[2]). Another reason why the data on Twitter may be interesting for companies is because people talk about brands nearly 20 percent of the time (Jansen et. al 2009, 3859).

Twitter was specifically chosen by us for its Streaming API and the amount of data this API gives every day. Furthermore, due to the size of Twitter big sample sizes are easy to obtain, which could yield interesting statistical significant results. On another note, the service is used by people from different countries[3]. This could mean that samples, taken on Twitter, have more variety[4]. For our study, we are interested in tweets that contain mentions as well as emoticons. This is because we want to determine the negative or positive sentiment (i.e. opinion) of Twitter users via emoticons.

To be more specific, this paper tries to answer the question if it is possible to get a positive or negative opinion from Twitter users about the users they mention via emoticons in their tweets[5]. We hypothesize that this is possible. In this paper we will give empirical support for our hypothesis in two parts. We refer to this hypothesis claim 1.

---

[1] Kazeniac, Andy (February 9, 2009). "Social Networks: Facebook Takes Over Top Spot, Twitter Climbs". Compete Pulse (blog of compete.com). Retrieved June 14, 2012. http://blog.compete.com/2009/02/09/facebook-myspace-twitter-social-network/

[2] We believe that this post will give you the general idea about what Twitter bots are. http://blog.stratepedia.org/2011/06/23/what-are-twitter-bots/

[3] This information was retrieved from the blog http://rossdawsonblog.com/weblog/archives/2012/02/which-countries-have-the-most-twitter-users-per-capita.html

[4] A problem in psychology, for instance, is that a lot of psychology students are tested in psychology research, which skews the data to young people. Furthermore, most of the time they are American. We have no hard proof for this, but we heard it from several psychology students and teachers on our own university.

[5] *An important assumption here is that an emoticon is directed towards the user that is mentioned. For example, the tweet "I like you :) @user" implies that the :) emoticon is directed towards @user, according to our assumption.*

The first part of this paper describes what we programmed with Python, Java and R. We needed to do this in order to collect, parse and process our Twitter data. Collecting data was done via Python and the Tweepy library, which implements the Twitter Streaming API. Our Java program searched and parsed all the tweets. These tweets contained emoticons and mentions to Twitter users. Finally, the data was partly processed with Java via the parsing program and, in a later stage, statistically processed with R.

The second part of this paper shows how users of Twitter, on average, use positive and negative emoticons at the people they mention. Furthermore, we will compare this result with previous results found in the field of sentiment analysis and positive psychology (see P/N ratio in the appendix). Moreover, we will discuss how our methodology (i.e. our programs) could indicate a general opinion about users who are mentioned in tweets. Throughout this paper, the perspectives from the fields of sentiment analysis and psychology in general will be used.

The comparison of our result with sentiment analysis is a finding of Bifet and Frank. One of the things they found was that the mean of the average relative frequency between positive and negative emoticons of all tweets on Twitter was 0.85 (Bifet and Frank 2010, 11). This means that 85% of all the emoticons used in tweets are positive. To support our first hypothesis, we furthermore hypothesize that our results align with the findings found by Bifet and Frank (2010, 11). With alignment we mean that our result should be near 0.85. Because of this we assume that Twitter users behave the same with their use of an emoticon regardless of whether they mention another Twitter user. We refer to this extension of the first hypothesis claim 2.

The comparison of our result with positive psychology is a finding of Fredrickson and Losada. They found that a P/N ratio in the range of 0.744 to 0.918 makes a community or individual flourish. In all other cases, the community or individual will languish. To extend our first hypothesis, we furthermore hypothesize that our result indicate that Twitter is a flourishing community in general (Fredrickson and Losada 2005, 678). We refer to this hypothesis as claim 3.

Besides looking at these 3 claims, this paper tries to uncover some properties (see section 4) related to the subject of mentioned users together with an emoticon in one tweet. The claims will be rejected or not rejected in our conclusion.[6]

1.1 Contributions

Firstly, the contributions of individual sections will be shown. Secondly, we present the unique contributions made by this paper. Finally, we will state to whom the results of this paper may be useful for.

1. We present an algorithm to collect emoticons and classify them into positive or negative emoticons. For this contribution read section 3.1.2 and 3.1.3.
2. A small discussion and experiment about how useful emoticons are as sentiment in general. For this contribution read section 3.3.
3. We explore how this method could be useful in determining an opinion about users that are mentioned. For this contribution read the results in section 4
4. We analyzed a small set of outliers to look for possible challenges for sentiment analysis to solve. For this contribution read section 4.5.
5. We attempt to link the results of the data to positive psychology. For this contribution read section read the results (section 4) and the conclusion (section 5).

[6] We use this phrasing because not rejecting a claim is subtly different than accepting a claim.

The unique contributions of this paper are contribution 2, 4 and 5. Contributions 1 and 3 have probably already been done before in some sort of fashion by researchers.

The interdisciplinarity of this paper provides some unique contributions to science in general. This is because a lot of papers in the field of sentiment analysis use only literature in computer science. Most of these researchers assume psychological related phenomena (see section 2). Furthermore, none of the papers that we found linked their findings to related findings in psychology.

Our analysis could be useful for heavy Twitter users who have an interest in knowing the opinion of the users on Twitter that mention them. Furthermore, it could be useful for some fields of science (e.g. psychology), to see their results linked to results in sentiment analysis.

1.2 Organization of this Paper

Section 0 and 1 contains the foreword and the introduction.

In section 2 the reader is able to find related work on the field of sentiment analysis and psychology, with an emphasis on positive psychology.

Section 3 contains the methodology of collecting, parsing and processing the data. It explains how a tweet is parsed, emoticons are recognized and evaluated as positive or negative. Furthermore, it also presents to what extent emoticons could be used as sentiment for short texts.

The results are found in section 4. Some of our results will be shown next to related results in sentiment analysis and positive psychology, properties of mentions and emoticons in general will be shown.

We conclude our results in section 5.

The discussion discusses the limitations of the results and methodology, it also discusses future work to be done (section 6).

The references are found in section 7. It is important to understand that this section only contains scientific sources (and school books or books from professors). All non-scientific sources throughout this paper are presented in footnotes.

In section 8 there are 4 appendices: a list of all the hard coded emoticons (appendix A), the 50 tweets which were used to manually classify positive or negative sentiment (appendix B), a glossary of all the terms which we find important to explain (appendix C) and further instructions to find all our source code, so that it is possible to repeat (or adapt) our experiment (appendix D).

2 Related Work

This section seeks to inform the reader about some related work, which are (zou kunnen, of is het echt zo? Anders, gebruik 'are' ipv 'could be') closely or distantly related to this paper. It is important to note that the field that is being studied here is predominantly sentiment analysis. However, due to its interdisciplinary nature, relevant work from the field of psychology is being looked at as well.

The field of sentiment analysis is a relatively new field. It progressed to such an extent that companies are now able to, for instance, predict: stock markets, brand sentiment and financial sentiment[7].

2.1 The field of sentiment analysis

[7] An example of such a company is www.sntmnt.com

Recently, British researchers of Bristol looked at a large scale analysis of social media content in order to discover macro-scale patterns in public opinion and sentiment. Their results showed that they could detect large scale events in the United Kingdom through their techniques. They could, for example, see the riots of summer 2011 that took place in various UK cities, back in their graphs (Landsdall-Welfare et. al 2012, 1).

Another example of sentiment analysis is detecting online service availability (e.g. of gmail and Facebook) via Twitter. A part of how they did this is by searching "via a small number of lexical features such as the phrase 'is down' and 'fail' has tags (e.g., '#gfail'), to signal an outage" (Motoyama et. al 2010, 1). That specific part of the article inspired our current methodology (searching for key features in a text).

With regards to detecting a specific emotion, a paper described that detecting sarcasm has proven to be not so successful via machine learning (González-Ibáñez 2011, 1). Furthermore, other researchers suggested the need for algorithms which could find such sentiments in a text (Read 2005, 48).

Finally, this paper shaped our methodology of evaluating emoticons in more detail. The authors decided to classify sentiment as positive, negative or neutral sentiment via the use of emoticons, which is almost the same approach as our paper (Pak & Paroubek 2010, 1321).

For further reading on related work in this field, the survey of Bo Pang and Lilian Lee gives a comprehensive overview of the field of opinion mining and sentiment analysis. It gives an overview of the applications, general challenges, classification and extraction, summarization of opinion mining and sentiment analysis. It furthermore states the broader applications, publicly available resources and it defines key terms used (Pang and Lee 2008)[8].

2.2 The field of psychology

Parts of researching happiness (and psychology in general) are related to researching sentiment. However, the papers that we read showed that sentiment analysis and psychology in general are not cited in the papers of the respective other field. There are some exceptions, such as González-Ibáñez (2011, 1).

In Buthan, the notion of happiness is so important that they have their own Gross National Happiness[9]. Buthan has the 8th highest subjective well-being and one can conclude that it is the only country in the top 20 happiest country that is poor (White 2007, 20). Even though this work is not directly related, it does show that our algorithm might contribute to better government policy.

Important work for our paper is the research finding about flourishing individuals and groups. Researchers from positive psychology state that a positivity ratio between 2.9 and 11.6 should be maintained (Fredrickson and Losada 2005, 684) in order for humans to flourish. This is the research finding that our results will be linked to.

Another important idea is to understand to what extent emoticons could be seen as emotions. Fortunately, research has been done in this area. A fMRI study suggests that "emoticons convey emotions without the cognition of faces" (Yuasa 2006, 1565). From a behavioral standpoint, another study shows that "to a large extent, emoticons serve the same functions as actual nonverbal behavior" (Derks, Bos, and Grumbkow 2008, 379).

Finally, according to a study in cyberpsychology and behavior emoticons define

---

[8] Special thanks go to Ivar Vermeulen, Christian Burgers, Peter Kerkhof, Wouter van Atteveldt of the course Marketingcommunicatie 2.0 at the VU University. They made us aware of the existence of this survey.

[9] http://www.grossnationalhappiness.com/

sentiment in a lot of cases (Lo 2008, 595). A very important implication of this article is that humans define sentiment different if an emoticon is in the message. For example, “no” might be considered as a neutral or negative message, whereas “no :)” might be considered a more positive message. This is interesting, since some computer scientists try to define sentiment via n-grams (i.e. words) and use tweets with emoticons only as their training set (see section 2.1). The article of Lo slightly suggests that text does not really define sentiment, or at least differently. In section 3.3 we conducted our own small experiment to see if we could refute the claim of Lo.

3 Methodology

This section shows our methodology of collecting data and parsing (section 3.1) and the processing of the data (section 3.2). It furthermore discusses the extent of how useful emoticons could be seen as sentiment in general (section 3.3).

3.1 Data Collection and Parsing

For our data collection program, we used the Twitter Streaming API to collect over 12.5 million tweets[10], the file size was 1.4 gigabytes of ASCII encoded text. The tweets were captured for a timespan of one week from 25-05-2012 to 01-06-2012. These tweets are a random sample of 1 percent of all tweets, sampled by Twitter itself[11]. This means that the data is representative for all tweets on Twitter. The data collection program[12] is programmed in Python.

With the collected tweets a dataset is made by our Java program of the following format:

- number of times the user is mentioned
- number of times the mentioned user occurred together with a emoticon in the same tweet
- the number of positive emoticons for the mentioned user
- the number of negative emoticons for the mentioned user

This data will be used to:

- determine the average relative frequency between positive and negative emoticons of mentioned users
- compare it to the relative frequency between positive and negative emoticons of tweets on Twitter in general thanks to Bifet and Frank (2010, 11)
- compare it to the P/N ratio (Fredrickson and Losada 2005, 678)
- understand why certain people are having an extreme relative frequency between positive and negative emoticons

This will be presented in the results section, as of now, we will look at the algorithms programmed by us, since these algorithms are part of the methodology.

---

[10] The real number is 12,753,921. However, the program counts the number of new lines, not the number of tweets. But most tweets (we assume >98%) take one line. The lowest bound we calculated could be 6 million tweets, but it is a lot more likely that the real number is higher than 12.5 million, because when we manually looked into the log we rarely saw a tweet taking up more new lines than one.

[11] Everything is explained at https://dev.twitter.com/docs/api/1/get/statuses/sample. If you log in with your Twitter account on Internet Explorer on the link that is on that page, which is https://stream.twitter.com/1/statuses/sample.json, then the whole Twitter stream appears in front of you in JSON format.

[12] 21 Recipes for Mining Twitter by Matthew A. Russell, O’Reilly Media, recipe 8 was used and the last line was adapted to capture everything that the normal access level of the Twitter Streaming API provides.

3.1.1 The data collection program

The data collection program is a python script, this script is able to access the Twitter Streaming API. The script was intended to filter on specific subjects, such as words or users. The change that we made to this script is to change the last line from (pseudo-code is in quotation marks):
"only receive tweets which passes the filter based on a specific set of parameters"
to
"don't filter anything, receive every tweet that you give to us from the Streaming API".

3.1.2 The algorithm for parsing the retrieved Twitter data

The purpose of this section is to explain the general idea of parsing retrieved Twitter data.

The algorithm to parse the tweets, which are contained in a text file is written in pseudo-code. The following style conventions are used in the pseudo-code.

Flow control constructs, such as a **<u>while</u>** loop are bold and underlined.
Commands, such as **do something** are bold.
Variables are in a format with hyphens and are self-explanatory, such as this-variable-does-nothing.
Italic words are variables saved for output.
CAPSLOCKED WORDS are header titles for the output.

<u>The algorithm for parsing the retrieved Twitter data</u>

open text file
scan text file line for line

<u>while</u> the scanner did not reach the end of file

  <u>if</u> the tweet contains a mentioned user
    <u>then</u> save the username of the mentioned user as a lower-case unique key if possible

    retrieve the number-of-times-this-user-is-mentioned-counter
    increment this counter by one
    save this counter as the *first value*

    <u>if</u> the tweet contains one or more emoticons
      <u>then</u> retrieve the mentioned-user-hasEmoticons-counter of the key
      increment this counter by one
      save this counter as the *second value*

```
                    evaluate whether the emoticons are positive or negative
                    retrieve the-counter-of-the-positive-emoticons
                    increment that counter by the number of emoticons that are
                    evaluated as positive
                    save this counter as the third value

                    retrieve the-counter-of-the-negative-emoticons
                    increment that counter by the number of emoticons that are
                    evaluated as negative
                    save this counter as the fourth value

                    put the key value pairs (4 values per key) in a data structure
                    that stores key value pairs as <KEY, VALUE>

            end if
      end if
end while

save to text file in the following format:
      (first value) (second "")           (third "")    (fourth "")
KEY   #TWEETS   #TWEETS_AND_SMILIES   #POS   #NEG
```

Explanation of the format:

- The key stands for the user name that is mentioned in all the tweets it is mentioned in
- The #tweets stands for the number of tweets that the user is mentioned in without emoticons
- The #tweets_and_smilies stands for the number of tweets that the user is mentioned in and a specific emoticon was detected in the tweets as well
- The #pos stands for the amount of positive emoticons that were found in tweets in combination with this username
- The #neg stands for the amount of negative emoticons that were found in tweets in combination with this username

The keys are only saved in lowercase because a Twitter username is not case sensitive, but it can be displayed case sensitive.

3.1.3 How emoticons were evaluated

<u>Determining what emoticons to search for</u>

Due to the courtesy of infochimps, we obtained a list which displayed how many times an emoticon was used (i.e. the frequency) on Twitter from March 2006 to November 2009. "The emoticon data comes from analysis on the full set of tweets during that time period, which is 35 million users, over 500 million tweets, and more than 1 billion relationships between users[13]".

Since all the tweets from March 2006 to November 2009 were analyzed we assume that the behaviour of the population has not changed much 32 months later. This is because although the behaviour of the population could have changed a little - 3 years later - we do not have enough reason to believe that the population is a lot different in its behaviour than three years ago. So we assume that the use of emoticons of Twitter users have roughly stayed the same. Because of this we assume that we have an emoticon frequency count of every tweet on a scale of the current population, thanks to the dataset of infochimps. This makes it possible to hard code the most used emoticons in our Java classes.

In total, the file from infochimps contains about 1500 unique emoticons. Frequency analysis in the program R showed that 75 of the most used emoticons accounted for 99% of the total usage of all emoticons combined. Furthermore, it follows Zipfs law as is shown in the following graph.

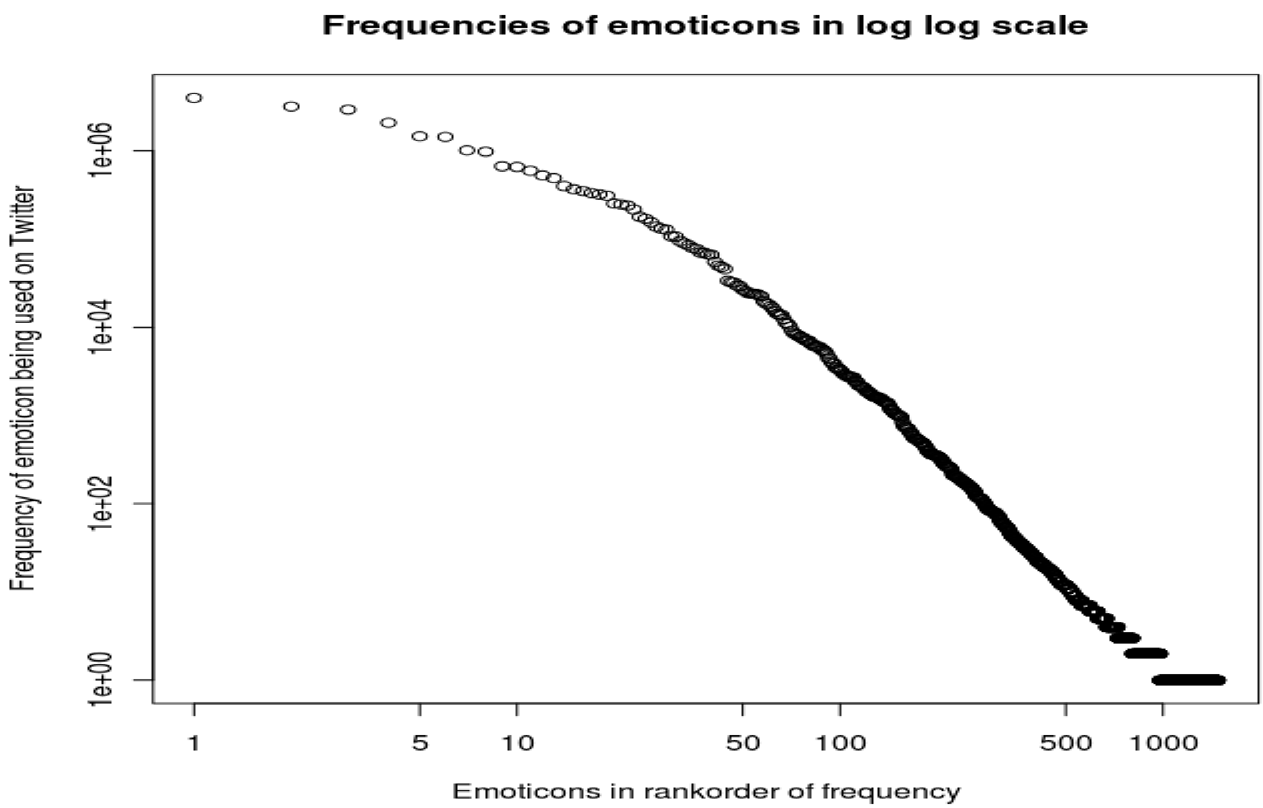

Figure 1 - Frequencies of emoticons in log log scale: all 1500 emoticons have been plotted, which is what the x-axis represent. The y-axis represents the frequency of how many times it is being used. For the curious minded the :) emoticon is on the first place, the :d emoticon is on the second place and the :(emoticon is on the third place.

The classification of whether an emoticon is positive or negative was done by ourselves. A list of all the emoticons that we searched for is in the appendix. For some of these emoticons it is open for discussion to what extent they are positive or negative. We did our best to divide them as

[13] Obtained from http://www.infochimps.com/datasets/twitter-census-conversation-metrics-one-year-of-urls-hashtags-sm--3

clearly as possible.

Evaluating the emoticons
With two hard coded lists of emoticons it became fairly straightforward to evaluate whether emoticons are positive or negative. In every tweet our Java program would check if the strings would match one of our emoticons. If this was the case, then a counter would be incremented, as explained in 3.1.2.

3.1.4 Encountered limitations of the algorithm that we programmed

Hard coded emoticons
Without machine learning (or related) algorithms it is impossible to recognize all emoticons. Firstly, this is because it is not clear whether an emoticon is meant as an emotion or is part of a text. Examples of what we mean are ":others", which could be mistaken for the :o emoticon or "(= some explanation)", which could be mistaken for the (= emoticon. As a result such emoticons were excluded. Secondly, since emoticons are hard coded in our program, not all emoticons contained in the tweets on Twitter will be seen by our program.

The way tweets are recognized
Our python script saves more than just tweets. It also contains the username which created the tweet, date created and the source where the tweet is sent from. However, it does not contain metadata, so our Java program has to recognize a tweet. For our program we defined a tweet as whenever there is a mention in a line of text.

This brings one limitation, which is that usernames that are mentioned could turn out to be false. This is because it is possible on Twitter to type in a username and add some characters after it. Twitter recognizes which part of the string is a username, our program does not.

General bugs
In our result set we saw an entry of the username @. This cannot be a username, so we deleted it from our result set.

3.2 Data Processing

With the results from our Java program we calculated the following with R:

- the mean relative frequency between positive and negative emoticons
- the relative frequency of a username being only mentioned or being mentioned and have at least one emoticon in one of the tweets it was mentioned
- Plots of:
  - The frequency of being mentioned as a user on Twitter with or without emoticon (section 4.2)
  - The relative frequency that users are mentioned with an emoticon within the same tweet. For example, Justin Bieber is mentioned around 15 percent with an emoticon and 85 percent without one (section 4.3)
  - The distribution of the relative frequency between positive and negative emoticons of mentioned users (section 4.4), note that the mean value of this graph is in section 4.1

3.3 The extent of how useful emoticons are as sentiment in general

When we look at the research paper that started the idea of using classified emoticons to approximate classified sentiment in general (Read 2005, 45), we wondered ourselves the following question: to what extent could emoticons be used to label a short text as sentiment? This question was raised, because sometimes tweets could contain sarcasm or other forms of mixed sentiment, which an emoticon would not represent (see examples below). Furthermore, the methodology described in section 3.1 and 3.2 would not make sense if this assumption turned out to be mostly false. Moreover, Read did not give any scientific sources or explanation for why he made that assumption. Other papers which we read on this subject are from Go et. al (2009) and Pak and Paroubek (2005), they did the same as Read.

<u>Examples of tweets with wrong classified sentiment</u>
To humans, this is positive tweet because of the word "LOL". In other papers this would be classified as negative sentiment by computer programs and used as training data for machine learning algorithms.

> RT @GG_sargeant: Scar "If they were a cheese cake they would be the biscuit the lowest of the low" Me"but I like the biscuit :(" LOL

A mixed tweet, because the :(only applies to "But we just opened". In other papers this would be classified as negative sentiment as training data for machine learning algorithms.

> @KennedyAmor righttttttt lol ! But we just opened :(

An interesting note is that the abbreviation LOL could be considered as an emoticon as well.

3.3.1 Emoticons are emotions and indicate more sentiment than text alone

According to different researchers emoticons serve almost the same function as emotions (Derks, Bos, and Grumbkow 2008, 379) (Yuasa 2006, 1565). We suspect that facial expressions in face to face interaction define the sentiment of a message more than the words themself do. However, we could not find this in the literature.

Nevertheless, it is interesting to note that a lot of tweets contain a different sentiment if an emoticon is added to the text (Lo 2008, 595). This could mean that emoticons explain sentiment of a tweet at a more fundamental level than words.

Because of this we conclude that it is harder to define sentiment via words than via emoticons. So according to us training classifiers on the classification of emoticons alone does make sense. To empirically support this idea we carried out an experiment, because as stated before, the literature is sparse on this question, and no paper that we read mentioned the paper of Lo (2008).

3.3.2 Our Experiment

We looked at 25 tweets which contained the :) emoticon and at 25 tweets which contained the :(emoticon. This is because these emoticons are the most clearest and most used emoticons in their meaning and they lie closest to the division we made with our main experiment (positive sentiment versus negative sentiment). The tweets could also contain other emoticons by chance. All the 50 tweets are in the appendix B.

The way we collected the data was that we opened our 1.4 gigabyte text file, typed in ctrl + f and searched for the string ":)". The first 25 tweets that we understood (no matter the language) were saved for further evaluation. The same method applies to the negative emoticon ":(".

The tweets that contained the :) emoticon were all classified as having positive sentiment. The tweets that contained the :(emoticon had 22 tweets which were classified as having negative sentiment, two as having mixed sentiment (i.e. positive and negative) and one as positive sentiment. The three tweets that were noisy are in section 6.2 given as examples.

We performed a sign-test to see whether the 25 tweets with negative emoticons could at least classify sentiment above chance level. This was the case with $p < 0.001$[14]. It is obvious then that the same applies towards the 25 tweets which were containing the positive emoticons.

We found that emoticons classify positive and negative sentiments above chance level. Of course, more research could be done to what extent emoticons classify tweets.

4 Results

Initial analysis showed that there are a little over 4.4 million usernames that are mentioned in tweets. 650,000 of these usernames had at least one tweet which contained an emoticon. On average, there were 1.4 emoticons per tweet. It is important to note that our set of results only contained users that had at least one emoticon, either classified as positive or classified as negative. This means that the result set contained 650,000 usernames.

4.1 Results with regards to the mean

The mean of the average relative frequency between positive and negative emoticons of mentioned users will be listed next to 3 other relevant means. The results are in table 1.

Our result: Mean of average relative frequency between positive and negative emoticons of mentioned users[15] (1 = completely positive, 0 = completely negative)	0.834
P/N Ratio (Fredrickson and Losada 2005, 678) P/N Ratio converted to percent	2.9 to 11.6 0.744 to 0.918
Mean of average relative frequency between positive and negative emoticons of all tweets on Twitter (Bifet and Frank 2010, 11)	0.85
Mean of average relative frequency between positive and negative emoticons from the file supplied by infochimps (see section 3.1.3)	0.865

Table 1 - comparison between means: since the first, the third and the fourth mean are from Twitter, and are calculated from millions of tweets, the differences between these means are statistically significant

[14] You could do this yourself at http://www.fon.hum.uva.nl/Service/Statistics/Sign_Test.html for n+ type in 22 and for n- type in 3

[15] derived from 12.5 million tweets

The figures in the table do not differ much from each other. The third and the fourth mean are means from the same type of data, which is the relative frequency between positive and negative emoticons over all tweets collected. Note that there still is a 0.015 difference.

4.2 Mentions on Twitter and their indegree distribution

Next, we will look at the indegree of Twitter usernames who are mentioned. What is meant by indegree is, for example, if Bob was only mentioned by Alan, then Bob has an indegree of one. As is observable from figure 2, the indegree looks like it is scale-free which means that there are hubs. This furthermore means the distribution of mentions follows Zipfs law. However, the ten most mentioned usernames do not follow the pattern of being scale free.

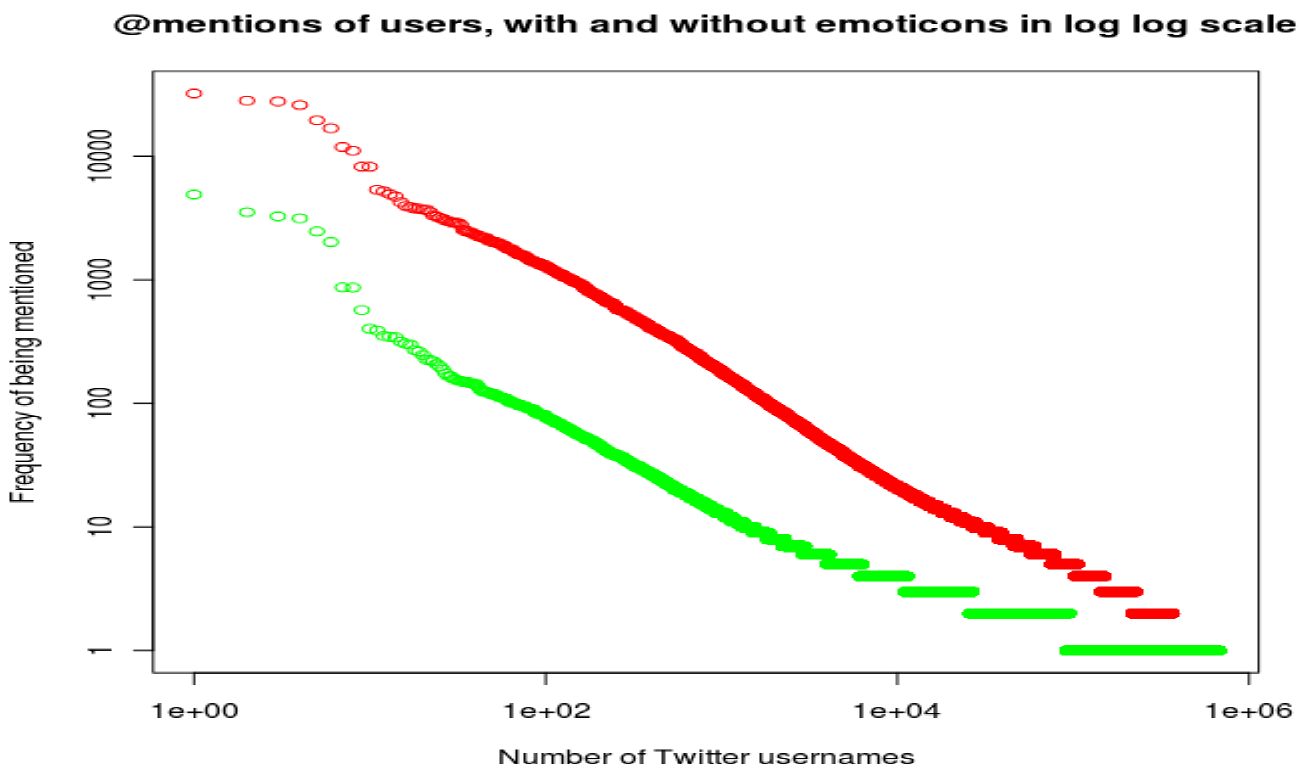

Figure 2 - usernames that are mentioned: green is the line that represents mentioned usernames that also contain emoticons in the tweets which they have been mentioned in and red represents all mentions.

In figure 2, the two lines almost follows the same pattern, except for the tails. There are over 650,000 data points processed for this graph.

4.3 Percentage of times that users are mentioned with an emoticon

On enough occasions, users will mention a person in their tweet but not type an emoticon in the same tweet. Figure 3 will show the relative frequency of how many times this is the case.

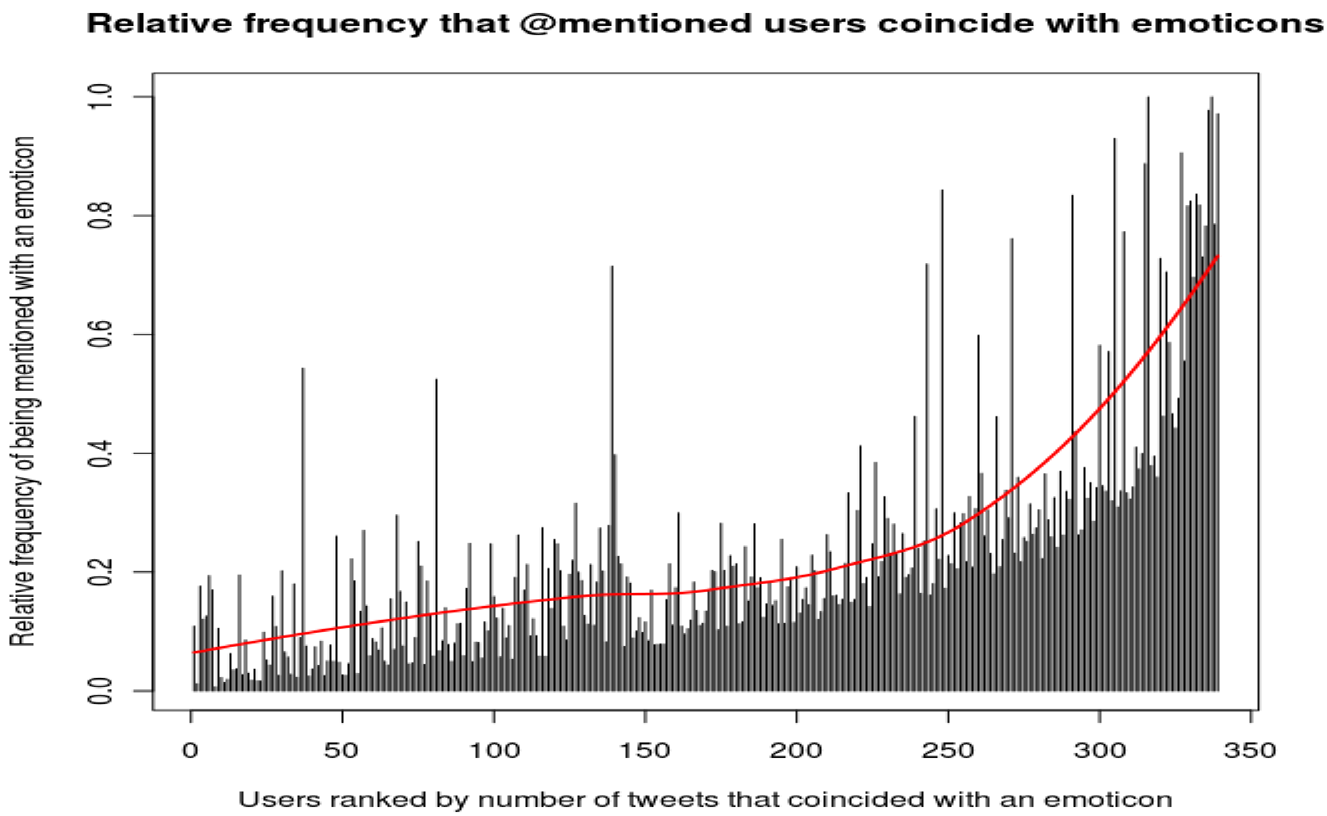

Figure 3 - Relative frequency that @mentioned users coincide with emoticons: the y-axis of this graph displays the (relative) times a user is mentioned together with an emoticon in a tweet. The x-axis represents the users themself, they are ranked on the total amount of mentions they have gotten in total. The red line represents the general trend done with the predict(loess(y~x)) function in R.

The trend in the graph of figure 3 goes up. However, the graph also shows a lot of exceptions to this trend. There are, for example, some bars which are a lot higher than their surrounding bars.

The general conclusion is that the number of mentions a user has and the relative frequency of emoticons in a mentioned tweet are negatively correlated. So for example, if person A is mentioned thirty times (lower right corner of the graph), then he probably has a higher relative frequency of emoticons when he is mentioned in a tweet compared to someone who is mentioned fifty thousand times (lower left corner of the graph).

4.4 Distribution of emoticon sentiment among mentioned users

We already showed that our average relative frequency between positive and negative emoticons of mentioned users is 0.834. Initially, we thought this was normally distributed. To test this we plotted a density histogram (figure 4). We only used keys from users that had 30 or more mentions that coincided with an emoticon in a tweet. We did this in order to ensure the central limit theorem[16] would be applicable.

---

[16] A lot of students have trouble with this theorem. If you do not know what it is we encourage you to watch http://www.khanacademy.org/math/statistics/v/central-limit-theorem, it is explained there in a wonderful way.

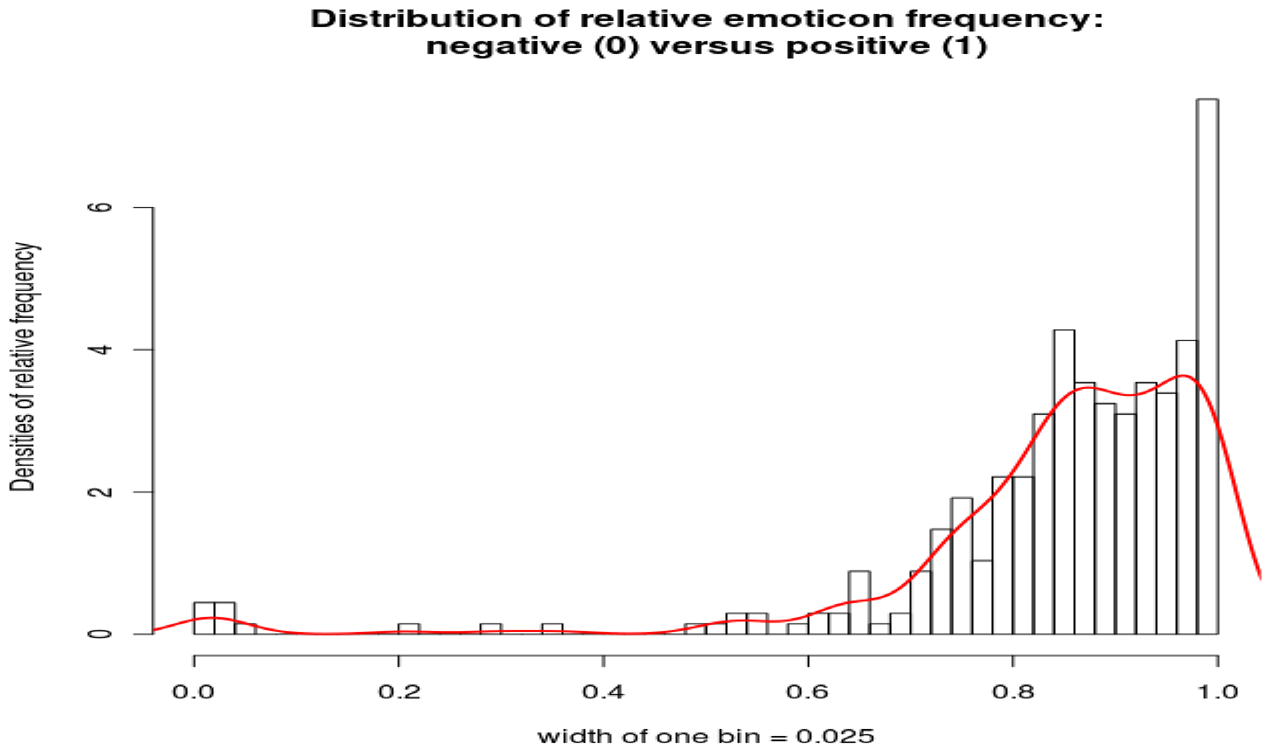

Figure 4 - the y-axis represents the density of the relative frequency and the x-axis shows the bins of the histogram. There are 40 bins in total and 339 data points have been used to draw this histogram. The red line is the trend line done with the density() function in R. It was not possible for us to draw the trend line with a normal (frequency based) histogram.

The histogram in figure 4 shows that the relative frequencies among the 339 highest mentioned users with emoticons are not normally distributed.

4.5 Examining the outliers - a few case studies

In the data we found three types of outliers:

- a low relative frequency between positive and negative emoticons (x > 0.99)
- a high relative frequency between positive and negative emoticons (x > 0.99)
- a high percentage to coincide with an emoticon if the username is mentioned

Note that x represents the value of the Twitter user in our data set.

We are going to discuss the first two outliers in a case study for which one example is taken. We do this in order to see what possible explanations there could be for such outliers.

4.5.1 Case Study 1 - one Twitter account that had a relative frequency lower than 0.01

When we filtered Twitter users that had a relative frequency lower than 0.01 only one username returned with a relative frequency of 0.007. This username has 2 emoticons, which were classified as positive, and 197 emoticons, which were classified as negative.

This could be a quite shocking result for the user. So we investigated why this is and read every tweet of him where he was mentioned in our result set.

It turns out that a user account called Louis_Tomlinson (a member of the boyband

OneDirection) had a bad experience with this Twitter user. The Twitter user represented an Italian restaurant called trecolori47. In our result set we saw that a lot of fans of Louis_Tomlinson retweeted his message which contained two :(emoticons. The message is in figure 5. As one can infer having 15818 retweets of this message will also show in a sample that is roughly one percent of that size (i.e. it will be a noticeable event in our result set).

Figure 5 - The tweet of Louis_Tomlinson, talking about trecolori47: a tweet that could hurt your business.

It is interesting to note that this person has almost 4 million followers. The lesson that we learned through this particular case study is: never treat your opinion leaders badly.

In marketing literature there is a concept known which is called opinion leaders. An opinion leader is "a person within a reference group who, because of special skills, knowledge, personality, or other characteristics, exerts social influence on others." (Kotler 2010, 165)

4.5.2 Case Study 2 - one Twitter account that had a relative frequency higher than 0.98

It is interesting to note that when we filtered with R to show results with a relative frequency higher than 0.99, 2 usernames returned out of the 650,000. The languages that these users spoke in were, however, not English or Dutch. So there was no way for us to understand their tweets. This is why we filtered on the relative frequency of 0.98. Still, less than 10 results returned.

Admittedly, we chose the username that we liked the most and since we would like to go to the city Toronto we chose the username 1DTorontooo out of the nine usernames. This should not be of any influence, however, because it is a case study, not a random sample that is an argument to the generalization of a theory.

So the username 1DTorontooo had a relative frequency above 0.98. We looked in the result set and the Twitter profile of the user and we saw that the part "1d" stands for the band OneDirection. It is interesting to observe that two outliers are related to the same boyband.

What happened is that 1DTorontooo tweeted that she found a website, which had a phone conversation of Lous_Tomlinson and the restaurant Tre Colori. Louis_Tomlinson tweeted to her "that is fake babe :)" and that tweeted message was retweeted thousands of times, see figure 6 for the full conversation.

Initially it seems like a rare coincidence that two case studies go about the same band, even if we would have read all the usernames their stories and picked this one. But we would like to point out that Twitter is a medium which focuses on actuality. So in a sense, this could be more normal than one might expect.

To conclude, this refers back to the marketing literature. The same concept (i.e. that of opinion leaders) applies, but now it has a positive sentiment, as opposed to a negative one for the restaurant. So the user 1DTorontooo would be likely to experience reputation gain, instead of reputation loss.

Figure 6 - The tweet of Alexa (1DTorontooo) being replied to by Lous Tomlinson: a tweet that will give Alexa more followers.

5 Conclusion

Microblogging services like Twitter became a lot more popular since they started. 340 million tweets are sent every day by 140 million users[17]. The data Twitter generates on a daily basis makes it attractive for sentiment analysis.

In our paper, we presented two algorithms: one to parse tweets and one to evaluate the emoticons as positive or negative. We presented the limitations of the algorithms as well.

The results gathered by our programs, indicate that from a psychological point of view Twitter is a flourishing community on average. This is because the mean that we obtained (0.834, section 4.1) is in the range of P/N ratio, which is 0.744 to 0.918. The mean furthermore indicates that users are generally positive towards each other. Moreover, Bifet and Frank (2010, 11) found that Twitter users use a lot more positive emoticons than negative in general (0.85, section 4.1). So Bifet and Frank (2010) indicated that Twitter users are quite positive in general.

All the results, including our own experiment, indicate that they have no big differences with each other, therefore claim 2 and 3 are not rejected (see section 1). On another note, it might mean that the limitations of our algorithm do not significantly impact our results. From that perspective, our algorithm works quite well, despite its limitations. This means that our algorithm is able to determine positive or negative sentiment in a relevant way about Twitter users, to a big extent. Thus, claim 1 is not rejected.

On another note, we showed some properties related to users who mention other users via emoticons (mainly section 4.2 to 4.5), which was mainly done to uncover some properties of the subject of this paper.

[17] http://blog.twitter.com/2012/03/twitter-turns-six.html

Finally, we conclude that our algorithm works well enough and is useful to the people who intend to use this form of sentiment analysis. The users that will benefit the most are users that tend to be quite popular, but not too popular (section 4.3). These users on Twitter will be mentioned together with an emoticon a lot of times and therefore could see their own specific relative frequency as an approximation of how people feel about them.

6 Discussion and Future Work

In the discussion we will discuss our results, the limitations of our research in general and future work.

6.1 Discussion of the results

The means in section 4.1 are significantly different. This might be because of different methodologies. It is interesting to observe that the values lie close to each other.

All the means are within the interval proposed by Fredrickson and Losada. If this was not the case, then future research should have been done because it would implicitly falsify the P/N ratio of Fredrickson and Losada. However, it has to be questioned: to what extent is the P/N ratio the same as a relative frequency between positive and negative emoticons? We are fairly confident that they are roughly the same, but since we are not psychologists we are not totally sure.

The results that we found with the percentage of times that users are mentioned with an emoticon are surprising. We did not expect to find that Twitter users who are quite actively mentioned (around thirty times), are mentioned a lot more times with emoticons than famous users. This is surprising because in reality these users are mentioned, on average, 3000 times (since our dataset is 1% of all the data of the week we captured). It might be that a specific segment of people are mentioning these users. Such a segment would mention these users in a more informal manner, because they are all more or less the same kind of person.

6.2 Limitations of our research

First of all, since we were fairly new to these types of experiments and with programming with Python in this way, we could not capture tweets in UTF-8 encoding, but in ASCII encoding. This means that the sample size is skewed towards Twitter users who do not tweet in languages that rely on any UTF-8 encoding. These tweets did not come into our data set, due to an exception made by Python.

We had one difficulty with data processing. Originally we captured a little over 5 gigabytes of data. However, the other text files would not open and there were some programming difficulties. At the time we were not well versed enough with the command-line of linux, so it could not use these other data sources. We had difficulties copying and pasting the text in another text file, because with such a huge amount of data, programs stall and crash. Furthermore, the time limit was a bit short to fix our Java program in such a way that it could read all the text files, so we decided to leave that as future work.

Another limitation is that everyone uses emoticons in tweets whenever they like. For instance, a Twitter account which has to use formal language would use less emoticons than an account that is allowed to be informal. This means that the representivity of the data is skewed even more because we do not differentiate between formal and informal Twitter users.

With regards in our literature study, we found some papers almost after this paper was finished. An example is the paper of Lo (Lo 2008). Before we knew that this paper existed we already carried out our experiment in our methodology, which is related to his paper. If we knew the existence of his paper earlier, then we would not have carried out this experiment because it would be less important.

Our biggest limitation are small errors in evaluating (hard coded) emoticons, which were detected too late after everything was finished. A few of those errors are: we do not have the ;-) emoticon in the positive emoticon set and we have the [: emoticon in the positive and the negative set. This is observable in appendix A.

Limitations of the 2 algorithms are in section 3.1.4. These were put there to fully explain the algorithm.

6.3 Future Work

A lot of difficulties that we discovered are eligible subjects for future work. There are a few subjects which have not been discussed in this paper, but are also important for future research, these subjects will be discussed here.

First of all, everything that we read did not contain any attempt to answer how computers should detect directionality of classified sentiment. For example, the message *"I love myself"* has a different directionality than *"I love you."* We found another example of a different kind in our data set which was: *"@Louis_Tomlinson @trecolori47 I am sry:(but how was the rest of your day"* Future research should be in finding a definition of what directionality means and how machine learning algorithms could detect these directionalities.

Another future research question that we would like to emphasize is to what extent are words in tweets important to classify a sentiment? The reason we want to emphasize this is because in some cases everyone can be wrong in evaluating the sentiment of a piece of text. For instance, if someone says *"I really love this teacher!"*, then everyone would believe this to be positive. However, what if this person is saying this to a friend and he knows how much this teacher is hated by everyone? An outsider without this context could not classify this message correctly because of missing information, the right context. It begs the question to what extent could computers be possibly right by classifying the sentiment of a piece of text, provided they have the right context?

These two subjects are, according to us, important future research topics for future psychologists and sentiment analysts.

7 References

Bifet, A., and E. Frank. 2010. “Sentiment Knowledge Discovery in Twitter Streaming Data.” In Discovery Science, 1–15.

Compton, William C. 2005. An introduction to positive psychology. Australia; Belmont, CA: Thomson/Wadsworth.

Derks, D., A. E. R. Bos, and J. von Grumbkow. 2007. “Emoticons and Online Message Interpretation.” Social Science Computer Review 26 (3) (December 3): 379–388. doi:10.1177/0894439307311611.

Fischer, Agneta. 2010. De zin en onzin van emoties. Amsterdam: Bakker.

Fredrickson, B. L., and M. F. Losada. 2005. “Positive Affect and the Complex Dynamics of Human Flourishing.” American Psychologist 60 (7): 678.

Go, A., L. Huang, and R. Bhayani. 2008. “Twitter Sentiment Analysis.” Final Projects from CS224N for Spring 2009.

González-Ibáñez, R., S. Muresan, and N. Wacholder. 2011. “Identifying Sarcasm in Twitter: a Closer Look.” In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Portland, Oregon.

Jansen, Bernard J., Mimi Zhang, Kate Sobel, and Abdur Chowdury. 2009. “Micro-blogging as Online Word of Mouth Branding.” In , 3859. ACM Press. doi:10.1145/1520340.1520584. http://portal.acm.org/citation.cfm?doid=1520340.1520584.

Kotler, Philip. 2010. Principles of marketing with mymarketinglab pack. [S.l.]: Pearson Education Ltd.

Lansdall-Welfare, Thomas, Vasileios Lampos, and Nello Cristianini. 2012. “Effects of the Recession on Public Mood in the UK.” In , 1221. ACM Press. doi:10.1145/2187980.2188264. http://dl.acm.org/citation.cfm?doid=2187980.2188264.

Lo, Shao-Kang. 2008. “The Nonverbal Communication Functions of Emoticons in Computer-Mediated Communication.” CyberPsychology & Behavior 11 (5) (October): 595–597. doi:10.1089/cpb.2007.0132.

Mehta, Rushabh, Dhaval Mehta, Disha Chheda, Charmi Shah, and Pramila Chawan. 2012. “Sentiment Analysis and Influence Tracking Using Twitter.” International Journal of Advanced Research in Computer Science and Electronics Engineering (IJARCSEE) 1 (2). http://ijarcsee.org/index.php/IJARCSEE/article/view/32.

Motoyama, Marti, Brendan Meeder, Kirill Levchenko, Geoffrey M. Voelker, and Stefan Savage. 2010. “Measuring Online Service Availability Using Twitter.” In Proceedings of the 3rd Conference on Online Social Networks, 13–13. WOSN’10. Berkeley, CA, USA: USENIX Association. http://dl.acm.org/citation.cfm?id=1863190.1863203.

Myers, David G. 2010. Social psychology. London; Boston: McGraw-Hill Higher Education.

Pak, A., and P. Paroubek. 2010. “Twitter as a Corpus for Sentiment Analysis and Opinion Mining.” In Proceedings of LREC. Vol. 2010.

Pang, Bo, and Lillian Lee. 2008. “Opinion Mining and Sentiment Analysis.” Foundations and Trends® in Information Retrieval 2 (1–2): 1–135. doi:10.1561/1500000011.

Read, J. 2005. “Using Emoticons to Reduce Dependency in Machine Learning Techniques for Sentiment Classification.” In Proceedings of the ACL Student Research Workshop, 43–48.

White, Adrian G. 2007. “A Global Projection of Subjective Well-being: A Challenge to Positive Psychology?” Psychtalk 56: 17–20.

Yuasa, Masahide, Keiichi Saito, and Naoki Mukawa. 2006. “Emoticons Convey Emotions Without Cognition of Faces.” In , 1565. ACM Press. doi:10.1145/1125451.1125737. http://dl.acm.org/citation.cfm?doid=1125451.1125737.

8 Appendices

Appendix A - Emoticons used

```
private final static String[] POSITIVE =
    {
    ":)", ":d", ";)",
    ":-)", "=)", ":p",
    "(:", "xd", "=d",
    "=]", ":]", ":o",
    "d:", ";d", "=p",
    ":-d", "^_^", ":o)",
    ":-p",
    //o: is omitted
    ";]", ";p", "(;",
    "=o", ";o)", "[:",
    //"(=" is omitted
    ":}", ":-o",
    ";o", ";-d",
    "d;",
    //d= is omitted
    "[=", "(^_^)", "(-:",
    "=}", "[;",
    //p: is omitted
    "o;",
    //:od is omitted
    ":-]",
    //:op is omitted
    ";}", ":-}",
    //do: is omitted
    ":=)", "{:", ":*)",
    ";-]", "(-;", "^_^.",
    "o(^_^)o", "*^_^*", ";=)",
    ":P"
```

Figure A - positive emoticons: Some emoticons were omitted because we assume that they could also be written by users as natural language a lot, for example, ":o" as in "... :others ..." or "(=" as in "(= some explanation)".

```java
private final static String[] NEGATIVE =
    {
    ":(", ":-(", ":/",
    "=(", "=/",
    "):", ":|", ":-/",
    ":[", ":@",
    ";(", ";/", ":o(",
    "/:", "=|", "[:",
    ");", ";-(", "|:",
    ":-|", "@:", ";@",
    ":{", ")=", "(^_^;)",
    ":*(", ";[", "(-_-)",
    ":o/", "(t_t)", ":-[",
    ":,(", "/=", "]:",
    "}:", ";-/", "=@",
    ";|", "={", ":-{",
    "/;"
    };
```

Figure B - negative emoticons: Some emoticons were omitted because we assume that they could also be written by users as natural language a lot, for example, ":o" as in "... :others ..." or "(=" as in "(= some explanation)".

Appendix B - Experiment 1 - 25 positive and 25 negative tweets and their classified sentiment

Positive :)

Text From Tweet	Sentiment By Human Interpretation
1. @iamalyy YaWelcome >:)< :) SayrealPlns 2012-05-25 14:05:10 Twitter for iPhone	positive
2. @jo_beth18 @Dj_Scooby @PrfctGntlmanO_o @SweetnothingSam :) THANKS JO !!!! <3 SharmaneJW 2012-05-25 14:05:10 web	positive
3. I run threw more trees than a koala her :) UhO_Troublee__ 2012-05-25 14:05:11 Twitter for Android	positive
4. :) LeynaW 2012-05-25 14:05:11 TweetCaster for Android	positive
5. @IvanSwarley Feliz dia del #OrgulloFriki !! :) jsp_993 2012-05-25 14:05:11 web	positive
6. @adamobeaudoin lend me your car for when you're gone :) jesscolatorti 2012-05-25 14:05:11 Twitter for iPhone	positive
7. 6 hours of work, then the lake for the weekend! :) DestinyHatley 2012-05-25 14:05:12 Twitter for Android	positive
8. @jilanah Jil! :) SuupermanYagurt 2012-05-25 14:05:12 web	positive
9. @ENTERSHIKARI What about playing all the songs? Nah, just kidding, it's gonna be great, no matter what songs you play :) mynameislorena 2012-05-25 14:05:15 Twitter for Android	positive
10. RT @XSTROLOGY: It is #Taurus nature to be a BOSS. But they are loved by their people anyway :) ekamiyaa 2012-05-25 14:05:17 Write Longer	positive
11. @BNEDirectioners PhilipPAYNEs ..uhm. i mean Philippines :)) Kk_1dibble 2012-05-25 14:05:18 web	positive
12. Thuis, was echt leuk vandaag :) Rosa5o1 2012-05-25 14:05:19 Twitter for iPad	positive
13. @Zara_ElOuardi1D :)! reach_my_dreams 2012-05-25 14:05:22 web	positive
14. @iamNinjabi welcome aboard :) sEeR4t 2012-05-25 14:05:22 Twitter for iPhone	positive
15. @PaulHUFC get your freddy mercury on then :) #iwanttobreakfree Mikeasaurus85 2012-05-25 14:05:25 Twitter for iPhone	positive
16. @stylerfnbcgx0 thank you! :) arahlovesyou 2012-05-25 14:05:27 web	positive
17. @WETthatNIPPLE I was jussssst about to text you :P but Okaaay duckling :) _xShawdiiMaac 2012-05-25 14:05:30 Twitter for	positive

Android	
18. RT @OfficialYves: RETWEET kung excited kana sa #BuildABigDreamConcert :) stevenasotal 2012-05-25 14:05:30 web	positive
19. ayeee that's good so far! :) RT @JaniseeAileen 34 names so far on this list. HannanShine 2012-05-25 14:05:31 web	positive
20. @VivianCat4 @irishlovesyou @nixcelleontoria :) don't let me make u sleep late again tonight vivi! chalen89 2012-05-25 14:05:31 web	positive
21. :) FeelLikeeBuzzin 2012-05-25 14:05:34 Twitter for Android	positive
22. RT @SamiYusuf: Blessed Friday all :) Groover19 2012-05-25 14:05:34 web	positive
23. 3,500th tweet :) #brittney NIRA_sistable 2012-05-25 14:05:37 Twitter for Android	positive
24. @Becs_avfc :) don't I know...always with the teasing! Haha not long left till beer time. :) DanielAVFCRees 2012-05-25 14:05:40 Twitter for iPhone	positive
25. @sunnyDELIGHT92 next Friday , june 1 :)) Chelsea_Dionne 2012-05-25 14:05:40 TweetCaster for Android	positive

Table A - 25 tweets that were selected for the :) emoticon: this was retrieved from our dataset, for more information email the author.

Negative :(

Text From Tweets	Sentiment By Human Interpretation
1. @KennedyAmor righttttttt lol ! But we just opened :(_ohhessjayy 2012-05-25 14:05:10 Twitter for iPhone	**positive and negative**
2. @thebumpiestpath @gypsi001 very dry. :(yes crisp. Tell me you have a trick up your sleeve! ;) http://t.co/Dp9Ckixh SaidByJeannie 2012-05-25 14:05:12 TweetCaster for Android	**positive and negative**
3. I always get a #Goodmorning text! But today I didn't :(wtf? Catt5268 2012-05-25 14:05:29 Twitter for iPhone	negative
4. I'm missing my sister's commencement today :(#wannagohome #6moredays monicashokar 2012-05-25 14:05:33 Twitter for iPhone	negative
5. Dont need to bother anymore lah.... :(CherylRainieLuv 2012-05-25 14:05:52 Twitter for Android	negative
6. Oh come on @Manilaconcerts, bring @onedirection to Manila! Directioners want to see them! :(abegailramiro 2012-05-25	negative

14:05:58 Mobile Web	
7. my foot still hurts :(TweeTaToasty 2012-05-25 14:06:08 Twitter for Android	negative
8. @nielskooiker ohja :(iknow :(:(: :(: :(:(NoryKooiker 2012-05-25 14:06:12 web	negative
9. @siobhannnx17 noo!! Okay :(erinfitzgibbons 2012-05-25 14:06:14 Twitter for iPhone	negative
10. Quiero playaa :(LisJuez 2012-05-25 14:06:15 UberSocial for BlackBerry	negative
11. I need some sunshine in my life. *sigh :(ambaam_ 2012-05-25 14:06:58 UberSocial for BlackBerry	negative
12. RT @RikkuZZZ: You : Mom Can I Go ?!:(Mom : Ask Your Dad ... ~o) You : Dad Can I Go ?! :| Dad : Ask Your Mom ... ~o) NONOB220 2012-05-25 14:07:03 web	negative
13. Can't sleep cs I'm so hungry :(annmgrr 2012-05-25 14:07:05 Twitter for iPhone	negative
14. @denese_ong nvm then :(kentonnx 2012-05-25 14:07:16 Twitter for iPhone	negative
15. getting sick :(#noo jmdoyle94 2012-05-25 14:07:18 web	negative
16. @_dnsx ik oook :(POIS0NED 2012-05-25 14:07:25 web	negative
17. @AelinorGreyjoy I hope your dog catches fire. >:(Ser_Darkstar 2012-05-25 14:07:42 web	negative
18. @irun_uu_hoes @n_denae32 ohkay...so ii takee tht gigi aint gunna miss mehh: :(Yana_Shanta 2012-05-25 14:07:43 Mobile Web	negative
19. Whole family eating ice cream... But except me :(_sweet_corn_ 2012-05-25 14:08:09 Twitter for iPhone	negative
20. Can i still call you my baby? :(Liaoww 2012-05-25 14:08:36 TweetCaster for Android	negative
21. RT @GG_sargeant: Scar "If they were a cheese cake they would be the biscuit the lowest of the low" Me"but I like the biscuit :(" LOL <3 jess_pope1 2012-05-25 14:08:38 Twitter for iPhone	**positive**
22. My beer is going to be warm :(CalumJamesSmith 2012-05-25 14:08:45 Twitter for iPhone	negative
23. That moment when all you do is cry because the seniors are leaving :(#imgonnamissyouguys #<3youthatnesbittboy <3 MollyCondon9 2012-05-25 14:08:45 txt	negative
24. RT @_Anhum_: OK winter come back. NOW. Please. :(ZainMrk 2012-05-25 14:08:51 Twitter for iPhone	negative
25. What sucks is that I like this girl. :(DannysEyepatch 2012-05-25 14:09:01 twicca	negative

Table B - 25 tweets that were selected for the :(emoticon: this was retrieved from our dataset, for more information email the author.

Appendix C - Explained Terms
The reason why we have explained terms is because some proofreaders had difficulty with some the following terms [18].

Emoticons: examples are :), :(, :-), xD and :-P

Tweets: Users on Twitter are allowed to send message. These messages are allowed to be 140 characters long or shorter. The message will be posted on their personal profile page and of the Twitter stream of everyone who follows them.

Mention or @mention: a referred user in a tweet. An example of an @mentioned in a tweet is, for example, hahahaha I luv you xoxo @JustinBieber.

Following a Twitter user: when you follow a Twitter user, you receive his tweets on your Twitter stream.

Twitter Stream: A place where all the tweets of the users who you are following appear.

Positive emoticons: emoticons which are positive such as :-), :), :-D, :P, ;-), ;-d,(^_^), =}

Negative emoticons: emoticons which display negative emotions such as :-(, :-{, :@

Positive psychology: To avoid any confusion, positive psychology is a part of psychology. It "studies what people do right and how they manage to do it." Positive psychology is intended as a complement on the traditional research areas of psychology. Martin Seligman was the founding father of it (Compton 2005, 3).

Subjective well-being: this is a construct by psychologists which measures happiness.

P/N ratio: a term used by Fredrickson and Losada to indicate the ratio between positive and negative affect (Fredrickson and Losada 2005, 678).

Affect: affect is a feeling or emotion. It is distinguished from a cognition, thought or action[19].

Flourishing: "to flourish means to live within an optimal range of human functioning, one that connotes goodness, generativity, growth, and resilience. " (Fredrickson and Losada 2005, 678)

Languishing: the opposite of flourishing, see Fredrickson and Losada (2005) for a more precise definition.

---

[18] The proofreaders of this paper were in their second or third year of their bachelor program in: Lifestyle Informatics, Business Informatics, International Business Administration, Anthropology and Psychology. They did not read the whole paper, most of them read the abstract only.

[19] http://www.edpsycinteractive.org/topics/affect/affsys.html

Positive Sentiment: a positive sentiment is any emotion or feeling of the positive emotions or feelings. It is positive affect[20]. This is derived from Mehta et. al (2012, 73), they defined the term sentiment analysis.

Negative Sentiment: a negative sentiment is any emotion or feeling of the negative emotions or feelings. It is negative affect. This is derived from Mehta et. al (2012, 73), they defined the term sentiment analysis.

Sentiment (the term used in psychology): this term will not be used in this paper, but it is important to make the distinction that the word sentiment is *not* the same word as is used in psychology. Sentiment is related to long-term feelings such as love, companionship and trust (Fischer 2010, 14).

Infochimps: infochimps is a website that has a lot of datasets on, for example, Twitter.

[20] http://www.saaip.org/program.html this page is about sentiment analysis from an AI and psychological perspective and the text on the page claims that there is no formal definition of “sentiment” and “affect”. Therefore, we defined it to what we believe it is.

Appendix D - Repeat our experiments yourself

on the blog of Melvin Roest (http://appinez.com/?p=86) it is possible to download our Python, Java and R source files in order to recreate the experiment.

Unfortunately, we are not allowed to put the data set online, so you will have to capture your own. Note that you do need your own access key, app secret, access token and one other thing for the python source code.

It is easily findable in the Twitter documentation what you specifically need to do and in the source code only four string constants need to be changed.