

Detection and Visualization of Emotions in an Affect-Aware City

Benjamin Guthier¹, Rajwa Alharthi¹, Rana Abaalkhail¹, Abdulmotaleb El Saddik^{1,2}

¹Multimedia Communications Research Laboratory, University of Ottawa, Canada

²Division of Engineering, New York University Abu Dhabi, UAE

{bguthier, ralha081, rabaa006, elsaddik}@uottawa.ca

ABSTRACT

Smart cities use various deployed sensors and aggregate their data to create a big picture of the live state of the city. This live state can be enhanced by incorporating the affective states of the citizens. In this work, we automatically detect the emotions of the city's inhabitants from geo-tagged posts on the social network Twitter. Emotions are represented as four-dimensional vectors of pleasantness, arousal, dominance and unpredictability. In a training phase, emotion-word hashtags in the messages are used as the ground truth emotion contained in a message. A neural network is trained by using the presence of words, hashtags and emoticons in the messages as features. During the live phase, these features are extracted from new geo-tagged Twitter messages and given as input to the neural network. This allows the estimation of a four-dimensional emotion vector for a new message. The detected emotions are aggregated over space and time and visualized on a map of the city.

Categories and Subject Descriptors

I.2.7 [Artificial Intelligence]: Natural Language Processing—*Text analysis*

General Terms

Algorithms

Keywords

Affect-Aware City; Emotion Detection; Twitter; Emotion-word Hashtag; Dimensional Emotional Model

1. INTRODUCTION

A smart city's main objective is to increase the quality of life for its citizens and to make the city more attractive, lively and greener [12]. To meet this goal, physical sensors are deployed throughout the city to monitor various aspects such as environmental parameters (weather, pollution, etc.),

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM'14, November 3–7, 2014, Orlando, Florida, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3126-5/14/11 ...\$15.00.

<http://dx.doi.org/10.1145/2661704.2661708>.

traffic and the consumption of resources. These sensors are interconnected into a computing framework to create a big picture of the live state of the city. This live state, however, only includes measurable quantities and disregards how the citizens actually feel. It is likely that there exist correlations between the affective states of the citizens and relevant statistics like well-being of the city's inhabitants or quality of living. When urban planners use the collected data to optimize parts of the city, the affective state of the inhabitants can thus serve as valuable implicit feedback. This gives rise to our vision of an "Affect-Aware City" with the ability to understand and utilize the affective states of its citizens to enable improved planning and decision making [10]. Relevant affective states in this scenario include emotion, feelings, personality, mood and subjective well-being. In this paper, we focus specifically on emotion.

Affective states can be detected from various sources that are available in a smart city. Examples are hard sensors such as cameras and environment sensors, or soft sensors such as information extracted from social networks or smartphones. Online social networks are used frequently by millions of people to share their current thoughts, observations and opinions. Social networks have been shown to provide more reliable and trustworthy information than traditional methods like questionnaires and other sources [16]. Moreover, social networks provide a huge amount of publicly accessible data without requiring any specialized equipment or particular participation of the citizens. Twitter is a popular example of such a social network where individuals send short messages called *tweets* about what they are feeling and thinking. The result is a large corpus of public real-time data that is rich with emotional content. Twitter has been used successfully in studies about disasters and emergencies [13], public health monitoring [2], events exploration [15], and more recently, in sentiment and emotion analysis. It is also possible to add meta data including location and time when posting a message. The location can be inferred from internet access points that are being used, directly given as geographical coordinates (latitude and longitude) if the tweet is sent from a mobile phone, or provided as additional user information. For these reasons, we decided to use Twitter as a source of geo-tagged affective data in our work.

Hashtags are keywords prefixed with the # symbol used by Twitter users as a way to categorize their tweets and to search for tweets that have relevant topics. They are also frequently used as a form of user-added meta data to express other aspects, including emotion, that cannot be expressed by the tweet directly. We use emotion-word hashtags as

search criterion to obtain emotion-annotated textual data from Twitter. This data is used to train a neural network to detect emotion from a number of textual features. In the online phase of our system, geo-tagged tweets are obtained from Twitter. The neural network then automatically detects the emotion contained in the tweets. The result is geo-tagged emotion data that is aggregated and visualized on a map of the city. The neural network is necessary, because the percentage of tweets containing both location data and emotion-word hashtags at the same time is too low to provide a sufficient amount of data.

This paper is structured as follows. Section 2 discusses the related work in the field of affect detection from social networks. Our system to detect and visualize emotion is presented in Section 3. Section 4 contains the experimental results of the training phase of our neural network. Conclusions of the paper and an outlook to future work are given in Section 5.

2. RELATED WORK

The detection of affect from social networks, especially from Twitter, has recently received widespread interest. The types of affective states being detected range from pure sentiment, which is either positive or negative, to mood and emotion. The latter is the most closely related to our work.

A method for *sentiment analysis* has been proposed by Go et al. [8]. They automatically classify the sentiment of any Twitter message into either positive or negative by using distant supervision learning. Twitter messages containing emoticons are used as labeled dataset to train different classifiers. The classifiers are: keyword-based, Naive Bayes, maximum entropy, and support vector machines. The best classifier reached an accuracy of 83%.

In another approach, the sentiments of tweets from 573 real Twitter profiles from London are extracted and analyzed using word count and maximum entropy classifiers. After comparing the result to the Index of Multiple Deprivation obtained from the UK National Statistics, it was proven that the sentiment of the community’s tweets can reveal its socio-economic well-being [20].

To the best of our knowledge, all works in the field of *emotion detection* from social networks use Ekman’s discrete model to represent emotion [5]. The model uses six basic emotion labels that are universal among different cultures. The basic emotions are anger, disgust, fear, happiness, sadness and surprise. More labels are added in other works as needed. There is an ongoing controversy in the field of psychology about how adequately discrete labels categorize emotion [11]. As an alternative, dimensional models of emotion have been proposed, where an emotion is expressed by a number for each dimension. In its most simple form, the two dimensions *pleasantness* (P) and *arousal* (A) are being used. Pleasantness corresponds to the common intuition of good and bad emotions. High pleasantness is indicative of a situation that is in agreement with a person’s own goals. Arousal denotes the level of physiological change the emotion causes in a person. High levels of arousal are characterized by a faster heartbeat and the readiness for action which is also found in emotions like stress and anger. This is in contrast to low arousal states like contentment [22]. In our work, we use the four-dimensional model proposed by Fontaine et al. which adds the two dimensions *dominance* (D) and *unpredictability* (U). The former is characterized by the feeling of

control, the impulse to act and changes in speech. Examples for emotions with high dominance are pride and anger, which are in contrast to low dominance emotions like sadness and shame. Unpredictability describes the novelty of the emotion eliciting stimulus. This dimension opposes behaviors like jaw dropping and raised eyebrows to the feeling of expectedness and familiarity [7]. The paper by Fontaine et al. lists 144 statements and their relationship to the four dimensions which is helpful for their interpretation.

In this paper, we refer to an emotion represented in the four-dimensional emotion model as 4D emotion vector or PADU values. We believe that using a more fine-grained dimensional model to represent emotion is less “black and white” than using coarse emotion labels as it allows blends of different emotional states. Another benefit of a dimensional model is its natural inclusion of the neutral emotional state (zeros in all dimensions) and any value between neutral and full emotion. The challenges brought forth by the fuzzy boundary between neutral and emotional text messages have been previously identified in [1]. Note that assigning a discrete emotion to a textual message is a problem of classification, whereas assigning a 4D vector requires regression. This means that our scenario belongs to a different family of machine learning problems than the related work, and existing techniques cannot be applied easily.

In [6], a learning approach has been proposed that combines textual sentiment and affect analysis techniques in a hierarchical classification scheme. The hierarchy consists of three levels: neutrality, polarity and emotions. In the first step, data is classified based on neutrality versus emotionality. Next, the polarity of the emotional messages is inferred (sentiment analysis). The messages with negative polarity are further divided into five basic emotions.

Nagarsekar et al. [19] classified tweets into the six basic emotions using supervised learning with two different machine learning algorithms. Both classifiers were trained and tested on three datasets that varied in their distribution of emotions. The maximum accuracy of 83% was achieved for a Naive Bayes classifier trained on a dataset that had an equal distribution of emotions.

The authors of [21] collected a corpus of emotional tweets using 14 popular topics with their corresponding list of hashtags. With the aid of a developed annotation tool, they manually labeled 7000 tweets with a finer-grained set of Ekman’s six basic emotions plus love.

To avoid the high effort of manual annotation, [17] and [18] created a large corpus of emotional tweets by specifically searching for Twitter messages containing one of the six emotion words as hashtags. This corpus is used to train and cross-validate support vector machines that classify the six basic emotions. Their experiments demonstrate that self-labeled hashtag annotations created from thousands of users are consistent and can be effectively used for automatic emotion detection. Also using the approach of harnessing emotion-word hashtags, the authors of [25] automatically created a large and high quality emotion-labeled dataset covering Ekman’s six basic emotions plus thankfulness. After collecting 5 million tweets from 131 twitter emotional hashtags, they applied a set of developed heuristics to filter out all but the emotion-relevant tweets (about 2.5 million). The two machine learning algorithms LIBLINEAR and Multinomial Naive Bayes were then used to identify emotion and to study the effectiveness of various features

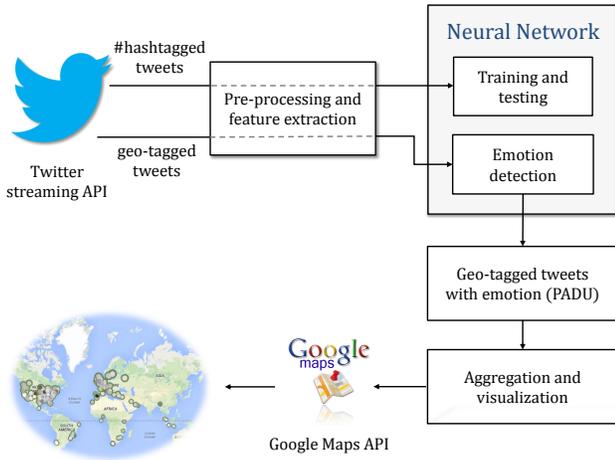


Figure 1: Overview of the proposed system to detect and visualize emotions from geo-tagged Twitter messages.

including n-grams, n-gram positions, emotion lexicons, and part-of-speech. By combining these features, the classifier performance was improved and their approach achieved an accuracy of 65.57%. Due to their successful use in creating training data for machine learning algorithms, we decided to also use emotion-word hashtags to train the neural network in our work.

The Twitter hashtags approach has also been used to *detect mood*. For instance in [4], a systematic method to identify different moods from social networks is presented. 203 mood hashtags were used to collect messages. They were analyzed for affective experience in the two dimensions valence and activation and subsequently used to train automatic classifiers. The classifiers reached an accuracy of 83%. The authors of [3] studied how public mood is affected by popular events of social, economic and/or cultural importance that are occurring at the same time. They performed a sentiment analysis of public tweets using an extended version of the Profile of Mood States (POMS) to extract six dimensions of mood (tension, depression, anger, vigor, fatigue and confusion). Instead of using machine-learning algorithms, each tweet is scored based on key terms from the set of POMS adjectives. Similar to our idea of characterizing places based on emotion, [14] introduced a mobile interface that harvests real-time and geo-located information generated by the users of multiple social networks. Natural language analysis techniques are used to automatically infer the emotion in the text. A compass guides the users towards areas of the selected emotion.

3. DETECTION AND VISUALIZATION OF EMOTION

In this section, we present our system to automatically detect and visualize emotions from Twitter messages. See Figure 1 for an overview of the system. It works in two different phases: the offline training phase and the online phase. In the offline training phase, we collect messages from Twitter that contain emotion-word hashtags. The messages are pre-processed and features are extracted from them. The features, together with the emotion derived from the emotion-word hashtags are used to train a neural network. In the

- I was actually falling asleep, but I stayed up #proud
- There’s something beautiful about happy people :) #smile #joy
- Aww that is so sad :(#disgusting #sad
- So this just showed up via UPS #surprise
- I hope you have a fantastic time! #jealous
- Safe travels. Have fun!! #jealous
- Argentina won on penalties #angry

Figure 2: Examples of Twitter messages that contain emotion word hashtags. The last three examples illustrate some of the difficulties of emotion detection. The negative emotion jealousy is used in a positive context as a compliment. Without the hashtag #angry, the pleasantness dimension of the last message is ambiguous.

online phase, the system collects live geo-tagged tweets from the area of interest, e.g., a city. The trained neural network is used to detect emotion in these new tweets. The geo-tagged emotion data is then aggregated and visualized on a map. Details of each of the processing steps are given in the following.

3.1 Obtaining Training Data

In order to produce the training data for our neural network, we use the approach presented in [17]. Users of Twitter sometimes add hashtags containing emotion-words to their messages to express their current feelings. Figure 2 shows some examples of such Twitter messages. It has been shown that these emotion-word hashtags are consistent and reliable. They correspond well with the labels a trained external annotator would assign to the message. Emotion-word hashtags are thus well-suited for the use as training data for a machine learning algorithm. For the reasons mentioned in Section 2, we choose the four-dimensional model of emotion to represent the emotion in a tweet, rather than the commonly used discrete model. The paper by Fontaine et al. contains a list of 24 emotion-words with their corresponding four-dimensional emotion vector [7]. An excerpt of this list is given in Table 1. We use the list as search terms to obtain tweets containing the desired emotion-word hashtags. The emotion in a tweet is then the four-dimensional emotion vector corresponding to the emotion-word, or their average if multiple emotion-words are present in a message. In addition, we use 18 variations of the terms in the list that are more commonly found in Twitter messages than the original ones. For example, Twitter users are more likely to use the adjective *angry* as a hashtag than the original noun *anger* contained in the list. Five extra emotion-words were added that are commonly found in tweets, but are not contained in the original list (e.g., nervous, smile). Their 4D emotion vector was derived from the table of features given in [7].

3.2 Pre-Processing

We use the filter endpoint of the public stream Twitter API to download the messages from Twitter. The API provides a stream of English tweets that is filtered for the 47 emotion-words. At this point, we do not require the mes-

Emotion	P	A	D	U
#anger	-1.1	1.4	1.3	0.1
#contempt	-1.5	-1.0	1.2	-0.3
#contentment	1.2	-1.2	0.7	-0.7
#disgust	-1.1	-0.9	0.2	1.6
#jealousy	-1.0	0.5	0.6	-0.6
#pleasure	1.5	0.1	0.5	-0.2
#pride	0.9	-0.2	1.0	-0.8
#sadness	-0.3	-1.0	-1.5	-0.2
#shame	-0.2	0.3	-1.6	-0.8
#surprise	0.7	0.6	-0.1	3.2

Table 1: Subset of the emotion-word hashtags that are used as search terms to obtain training data for our neural network. A message containing one of these hashtags is assigned the four-dimensional emotion vector as given in the table. The emotion words and the PADU values are taken from [7].

sages to be geo-tagged. In addition to the message body, the tweets come with a time stamp, a unique ID, the ID of the sender and may contain location information. Hashtags, external links and usernames in the message body are marked separately, which is convenient for later processing.

Before they can be used as training data, the tweets need to be pre-processed. This procedure is identical for the training data and in the online phase. The first step of pre-processing is text normalization. Tweets are short text messages that are no longer than 140 characters and are often written on mobile devices. As a consequence, they typically contain a large number of abbreviations, colloquial expressions and non-standard words. The goal of text normalization is to remove the variety from the messages, which constitutes noise in the training data.

The English language does not contain any words where the same letter is repeated more than twice in a sequence. However, longer runs of the same letter are often used in tweets for emphasis (for example: “I am boooored!”). We replace such runs with a single occurrence of the repeated letter. In our experiments this heuristic more often led to the correct spelling of the word than replacing it with a double letter. Next, all contractions are expanded, like “aren’t” to “are not” and “they’ve” to “they have”. We also expand a list of common chat acronyms and abbreviations into their original long form (e.g., “rn” to “right now”). The same is done for colloquial expressions like “gonna” (“going to”). As a last normalization step, all messages are transformed into lowercase. In this fashion, a large portion of the variety in the text is removed.

A similar problem is the large variety of emoticons contained in the messages. An emoticon is a short sequence of symbols like :) or a single Unicode character. They resemble small pictures and are another widely used tool to express emotion in a text message. They are thus very relevant to our task. Due to the flexibility of Unicode, there exists a large variety of emoticons, with many of them expressing very similar concepts. In previous experiments, we identified 130 of the most commonly used emoticons and classified them into 30 general categories. Examples for these categories are: hearts, smiling, frowning, etc. The emoticons in the tweets are replaced with their category to unify them.

Twitter gives users the ability to forward tweets that were sent by someone else. Such forwarded messages are called *retweets*. Since a retweet repeats a message from someone other than the poster of the retweet, we ignore them for

emotion detection. They are easily identified by the “RT” prefix in the message. Similarly, we ignore messages with less than three English words and those containing overly inappropriate content.

The last pre-processing step is the tokenization of the message to simplify feature extraction. Our system identifies the following tokens: hashtag, word, emoticon, link, username and symbol. Hashtag and word tokens have the corresponding text string attached to them. The value of an emoticon token is one of the 30 categories. Links are URLs to various media. Username tokens are references to other users. Symbol tokens are punctuation and special characters. Links, usernames and symbols are ignored in the later processing.

3.3 Feature Extraction

From the pre-processed and tokenized messages, we extract 110 features. The features are the same for the training and the online phase. Each feature is a numerical value that is set to either 1 if an emoticon, word or hashtag is present in the message or 0 otherwise. The first 30 features denote the presence of emoticons from the 30 categories. The next 40 features are the presence of common words, and the remaining 40 are common hashtags. In the training and testing phase, the 47 emotion-word hashtags are excluded as features to not bias the results. Hashtags that are surrounded by words are counted as words, because they are not suitable as emotion labels. This is in line with the findings of [9] and [17].

The vector of 110 features for each tweet and its four-dimensional emotion vector derived from the emotion-word hashtags are then used as the training data for our neural network. The network is trained with the Levenberg-Marquardt algorithm. More details about the training of the network can be found in Section 4.

3.4 Online Phase

In the online phase, the system collects live data from Twitter. The process of obtaining the tweets, pre-processing and feature extraction is similar to the training phase. In the online phase, only geo-tagged tweets are collected, regardless of whether they contain emotion-word hashtags or not. Requiring both location and emotion hashtags at the same time would be too limiting to obtain enough data. Instead, the trained neural network is used to estimate a 4D emotion vector from the features extracted from each message. Like this, latitude, longitude and a four-dimensional emotion vector is available for each tweet which can then be visualized.

3.5 Aggregation and Visualization

In the live phase of our system, it obtains real-time geo-tagged tweets. They are annotated with a 4D PADU vector by the neural network. Depending on the size and the population of the chosen region of interest, the system will acquire in the order of one tweet per second for a large city or about 50 tweets per second for the entire world. Visualizing this steady stream of data on a map requires aggregation of tweets over space and time. Aggregation over space is accomplished by dividing the area to be visualized into rectangular cells. Based on the latitude and longitude of a tweet, it belongs to exactly one map cell. Each time a new tweet is sent from within a cell, its PADU values are averaged with all other tweets in the cell. In a study about the duration of

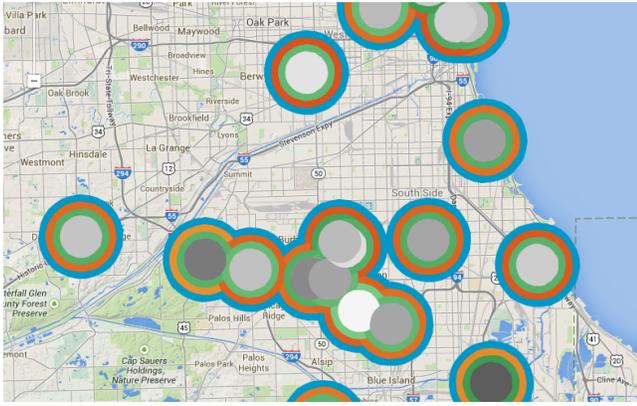


Figure 3: Visualization of geo-tagged emotional Twitter messages originating from Chicago¹.



Figure 4: Aggregated emotion detected from tweets from the USA¹.

emotional experience, 80% of the emotional episodes were found to last for less than 30 minutes [24]. We use this finding as the basis for our aggregation over time. The emotional tweets in a cell remain valid for 30 minutes. After this time period, they are discarded and no longer considered for averaging.

We decided to visualize the aggregated 4D emotion vector of each cell as four concentric disks. Their center point coincides with the center of the cell. The radii of the four disks differ and they are stacked like the towers of Hanoi (see Figure 3 for an example). The overall radius of the stack of disks is chosen in proportion to the amount of active tweets in the cell, so a bigger circle represents a higher number of tweets that were being sent. The order of the disks was chosen based on their importance as outlined in [7]. *Pleasantness* is considered to be the most meaningful of the four dimensions. It is thus visualized as the center disk, which is the most visible one. *Unpredictability* is considered to have the least significance and is shown in the outermost ring.

We created a color model to determine the color of each of the four disks representing the PADU dimensions. For each dimension disk, we chose a base color. Pleasantness is associated with the color grey, arousal is shown in green, dominance is red, and unpredictability is displayed in blue. The brightness of the color expresses the value of this dimension, with brighter colors representing higher values. All of the colors representing values in between the extremes are calculated by interpolation. Refer to Section 2 for the inter-



Figure 5: Overview of all 3 million emotional tweets used in this visualization example¹.

pretation of the four dimensions.

With the colors of the four disks calculated, the stacks corresponding to map cells that contain at least one tweet are then drawn on a map at the size corresponding to the number of tweets. We use the Google Maps API for the final drawing step. Figures 3, 4 and 5 show three example views. The data that was used for the images consists of 3 million tweets that were processed by the system over the course of 20 hours. The areas for visualization were chosen to be the city of Chicago in Figure 3, the USA in Figure 4 and the entire world in Figure 5. Depending on the zoom level, emotions can be explored at resolutions comparable to the size of entire cities, city districts or individual blocks of houses.

4. EXPERIMENTAL RESULTS

50,000 Twitter messages with emotion-word hashtags were obtained over the course of 36 hours to train the neural network. As described above, 110 features were extracted and a PADU emotion vector was assigned to each message. 8000 of the messages were reserved for validation. The rest was used for supervised training of a neural network with one hidden layer. The network has 110 input nodes, 10 hidden nodes and 4 output nodes.

After training, the network was used on the 8000 unused samples. The average mean square error (MSE) between the PADU values of the messages and those calculated by the network was 0.55. The achieved correlation was 0.51, which is very close to the inter-annotator agreement found in other works [23]. In one of our experiments predating the training of the neural network, three annotators manually assigned PADU values to 250 tweets. Only 42% average agreement was achieved in this experiment. The highest level agreement was reached for the pleasantness dimension (70%) and the lowest for arousal (only 17%). The low levels of agreement found throughout the literature are due to the inherent difficulty of assigning emotion values to text messages.

To determine the optimal number of hidden nodes for the network, we varied the number between 5 and 50. No significant improvement of the mean square error and the correlation was found. The maximum MSE difference between node numbers above 10 was 0.05. However, the training time increased by a factor of 10 at the same time.

¹Map Data ©2014 Google

In the last experiment, we analyzed the impact of the number of messages used for training on the error of the trained network. We increased the number of training messages from 10% to 100% of the available data. The MSE decreased from 0.63 to 0.55 when going from using 5000 messages for training to using all 50,000.

Note that due to the fundamental difference between the dimensional and the discrete model of emotion, it is difficult to compare our results to the related work. From the features of a text message, we obtain a 4D vector by regression. The resulting MSE of 0.55 denotes the average distance to the real coordinates of the message from the training data. Existing approaches classify a message into a discrete category. This classification can be either correct or wrong, and there exists no metric to describe the distance to the correct category. Likewise, clustering the 4D space into discrete subspaces would undo the benefits of using a continuous dimensional model to represent emotion.

5. CONCLUSIONS AND FUTURE WORK

We presented a system to detect emotions and visualizing them on a map. Twitter is used as data source. Emotions are detected from geo-tagged tweets by using a neural network. Our system is intended to be a tool for attaining a more complete view of the live state of a city that goes beyond what the physical sensors of a smart city can provide. It can assist decision making processes and urban planning by providing a form of implicit feedback from the citizens.

Our work is limited in that we require tweets to be geo-tagged by the sender. This limits the amount of usable messages, because adding location is an optional feature of Twitter. In future work, we would like to infer location by other means to overcome this limitation. Furthermore, bad spelling in tweets constitutes a problem for consistent feature extraction which may be tackled by employing a dictionary of the English language. We would also like to explore the possibility of tracking a user's emotion over time to increase the accuracy of the system.

6. REFERENCES

- [1] C. O. Alm. *Affect in text and speech*. VDM Verlag, 2009.
- [2] E. Aramaki, S. Maskawa, and M. Morita. Twitter catches the flu: detecting influenza epidemics using twitter. In *Proc. of the Conf. on Empirical Methods in Natural Language Processing*, pages 1568–1576, 2011.
- [3] J. Bollen, H. Mao, and A. Pepe. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *ICWSM*, 2011.
- [4] M. De Choudhury, S. Counts, and M. Gamon. Not all moods are created equal! Exploring human emotional states in social media. In *Sixth International AAAI Conference on Weblogs and Social Media*, 2012.
- [5] P. Ekman. Are there basic emotions? In *The nature of emotions: fundamental questions*. Oxford University Press, 1992.
- [6] A. A. A. Esmin, R. L. de Oliveira, and S. Matwin. Hierarchical classification approach to emotion recognition in twitter. In *Machine Learning and Applications*, volume 2, pages 381–385, 2012.
- [7] J. R. Fontaine, K. R. Scherer, E. B. Roesch, and P. C. Ellsworth. The world of emotions is not two dimensional. *Psychological science*, 18(12):1050–1057, 2007.
- [8] A. Go, R. Bhayani, and L. Huang. Twitter sentiment classification using distant supervision. *CS224N Project Report, Stanford*, pages 1–12, 2009.
- [9] R. González-Ibáñez, S. Muresan, and N. Wacholder. Identifying sarcasm in twitter: a closer look. In *Proc. of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, volume 2, pages 581–586, 2011.
- [10] B. Guthrie, R. Abaalkhail, R. Alharthi, and A. El Saddik. The affect-aware city. In *Int. Conf. on Computing, Networking and Communication*, 2015.
- [11] S. Hamann. Mapping discrete and dimensional emotions onto the brain: controversies and consensus. *TICS*, 16(9):458–466, 2012.
- [12] C. Harrison et al. Foundations for smarter cities. *IBM J. of Research and Development*, 54(4):1–16, 2010.
- [13] T. Hossmann et al. Twitter in disaster mode: Opportunistic communication and distribution of sensor data in emergencies. In *Proc. of the 3rd Extreme Conference on Communication*, 2011.
- [14] S. Iaconesi and O. Persico. An emotional compass. In *Emotion and Sentiment in Social and Expressive Media*.
- [15] A. Marcus et al. Twitinfo: aggregating and visualizing microblogs for event exploration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 227–236, 2011.
- [16] A. E. Marwick et al. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13(1):114–133, 2011.
- [17] S. Mohammad. #emotional tweets. In *Proc. of the 6th Int. Workshop on Semantic Evaluation*, pages 246–255. ACL, 2012.
- [18] S. Mohammad and S. Kiritchenko. Using hashtags to capture fine emotion categories from tweets. *Computational Intelligence*, 2014.
- [19] U. Nagarsekar et al. Emotion detection from "the SMS of the internet". In *Recent Advances in Intelligent Computational Systems*, pages 316–321, 2013.
- [20] D. Quercia, J. Ellis, L. Capra, and J. Crowcroft. Tracking gross community happiness from tweets. In *Proc. of the ACM Conf. on Computer Supported Cooperative Work*, pages 965–968, 2012.
- [21] K. Roberts, M. A. Roach, J. Johnson, J. Guthrie, and S. M. Harabagiu. Empatweet: Annotating and detecting emotions on twitter. In *LREC*, pages 3806–3813, 2012.
- [22] J. Russell. A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6):1161–1178, 1980.
- [23] C. Strapparava and R. Mihalcea. Semeval-2007 task 14: Affective text. In *Proc. of the 4th Int. Workshop on Semantic Evaluations*, pages 70–74. ACL, 2007.
- [24] P. Verduyn et al. The relation between event processing and the duration of emotional experience. *Emotion*, 11(1):20–28, 2011.
- [25] W. Wang, L. Chen, K. Thirunarayan, and A. P. Sheth. Harnessing twitter "big data" for automatic emotion identification. In *Int. Conf. on Social Computing*, pages 587–592, 2012.