

Spring 6-10-2017

# DETECTING PANIC POTENTIAL IN SOCIAL MEDIA TWEETS

Anuja Hariharan

*Karlsruhe Institute of Technology, Germany, anuja.hariharan@kit.edu*

Verena Dorner

*Karlsruhe Institute of Technology, verena.dorner@kit.edu*

Christof Weinhardt

*Karlsruhe Institute of Technology, weinhardt@kit.edu*

Georg W. Alpers

*University of Mannheim, alpers@uni-mannheim.de*

Follow this and additional works at: [http://aisel.aisnet.org/ecis2017\\_rip](http://aisel.aisnet.org/ecis2017_rip)

---

## Recommended Citation

Hariharan, Anuja; Dorner, Verena; Weinhardt, Christof; and Alpers, Georg W., (2017). "DETECTING PANIC POTENTIAL IN SOCIAL MEDIA TWEETS". In Proceedings of the 25th European Conference on Information Systems (ECIS), Guimarães, Portugal, June 5-10, 2017 (pp. 3181-3190). ISBN 978-0-9915567-0-0 Research-in-Progress Papers.  
[http://aisel.aisnet.org/ecis2017\\_rip/65](http://aisel.aisnet.org/ecis2017_rip/65)

This material is brought to you by the ECIS 2017 Proceedings at AIS Electronic Library (AISeL). It has been accepted for inclusion in Research-in-Progress Papers by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact [elibrary@aisnet.org](mailto:elibrary@aisnet.org).

# DETECTING PANIC POTENTIAL IN SOCIAL MEDIA TWEETS

*Research in Progress*

Hariharan, Anuja, Karlsruhe Institute of Technology, [anuja.hariharan@kit.edu](mailto:anuja.hariharan@kit.edu)

Liao, Wei, [williamlv211@gmail.com](mailto:williamlv211@gmail.com)

Dorner, Verena, Karlsruhe Institute of Technology, [verena.dorner@kit.edu](mailto:verena.dorner@kit.edu)

Weinhardt, Christof, Karlsruhe Institute of Technology, [weinhardt@kit.edu](mailto:weinhardt@kit.edu)

Alpers, Georg, Universität Mannheim, [alpers@uni-mannheim.de](mailto:alpers@uni-mannheim.de)

## Abstract

*A high degree of real-time interconnectedness can aid information transmission, particularly in disaster situations. However, it can have substantial negative consequences when information is emotionally laden and transmits these emotions, particularly the emotion of panic, to the individual across social media in an already grave situation. Prior research has shown that information laden with emotion spreads through social network faster than otherwise. Hence, we highlight the need to understand and curtail potentially panic-causing information, without compromising on good quality information from being available for effective crisis communication and management. With this research, we present the necessity of detecting the panic potential of social media messages, and aim to address two research questions: What are the features, and metrics necessary, to compute and evaluate the panic potential of a social media message (respectively)? Our planned analysis takes the case of the Munich shooting incident, 2016, based on user tweets immediately after the incident. Different features and evaluation metrics are proposed and discussed. The work aims to detect panic potential of messages in social media networks during disasters.*

*Keywords: Panic Detection, Disaster Communication Management, Social Media, Individual state*

## 1 Introduction

The role of online social media in disaster information management has been recognized heavily in recent events of terror, as well as during natural disasters (Simon et al. 2014; Palen, 2008). Information transmitted during crisis events is seldom void of emotions, and the use of social media particularly accelerates the spread of emotional content, in addition to its purely informational purpose (Woo et al. 2015). For instance, Stieglitz & Dang-Xuan (2013) observed that emotionally charged messages are likely to be substantially forwarded and retweeted faster particularly in crisis events. The phenomena of emotion contagion (Coviello et al. 2014) necessitates the use of discretion, when designing effective information management systems for crisis events. Novel challenges are being brought to social media by social bots (Ferrara et al. 2016), which are engineered to give the false impression that a piece of information, regardless of its accuracy, is highly popular and endorsed by many, thus exerting an influence on the public, against which social media hasn't yet developed antibodies. While there exist disaster frameworks that are currently being applied in the management of disaster communication (Reynolds & Seeger 2005; Houston, 2012), none of these frameworks, by design, take into account the emotional state of the user, and their potential impact on information con-

sumption and subsequent sharing behaviour. Particularly, in disaster situations, thoughtless (or even intentional) spreading of panic-inducing messages can deepen a crisis by increasing disorder for active and passive participants, as well as social groups of potential victims.

In this paper, we focus on defining and detecting, what may be termed as the “panic potential” of a social media message – the likelihood of a (social media) message to induce a state of panic in a reader. The defined panic potential is designed to be a preventive one that can potentially be used as an evaluation metric to curtail information flow in a panic situation. The purpose of the proposed panic potential detector differs from existing event-detection systems (Imran et al. 2015), which determine ex-post, whether an event has occurred or not, based on content processing, or geographical influx of twitter messages, or identification of frequencies of specific keywords. We present a preliminary prototype of a panic potential detection component, based on tweets extracted from the Munich shooting incident, July 2016.

## **2 Literature Background**

### **2.1 Technology support for individuals in social media for disaster situations**

The emergence of social media in disaster recovery management has been indispensable, especially in the last decade. Different forms of social media such as Facebook, Twitter, Google Plus, and specialized online websites, are being employed to signal and detect disasters (Huang et al. 2010), broadcast requests for help (Acar & Muraki, 2011), and identifying support and alert systems to respond to these requests (such as EARS, Avvenuti et al. 2014). The emergence of these online tools for disaster management has been inorganic, 1) that does not proactively consider the repercussions of the impact of social media as a communication channel in disasters, and 2) particularly with little consideration for the potential impact of social media on the general public’s emotional state, information consumption and sharing patterns. Hence, there arises a need to detect the emotional content that is being transmitted through social media, when applying a structured approach as suggested by disaster communication frameworks such as CERC (Crisis & Emergency Risk Communication, (Reynolds & Seeger 2005)) and of DCIF (Disaster Communication Intervention Framework), Houston 2012; Houston et al. 2014, for the reduction of potential turmoil, and to re-establish empathy, reassurance, as well as reduction of crisis-related uncertainty. Moi et al. (2015) use information techniques to process and analyze social media data, and transform the high volume of noisy data into a low volume of rich content that is useful to emergency personnel. While this approach focusses on assessing information quality by metrics such as timeliness, understand ability and believability, the incorporation of an emotional state metric in addition to information quality to channelize data streams would be necessary in the current scenario of information flow on the basis of emotional states of social media users. We next review existing methods to detect emotional content of a social media context, and how these can inform to detect the “panic potential” of a given message.

### **2.2 Detecting user perceptions of sentiment and emotions with text analysis**

Emotions as expressed by the poster of a message might be perceived and interpreted very differently by the receiver of this message (Barrett et al. 2007). This difference potentially arises due to individual differences in emotion vocabulary, structure and content of the conceptual system for emotion and emotion perception (Barrett et al. 2007). In the social media context, emotions can be conveyed through words, messages, and graphical messages, leaving scope for a gap arising between the emotion conceptualized, and the emotion perceived by the reader (Barrett et al. 2007, Niedenthal et al. 1997). Understanding this distinction requires an investigation of the neural mechanisms underlying message interpretation, which is beyond the scope of the current work. However, we begin by detect-

ing emotional content of messages using based on knowledge of previous messages, and independent knowledge of emotional words, which is used in machine learning techniques.

In terms of text data analysis, with the increasing interest in machine learning, sentiment classification has become a popular area in text mining. Sentiments are detected based on detecting the positive and negative words in a given social media message, or other possibly relevant information and features of the tweet (such as time of the day, length of the messages, etc.). Nagy and Stamberger (2012) used the SentiWordNet library to detect sentiments during a hurricane incident, and in addition, used a comprehensive list of emoticons, a sentiment based dictionary and a list of out-of-vocabulary words such as lol, wow, etc., to achieve nearly 95% accuracy of detection. Hence, the choice of dictionary and the choice of features was shown to be an important element to improve the accuracy of the detection system. Verma et al. (2011) studied four mass emergencies, to classify sentiments of tweets using Naïve Bayes and entropy maximization methods. They identify that in addition to objective information, using features of subjective information and personal/impersonal styles of communication, substantially improved classification performance, which makes these classifiers also usable across events of similar kinds. Halse (2016) investigate Twitter data in man-made and natural disasters, to reveal that the use of humor and type of emotion used, among others, influence the trustworthiness and usefulness of messages. Bruns & Stieglitz (2013) identify three kinds of metrics which are critical to analysing sentiments in tweets: user, temporal, and combined tweets, which are also applicable for disaster events. We refer to these metrics henceforth as features of tweets, which aid to successful classification of tweets based on their panic potential.

In addition to sentiment, classification of emotions is also a possible approach, to detect the emotional content of a message. Emotion classification gives fine grained information about the nature of a positive or a negative sentiment, and can be used directly to deduce the state of panic. For instance, emotions of fear and anxiety are likely to be associated with a state of panic, emotions of surprise and sadness possibly still relevant, and emotions of joy are possibly not related. Emotion model classification for panic is hence not a binary, but rather a multi-class classification problem. Since emotion classification in text requires more training inputs as well as computational effort, many researchers take advantage of acoustic (Yang et al. 2006) or visual features (Machajdik & Hanbury, 2010) to classify facial expressions. Kim et al. (2010) compared several unsupervised methods of classification, such as Latent Semantic Analysis (LSA) and focused, on identifying anger, fear, joy and sadness. Aman and Szpakowicz (2008) in contrast leveraged advantages of supervised methods i.e. SVM to classify sentences into one of Ekman (1992)'s six basic emotions, namely (i.e., surprise, fear, disgust, joy, sadness, anger). Ekman's emotions are skewed towards negative emotions, and contain only two potentially positive emotions: surprise and joy. However, the classification is a useful one to employ in studying panic disorders (Pine et al. 2005) in clinical applications, and hence we rely upon the Ekman model in our classification method.

### **2.3 Detecting panic potential of messages with text analysis**

Clark (1986)'s cognitive model of panic attacks, describes panic as a state of mind of people, when they are threatened by the presence or potential presence of fear-related phenomenal states. Clark's model can be thought of as an extension of the fear hypothesis, however panic-related sensations are not uniquely associated with fear. A considerable body of research has focused on the establishment of a measure of anxiety sensitivity – the Anxiety Sensitivity Index (ASI; Reiss et al. 1986), denoting the tendency to respond fearfully to symptoms of fear. These studies further show that panic patients, whether or not they have agoraphobia, seem to interpret ambiguous scenarios (potentially involving internal bodily sensations) in a threat-related and panic-related way and focus their attentional resources to interoceptive changes.

Panic and negative sentiment has been detected in the context of well-being (Honkela et al. 2012), health care event detection (Hripcsak et al. 2003), and in financial events to detect panic selling (Bozic & Seese, 2011). In the 2014 Ebola crisis, Lazard et al. (2015) demonstrated that text mining facilitates real-time investigations and uncover unique, potentially unpredicted, public opinions to inform rapid development of communication strategies and targeted messages to alleviate fears and confusion. Although there are several works examining the detection of events using twitter data, there is no work that examines detection of panic in a preventive manner – i.e., to curtail the (panic) information flow in the wake of disaster events, by computing the panic potential of a message. Hence, the first question we address is: *RQ1: What are the features necessary, to compute the panic potential of a social media message?*

Upon automatically detecting the panic potential of a message, subsequent questions arise, as to how to validate this panic potential, and which metrics would be suitable to perform this evaluation. Ground truth elicitation is a complex problem in this domain, given the scarcity of resources, as well as shortage of relevant data. Ground truth may be established by a combination of native speaker annotators. Evaluation of the detected panic potential may also be carried out by psychophysiological methods, which do not rely on perceived panic potential, but are based on the measured human responses to a given message, and their subsequent information sharing patterns. Hence, the second question addressed in this work is: *RQ2: What are the metrics necessary, to evaluate the panic potential of a social media message?*

In our current model, we focus upon panic potential of a message, taking the psychological definition of panic into account (as an extension of fear), and applying this in the domain of text mining. We examine metrics and features based on the sentiment and emotion classification methodologies in social media, as highlighted in earlier literature.

## 3 Methodology

### 3.1 Twitter Analysis of Munich Shootings

On 22<sup>nd</sup> of July 2016, a malicious shooting incident occurred near the Olympia shopping mall in Munich. 10 people were killed and 36 were injured in this shooting<sup>1</sup>. We chose the Munich shooting incident for analysis, since it was an example of an event where social media data was substantially used to spread information and opinions. We used Twitter as our data source for sentiment and emotion classification, and took advantage of the most commonly used hashtags #München, #Schießerei, #OEZ, #öffenetür and in addition used a language detector to acquire German tweets. In order to improve the effectiveness, only tweets within a week of the event, i.e. from 22 July to 28 July, were extracted. The data set was lemmatized and stop words, hashtags, references and hyperlinks were removed, and a spell corrector was applied to improve the readability.

Overall, 19033 German tweets were extracted, of which 89% were posted in the first two days, with a peak point at 19:00 h on 22<sup>nd</sup> July, the day of the incident, and the number of tweets decreasing with time. The earliest tweet occurred at 18:00, 22<sup>nd</sup> July 2016, which was possibly posted by a witness of the shooting. This indicates that in this kind of emergency social media reacts faster than traditional media. Yet traditional media and authorities tend to be more trustworthy as shown in Figure 1 - that authorities and traditional media generated the most retweets, as shown by the top 10 most retweeted accounts.

---

<sup>1</sup> <https://www.theguardian.com/world/2016/jul/22/munich-shopping-centre-evacuated-after-reported-shooting-germany>

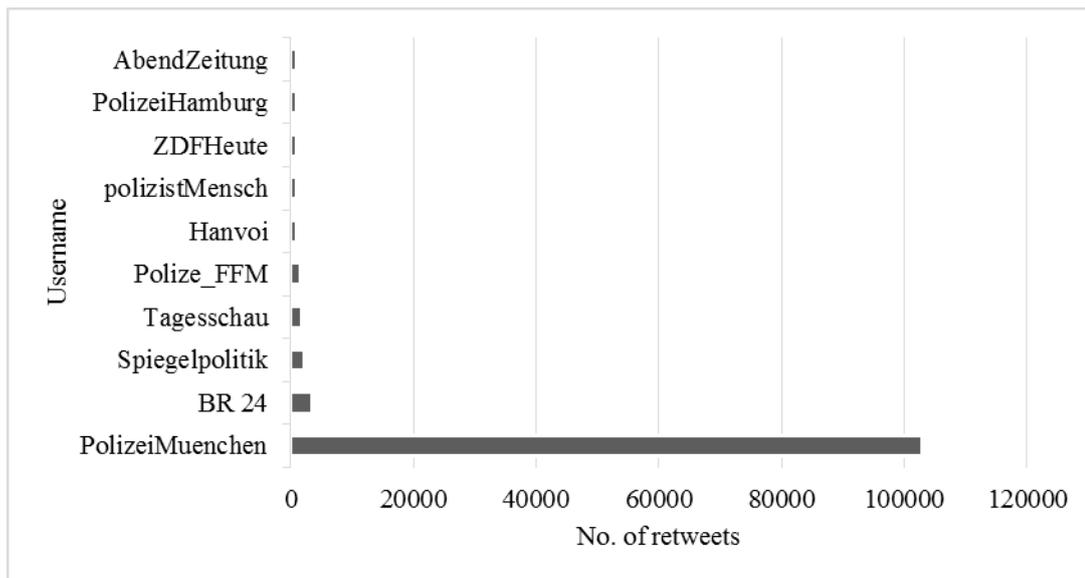


Figure 1. Top 10 username/accounts that were re-tweeted, in the first seven days

### 3.2 Features for the panic detection module

We randomly sampled 1000 tweets from the body of extracted tweets. Each Tweet contained more than 27 characters and hyperlinks were removed. These tweets were pre-processed (c.f. Appendix A for a detailed procedure), and independently annotated by 3 German native speakers, to provide a comparison of readers' perception with classifier-detected emotions (Figure 3, Labelling 1). Each tweet was rated from 0 to 100 for each of six emotions and from -100 to 100 to indicate positive/negative sentiment. The maximal rating of the emotion dimension was chosen to be the overall emotion label, otherwise the tweet was labelled as non-emotional. For sentiment, values greater than 30 were labelled as positive, and lesser than -30 were labelled as negative sentiment. 708 tweets are effectively used by the panic detection module, the other 282 were either irrelevant or were not readable for the annotators. All the data sets were classified into sentiment (positive/negative/no-sentiment) and one of the Ekman (1992)'s six emotions (i.e., surprise, fear, disgust, joy, sadness, anger) plus no-emotion with the help of naïve Bayes, Support Vector Machine, random forest and decision tree algorithms. The performance of the algorithms were compared based on precision, recall, F-score, and overall accuracy of classification.

It appears that when such a malicious incident as the Munich shooting occurs, people express their emotions intensively in the first few days before they start to think about their posts. However, the total number of emotional words, as well as the number of tweets was found to decrease with the days. Figure 2 illustrates how each emotion changes with each passing day in the first week, using the Naïve Bayes Classifier (Figure 3, Classification 1). Joy is the least emotion expressed on all days, which meets our expectation, and no-emotion is always the highest except the last day. No-emotion includes tweets which were annotated as informational tweets, without necessarily having an emotional content. An interesting finding is that the emotion fear was the highest on the first day, decreases on the next three days but increases in the last three days. Besides that, anger and surprise rise fast on the first three days, anger remains high while surprise drops rapidly from 25<sup>th</sup> to 26<sup>th</sup> of July.

Turning to the first research question, we randomly sampled 500 tweets from the above body of extracted tweets, omitting non-relevant tweets. These tweets were again independently annotated by German native speakers, to obtain ground truth about the panic level of a tweet. Participants hence rated the panic potential of a tweet, by answering the questions, "How likely do you think this tweet can incite panic?" Each tweet was rated as "None", "Low", or "High," to indicate no panic, low panic, high panic, respectively. These annotations were performed by 5 independent labellers, and a majority

voting was taken to determine the final panic potential of the tweet, and to ensure inter-rater reliability (Figure 3, Labelling 2). The data sets were classified into panic or no panic, and panic tweets were further classified as low and high panic with the help of naïve Bayes, Support Vector Machine, random forest and decision tree algorithms, and compared with the metrics of accuracy, precision, recall and F-score. Preliminary results showed that random forest and decision trees performed better than SVM and Naïve Bayes, achieving average classification accuracies of up to 70% (Figure 3, Classification 2) and Table 1.

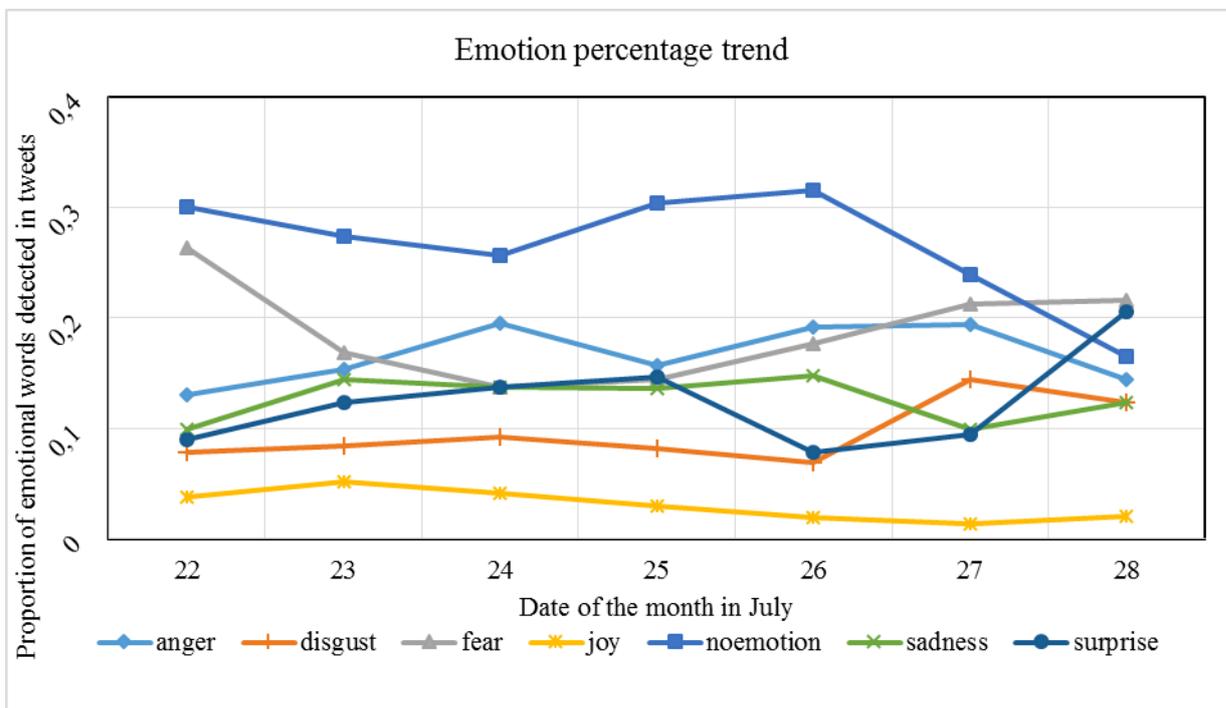


Figure 2. Emotion percentage trend on the first seven days

In addition to the labelled tweets, we consider additional features of the tweets (Figure 3, Classification 3). For instance, one expectation would be that the longer the tweets, the more probability it has to increase panic. In several tweets, this length was influenced by the length of hashtags, and the number of hashtags, where more hashtags were often used to increase the importance or attention drawn to a tweet, thus potentially increasing its panic level as well. In addition, punctuation marks that can be used to express emotions were considered, such as the number of exclamation marks, number of question marks. The presence of specific emoticons (such as :( and :)) was also checked, and added as features, although the proportion of tweets containing these emoticons was minimal, in the given dataset. The number of retweets of a given tweet was also considered, to verify whether it predicted the panic potential of a given tweet.

Finally, we expect that the number of positive and negative words, and the number of emotional words, as learnt from the above emotion and sentiment classifiers, is directly related to the overall panic potential of a tweet. Particularly, the number of negative words, and emotions of fear, might not only reflect fear in the writer, but also be capable of inciting fear in the reader, thus increasing the overall panic potential of the message (basing on Clark (1986)’s cognitive model for panic). Hence, these emotion and sentiment features are also used to train the classifier for detecting the panic potential. Preliminary results are presented in Table 1.

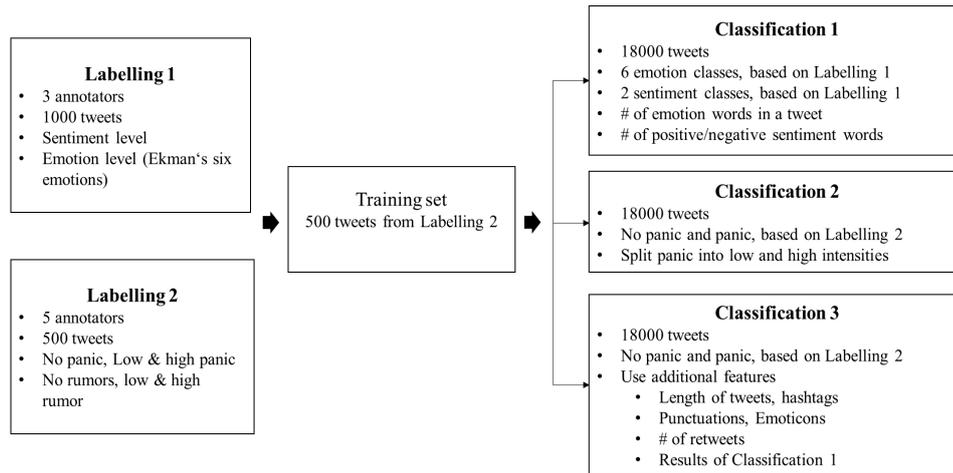


Figure 3. Proposed training and classification model

#### 4 Considerations for evaluation metrics for panic potential

Turning to the second research question, the proposed classifier (s) may be evaluated in several ways. To overcome differences in emotional, sentiment, and panic perceptions of annotators, a majority voting amongst annotator labels was taken, to establish ground truth. Collecting and creating a dictionary of words that impacted the votes of people, is another possibility to obtain ground truth. In addition to the bag of words, we propose to use additional dictionaries, such as TwitterMonitor, TwitInfo, Tweepent, etc. which have been used in the domain of detecting disaster events, based on Twitter data, multimedia information, and geospatial information (Imran et al. 2015). Gamifying the labelling process and create a cross-lingual disaster-specific dictionary remains a promising approach. In spite of different annotation and dictionary methods for improving the ground truth quality, and hence overall accuracy, the detection methods outlined above still rely heavily on the perceptions of the reader (in this case, the annotator), as well as the ability of the reader to access these emotional perceptions. Emotional perception research has shown that people vary in their ability to perceive their own emotions (Barrett et al. 2004). To judge how likely a message can incite panic in them, is a hypothetical situation, which might be often difficult for the participants to imagine and label. Hence, in addition to the annotated data for perceived panic potential, we propose to obtain ground truth by measuring the physiological effects of potential panic messages on the individual state, thus reducing differences in labelling due to emotion perceptions. Smartphone sensors could be used to gather real-time arousal data, to measure affective reactions to a given message (Rouast et al. 2017). To obtain less noisy data, collecting physiological data for specific messages in a controlled laboratory setting is another possibility.

The classifier detection accuracy is evaluated by means of standard classification metrics, such as precision, recall and average classification accuracy in comparison to the baseline accuracy. To construct and evaluate an event-independent panic potential detector, necessitates training data across different kinds of disaster situations (such as Verma et al. 2011). Second, cross-lingual semantic and text analysis poses a challenge that would impact both the choice of language of the classifier, as well as the required training sets (Balamurali, 2012). For the moment, we address this problem by using (German) native speakers for the annotation, and applying machine translation for computing the accuracy.

<i>Classification 2, Bag of words</i>	Avg. classifier accuracy	<i>Classification 3, Bag of words</i>	Avg. classifier accuracy
Algorithm		Algorithm	
Support Vector Machine	0,633	Support Vector Machine	0,633
Random Forest	<b>0,693</b>	Random Forest	0,673

Decision Tree	0,673	Decision Tree	0,683
<i>Classification 2, Bag of words +TFIDF</i>		<i>Classification 3, Bag of words + TFIDF</i>	
Support Vector Machine	0,633	Support Vector Machine	0,663
Random Forest	<b>0,694</b>	Random Forest	0,673
Decision Tree	0,663	Decision Tree	0,673
<i>Classification 2, word2vec</i>		<i>Classification 3, word2vec</i>	
Support Vector Machine	0,673	Support Vector Machine	0,612
Random Forest	0,643	Random Forest	0,653
Decision Tree	0,551	Decision Tree	0,653

Table 1. Results of Classifications 2 and 3, with Labelling 2 (Low, Medium, High Panic)

## 5 Conclusion & Future Work

Despite social media’s pivotal role in managing disaster-related information, the emergence of this media has generated the need to take the individual state into account. There is a dearth of literature which quantifies such metrics – such as panic potential, and rumour potential of messages, especially during disaster situations. In this work, we highlight the need to detect the panic potential of tweets, and propose a methodology to compute it learning from the emotions and sentiments expressed in Tweets. We present candidate features and evaluation metrics for detecting the panic potential of a tweet. We have now regarded the emotional perception of the reader of the message, which might differ from the emotional state of the poster, and how it is reflected in the message needs to be carefully understood in future work. In addition, differences in emotion perception and annotations would need to be validated by physiological methods for ground truth, and we expect the accuracy to improve with precise dictionary choices as well (including annotator dictionaries). Finally, the impact of panic potential information on subsequent information sharing behaviour, is a promising area of research.

The panic potential of a social media message being posted in a disaster situation is useful to strategize information sharing and forwarding patterns in disaster and crisis management, and subsequently regulating how it impacts the individual and societal emotional well-being. For instance, messages which seek or provide information about a particular geographical area, or status updates, are likely to be low in emotional content. Other messages with emotional content, need to be allowed to be forwarded with care, especially considering their impact on an already-panicked public. This could be achieved by setting different permission levels for messages to be re-tweeted, as well as providing warnings to individuals who share emotionally laden content during disasters. In this work, we emphasize on the necessity for social media providers and users to be *panic-aware*, and argue the need for users to share information with more consideration of the consequences, and detecting the panic potential of messages is a first step in this direction. In the future, remedial measures may also be taken on the basis of the panic potential (such as activating special disaster recovery modes in social media frameworks, which display the panic potential of messages to users, and recommend precautionary and responsible information sharing behaviour. The latter, however, needs to be designed with further experimentation, considering the impact on overall user experience and sharing behaviour. Taking the panic potential to a next degree, public services such as police and firefighters can monitor and react to unfolding disaster situations faster and more accurately, by assessing potential panic states of messages, to curtail public sentiment, and as a guide to managing information flow and credibility in a disaster situation. However, purely determining the decision to forward (or to not) messages and share information based on panic potential might have its limitations, and hence needs to be combined with source and information credibility as well as the context of sharing.

## References

- Acar, A., and Muraki, Y (2011). "Twitter for crisis communication: lessons learned from Japan's tsunami disaster," *International Journal of Web Based Communities* 7 (3), 392-402.
- Aman, S., and Szpakowicz, S (2008). "Using Roget's Thesaurus for Fine-grained Emotion Recognition," In *Third International Joint Conference on Natural Language Processing*, 312-318.
- Avvenuti, M., Cresci, S., Marchetti, A., Meletti, C., and Tesconi, M. (2014). "EARS (earthquake alert and report system): a real time decision support system for earthquake crisis management." In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* p.1749-1758
- Balamurali, A. R., A. Joshi, and P. Bhattacharyya (2012). "Cross-lingual sentiment analysis for Indian languages using linked wordnets." In *Proceedings of the International Conference on Computational Linguistics (COLING)*
- Barrett, L. F., Quigley, K. S., Bliss-Moreau, E., & Aronson, K. R. (2004). Interoceptive sensitivity and self-reports of emotional experience. *Journal of personality and social psychology*, 87(5), 684.
- Bozic, C., and Seese, D. (2011). "Neural networks for sentiment detection in financial text." In *Proceedings of the 14th International Business Research Conference*.
- Chen, R., and Sakamoto, Y. (2013). "Perspective matters: Sharing of crisis information in social media," In *System Sciences (HICSS), 2013 46th Hawaii International Conference*, 2033-2041.
- Clark, D. M. (1986). "A cognitive approach to panic," *Behavior research and therapy*, (24:4), 461-470.
- Coviello, L., Sohn, Y., Kramer, A. D., Marlow, C., Franceschetti, M., Christakis, N. A., and Fowler, J. H. (2014). "Detecting emotional contagion in massive social networks." *PloS one*, 9(3), e90315.
- Ekman, P. (1992). "An argument for basic emotions," *Cognition & Emotion*, 6 (3-4), 169-200.
- Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, 59(7), 96-104.
- Halse, S. E., Tapia, A., Squicciarini, A., & Caragea, C. (2016, May). Tweet Factors Influencing Trust and Usefulness during Both Man-Made and Natural Disasters. In *13th International Conference on Information Systems for Crisis Response and Management (ISCRAM) Rio de Janeiro, Brasil*.
- Honkela, T., Izzatdust, Z., and Lagus, K. (2012). "Text mining for wellbeing: Selecting stories using semantic and pragmatic features." In *International Conference on Artificial Neural Networks* (467-474). Springer Berlin Heidelberg.
- Houston, J. B. (2012). "Intervention across disaster phases," *Journal of Emergency Management*, 10 (4), 283.
- Hripsak, G., Bakken, S., Stetson, P. D., and Patel, V. L. (2003). "Mining complex clinical data for patient safety research: a framework for event discovery." *Journal of Biomedical Informatics*, 36(1), 120-130.
- Huang, C. M., Chan, E., and Hyder, A. A. (2010). "Web 2.0 and internet social networking: A new tool for disaster management? Lessons from Taiwan," *BMC medical informatics and decision making*, 10 (1), 57.
- Imran, M., Castillo, C., Diaz, F., and Vieweg, S. (2015). "Processing social media messages in mass emergency: A survey." *ACM Computing Surveys (CSUR)*, 47(4), 67.
- Kim, M. S., Valitutti, A., and R. a Calvo. (2010). "Evaluation of Unsupervised Emotion Models to

- Textual Affect Recognition,” In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, 62–70.
- Lazard, A. J., Scheinfeld, E., Bernhardt, J. M., Wilcox, G. B., and Suran, M. (2015). “Detecting themes of public concern: A text mining analysis of the Centers for Disease Control and Prevention's Ebola live Twitter chat.” *American journal of infection control*, 43(10), 1109-1111.
- Machajdik, J. and Hanbury, A. (2010). “Affective image classification using features inspired by psychology and art theory,” in *Proceedings of the 18th ACM international conference on Multimedia*, 83–92.
- Nagy, A., and Stamberger, J. (2012). “Crowd sentiment detection during disasters and crises.” In *Proceedings of the 9th International ISCRAM Conference*, 1-9.
- Palen, L. (2008). “Online social media in crisis events,” *Educause Quarterly*, 31 (3), 76-78.
- Pine, D. S., Klein, R. G., Mannuzza, S., Moulton, J. L., Lissek, S., Guardino, M., and Woldehawariat, G. (2005). “Face-emotion processing in offspring at risk for panic disorder.” *Journal of the American Academy of Child & Adolescent Psychiatry*, 44(7), 664-672.
- Reiss, S., Peterson, R. A., Gursky, D. M., and McNally, R. J. (1986). “Anxiety sensitivity, anxiety frequency and the prediction of fearfulness.” *Behaviour research and therapy*, 24(1), 1-8.
- Reynolds, B., and Seeger, M. W. (2005). “Crisis and emergency risk communication as an integrative model,” *Journal of Health Communication*, (10:1), pp. 43-55.
- Rouast, P. V., Adam, M. T., Cornforth, D. J., Lux, E., and Weinhardt, C. (2017). “Using Contactless Heart Rate Measurements for Real-Time Assessment of Affective States.” In *Information Systems and Neuroscience* (pp. 157-163). Springer International Publishing.
- Simon, A. Goldberg, L. Aharonson-Daniel, D. Leykin, and B. Adini. (2014) “Twitter in the cross fire—The use of social media in the Westgate Mall terror attack in Kenya” *PLoS One*, 9, e104136
- Stieglitz, S., and Dang-Xuan, L. (2013). “Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior,” *Journal of Management Information Systems*, 29 (4), 217-248.
- Verma, S., Vieweg, S., Corvey, W. J., Palen, L., Martin, J. H., Palmer, M., & Anderson, K. M. (2011, July). “Natural Language Processing to the Rescue? Extracting " Situational Awareness" Tweets During Mass Emergency.” In *International Conference on Web and Social Media*.
- Woo, H., Cho, Y., Shim, E., Lee, K., and Song, G. (2015). “Public trauma after the Sewol Ferry disaster: The role of social media in understanding the public mood,” *International Journal of Environmental Research and Public Health*, 12 (9), 10974-10983.
- Yang, Y.H., Liu,C.C., and Chen, H.H. (2006). “Music emotion classification: a fuzzy approach,” in *Proceedings of the 14th ACM International Conference on Multimedia*, 81–84.

## Appendix

### Procedure

To preprocess the text data, the Python package *nlk* was applied to provide the stopwords in German and *treetaggerwrapper* to lemmatize the text. *Sklearn* was used to transform the text into bag of words with *tf* and *tf-idf*. This package was also responsible for training, testing and evaluating. Beyond that, *numpy* and *pandas* built the data frames and matrices required for the classification. All pre-processed data sets were then transformed into Bag-of-words model together with term frequency–inverse document frequency to generate the vector space model of the data sets.