Pooling Tweets by Fine-Grained Emotions to Uncover Topic Trends in Social Media

Annika Schoene
IBM Research
IBM UK Ltd.
Daresbury, United Kingdom
annika.marie.schoene@ibm.com

Geeth de Mel
IBM Research
IBM UK Ltd.
Daresbury, United Kingdom
geeth.demel@uk.ibm.com

***Abstract*—In this paper, we present a lexicon-based sentiment analysis method that is used as an annotation scheme for identifying fine-grained emotions in social media topics. This methodology is based on Plutchik's wheel of emotion and Latent Dirichlet Allocation (LDA). We firstly annotate a tweet based on eight basic emotions and secondly we compute further eight *dyads* as a product of the basic emotions. We demonstrate that this lexicon-based approach achieves up to 78.53% ground truth accuracy when compared to human annotated data that is split into positive and negative polarities. Moreover, we investigate a novel means to identify trending topics in twitter data by utilizing LDA and focusing on fine-grained emotions associated with each tweet. We compare the most dominant emotions in social media as topics from an emotion-document pooling strategy and compare the results to an author-topic modeling strategy.**

***Index Terms*—Fine-Grained Emotions, Emotion Lexicons, Emotion Fusion, Topic Modelling, Plutchik**

I. Introduction

Given the advancement in the technology, cyber social spaces are fast having an impact on the inferences users—be it humans or autonomous systems—make about their surrounding context. We refer to a network of humans and intelligent autonomous systems as a cyber social space [1] where there are links among human societies, autonomous societies, and between human and autonomous societies. In such settings, when the power of artificial intelligence (AI) capability increases, the influence such systems will have on the network will be amplified; there are multitude of recent examples [2], [3], where autonomous systems were used to polarise a social view or an outcome. Thus, it is important to have means to effectively mine the sentiment—especially the emotional contagion— [4] of such components, thus the network. In order to infer emotions, sentiment analysis is used and has been defined as a sub-discipline of Natural Language Processing with the aim of *determining valence, emotions, and other affectual states from text* [5].

Within the field of sentiment analysis, Twitter has become an increasingly popular data source, with mainly machine and deep learning models being submitted to shared tasks [6]. Most of this work has either focused on detecting polarities [7] in a tweet or used data that has been annotated by human annotators for additional information such as semantic interpretations [8]. Detecting fine-grained emotions from micro-blogging platforms such as Twitter has a number of wide-ranging applications such as tracking public mood [9] However, this is a difficult task given that content and dynamics in such platforms varies among topics as well as in the same conversation [10]. Additionally, micro-blogging data has also been used in topic modeling to uncover topics in user conversations [11] as well as detecting connections between emotions and affective terms [12]; this involves analysis of emotions conveyed towards specific topics, and how emotions of groups change over time [13].

Motivated by this observation, we first explore a lexicon-based approach for fine-grained emotion detection and then propose means to detect topics using the emotions expressed in the annotated tweets. Specifically, we use Latent Dirichlet Allocation (LDA) [14] to detect topics encapsulated in each individual tweet and use insights from our emotion annotation strategy to pool tweets into emotion groups and show which emotions are most dominant in topics. Finally, we evaluate our emotion annotation strategy against an existing lexicon called SentiWord Net [15], which is frequently used in Sentiment Analysis task and is annotated for polarities. We then compare our topic-based emotion pooling scheme to another successful model called the Author-Topic Model (AT Model) [16].

II. Background

Much research has been conducted using basic emotion theory such as Ekman's six [17]. In order to perform fine-grained analysis of emotions, in this work, we use Plutchik's [18] eight basic emotions—and associated less intense emotions—in a wheel as shown by Figure 1. The key ideas of this work are that *(a)* basic emotions have polar opposites (e.g., *joy* vs *sadness*); *(b)* emotions can have different intensities (e.g., *rage* is more intense than *anger*); and *(c)* emotions can be mixed or have derivative states that are called *dyads* (e.g., *anger* and *disgust* together will result in *contempt*).

Work by Kiritchenko *et al.* has developed two lexicons for sentiments analysis for short informal messages and is based on supervised text classification methods using Support Vector Machines [19]. Moreover, Abdul-Mageed and Ungar has derived Plutchik's eight basic emotions from a self-collected dataset of 24 fine-grained emotions using gated recurrent neural networks [20].

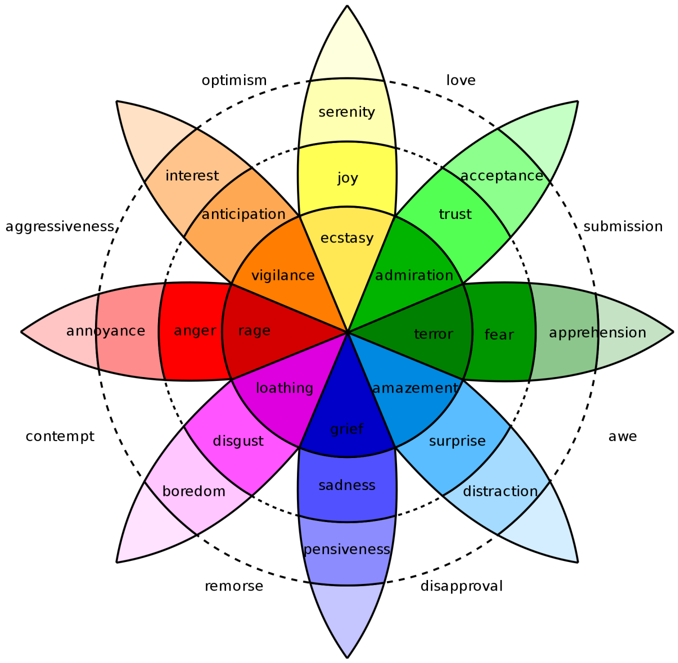

Fig. 1. Plutchik's Wheel of Emotion

a) Lexicon- and Knowledge- based Approaches: One of the most commonly used resources for lexicon-based sentiment analysis is WordNet Affect [21] and has annotations for affective states such as positive, negative and neutral polarities or emotions. Another popular lexicon is SentiWord Net [22] which also has annotations for polarities and has been used for a number of different studies such as derivation of prior polarities [15]. Taboada *et al.* developed a lexicon-based approach to infer sentiment from text [23]; their system is referred to as *SO-CAL* (Semantic orientation calculator).

b) Sentiment Analysis (SA) for Topic Modeling: A promising approach to detect emotions in topics is the use of LDA where the focus is on reviews (or blogs) with multiple sentiments and are calculated over the document [14]. Typically, the approach is aimed at detecting pre-determined sets of emotions [12] or Ekman's six basic emotions [24]. Thus, many SA models in this area amended the prior knowledge of the models with different lexicons or word lists and have therefore been deemed domain or topic-dependent as well as not easily transferable [25].

The work by Bao *et al.* uses LDA to discover topics from emotions and then generating affective terms from those topics [12]. The data for this study was taken from a popular Chinese social media platform called *Sina*, where users can rate entries based on eight given emojis. Kim *et al.* focused on detecting emotional influences and patterns in unannotated Twitter conversations using so called sentence-LDA based on Plutchik's emotion theory with a lexicon [26].

Meta-data increases the interpretability of topic models. Rosen-Zvi *et al.* have developed the author-topic model which is an extension to LDA where the model includes information about authors [16]. Mehrotra *et al.* has looked at pooling twitter topics using unmodified LDA models to aggregate tweets by hashtags [27]. This is mainly because, in micro-blogs, message quality is variable and language used is more colloquial when compared with traditional documents. Detecting topics in social media and their emotional content has a number of different applications—e.g., [28] have argued that due to constant political and social changes, there is a need to better understand the role of social media and how it influences individuals and groups [29] have proposed a statistical model for ideological discourse where it is hypothesized that some words are more frequently used when people speak about an ideological topic or hold a specific ideological perspective.

III. Approach

In order to infer fine-grained emotions and to indentify the emerging topics w.r.t. those emotions, we do the following: *(a)* we employ a lexicon-based approach for fine-grained emotion detection in tweets—both for text content and hashtags; *(b)* we then devise a strategy to infer added emotions so that additional information is available to deduce the polarity of the identified emotions; and *(c)* we employ means to detect topics using the emotions expressed in the annotated tweets so as to pool tweets into emotion groups and show which emotions are most dominant in topics. We used both the EmoLex—i.e., National Research Council Canada (NRC) Emotion Word lexicon—which was manually created by [30] and the automatically created NRC hashtag lexicon by [31] for our emotion annotation scheme.

A. Tweet Annotation

We annotated each tweet through the emotion lexicon for each word and as a result we have a number of basic emotions associated to each tweet. Then we analyzed which emotions were predominately used in each tweet through the Algorithm 1. Motivated by this observation, we developed three different strategies for finding the dominant emotion for each tweet. Every tweet that only had one basic emotion we kept this emotion as the main indicator. For two or more dominant emotions, we checked if there are any matches between the basic emotions which can create a less strong dyad as per Algorithm 2.

Moreover, for tweets with only two highest emotions, the algorithm searches for a match between the second highest emotion. For tweets that did not have one basic emotion or no matches, we used the emotions expressed by hashtags to detect the emotions. For example, let us assume a tweet *I hate dogs but love cats #truetalk*. After steps 2–9, the *emotionlist* contains {(anger:3),(love:1),(disgust:2), (trust:0)}. After steps 10–19, we have (anger:3), (love:1), and (disgust:2) as our three highest emotions. In the case of emotion values being zero in the above example, we would have no emotion annotation. However, we would then consider the relevant

```
input : Preprocessed Tweets
output: Tweets annotated for three most dominant
        basic emotions
/* Each word in a tweet is annotated
   through a lexicon                        */
emotionlist = []
for ( word in Tweet ) {
    if word in Lexicon then
        map word to emotion;
        add each emotion to emotionlist ;
    else
        continue;
    end
}
/* Count how many times emotions
   occur per tweet                          */
for ( emotion in emotionlist ) {
    if emotion value > 0 then
        count each emotion;
        return three highest emotions per tweet;
    else if emotion = 0 or emotion = same value
      then
        return None;
    else
        return None;
    end
}
```

Algorithm 1: Emotion Annotation

```
input : Tweets annotated for basic emotions
output: Tweets annotated for basic emotions and
        dyads
/* Dyads are calculated for tweets
   that have more than one basic
   emotion                                  */
for ( emotion in emotions ) {
    if emotion = 1 then
        return basic emotions;
    else if emotion ≤ 2 then
        return dyad;
    else
        return None;
    end
}
```

Algorithm 2: Calculating Dyads

hashtags for potential emotions through Algorithm 3. Additionally, if a tweet has more than six distinct emotions as the dominant emotion or no annotation at all we discarded the tweet—this is because it indicates that *(a)* the number of potential dyad matches increases *exponentially* so that there is no clear dominant emotion; or *(b)* none of the words used in a tweet are not present in the lexicon. Furthermore we argue that not all tweets are of the intensity of a basic emotion and there might be multiple basic emotions expressed in one tweet that are of similar *polarity* but the true sentiment of the tweet might be a less strong dyad emotion. Taking the above example, let's assume the *emotionlist* returns the following emotions: {(anger:3),(love:1),(disgust:3), (trust:0)}. In this case Algorithm 2, will check if there is a match between the two highest emotions *anger* and *disgust*, which would be true for this example and *contempt* is returned. Finally, if there is a negation present in the tweet, such as *not* we use Algorithm 4 to invert an emotion. An example of this would be where the original annotation for a tweet is *joy*, but due to the negation annotation is inverted to *sadness*.

B. Hashtag Annotation

We treated hashtags separately from the main tweet content as it has been shown in previous research to have emphasizing emotions or change the sentiment of presumably a *neutral* content [31]. Furthermore, it has been argued that hashtags can be seen as annotations provided by a user for topics and emotions [32]. Therefore, we employed a similar approach—i.e., after preprocessing tweets for emotion annotations, we used hashtags as a main indicator where the main tweet did not return a main basic emotion or dyad Algorithm 3.

```
input : Tweets with no emotion annotation
output: Tweets with emotion annotation through
        hashtag
/* Tweets are annotated for emotions
   through a hashtag lexicon                */
for ( hashtag in Tweet ) {
    if hashtag in Lexicon then
        map hashtag to emotion;
        return each emotion with value 1;
    else
        return None;
    end
}
```

Algorithm 3: Emotion annotation with hashtags

C. Dealing with Negations in Tweets

Negations play an important role when detecting emotions in textual data and previous work has inversed polarities for each tweet, however, there has been evidence that this technique is not always applicable [19]. For the purpose of this work, we still inversed emotions as this is a different approach compared to the polarity inversion in the previously mentioned work. We inversed the dominant emotion and dyad of each tweet if a negation occurred as per Algorithm 4 adhering to the axes of Plutchik's wheel of emotion [33].

```
/* Negations contains a list of
   negations in the English,
   emotionpairs is a dictionary with
   opposing emotions according to
   [33]                               */
input : Tweets annotated for basic emotions and
        dyads, negations, emotionpairs
output: Final tweet annotation
for ( word in Tweet ) {
    if word in negations then
        invert emotion according to emotionpairs;
        return inverted emotion;
    else
        do nothing;
    end
}
```

Algorithm 4: Negations in tweets

IV. Evaluation: Emotion Tagging

In order to evaluate our emotion tagging approach, especially the inference of dyads, we used three data sets—two of which already had some ground truth information in the form of polarities.

A. Data

We used three different Twitter datasets for our experiments. The first dataset is a collection of NBC's tweets [34] from Russian chat bots to influence the 2016 U.S. elections (henceforth referred to as Russia Data); this dataset has no annotation of polarities nor fine-grained sentiments. The second dataset contains tweets collected about the weather (henceforth referred to as Weather Data) which has been annotated for three different polarities—i.e., positive, negative and neutral—as well as a category called *not related to content* [35]. The last twitter dataset (henceforth referred to as 104K Data) has been collected based on query terms on varying topics where only positive and negative content was considered [36]. Figure 2 shows how many emotions are present in each dataset when the emotion annotation strategy is applied to the data sets.

B. Preprocessing

It has been noted by that extensive text preprocessing and labeling usually takes place prior to conducting experiments where longer documents are broken down into sentences [25]. We, therefore, applied a similar text preprocessing methods to our tweets, which included tokenization,lower casing all tokens,removing punctuation and stopwords. However, it is important to note that we excluded negations from the stopword removal process as negations in a tweet could inverse the overall emotions expressed in a tweet. Furthermore, we anonymised usernames, removed retweets and duplicates as well as web references.

The Weather and 104K Datasets were already preprocessed, thus there was not a significant drop in tweets after the preprocessing. Table I outlines general information on the datasets and gives an overview over the size of the datasets after each preprocessing step was performed. In order to take advantage of hashtags, we preprocessed hashtags separately by firstly splitting '#'from the full string and then split strings into possible words by using word lists; we used the first suggested option for further analysis.

Number of	Weather	Russia	104K
Original tweets	203482	1600000	1000
Preprocessed tweets	192558	1600000	1000
Tweets with annotations	133220	963844	737

TABLE I
DATASETS AFTER PRE -PROCESSING

In Figure 2 we demonstrate the distributions of basic emotions and dyads across the three datasets. In order to compare our emotion tagging approach and it's accuracy we have split emotions into two different polarities as depicted by Table II. Whilst we first considered to split emotions into three categories (positive, negative and neutral) we decided for the purpose and ease of this evaluation to split emotions into two categories, because there is some lack of clarity around some emotions associated polarity categories or orientations (i.e.: *surprise*) [37].

Polarity	Emotions
Positive	Joy, Trust, Love, Optimism
Negative	Sadness, Disgust,Remorse, Disapproval Submission, Fear, Awe, Surprise, Contempt, Anger,Aggressiveness, Anticipation

TABLE II
EMOTION SPLIT INTO POLARITIES

Table III shows the ground truth of the human annotated weather dataset and 104K dataset compared to our own emotion split as well as a comparison to SentiWordNet. We have found that for positive polarities we achieved similar accuracies, however for negative polarities we have found a distinct difference. For the Weather data accuracy is significantly higher compared to positive polarities and 104K data results. The results in Table III show that based on our emotion split our approach performs better on negative emotions and SentiWordNet better on positive emotions. Furthermore it can be seen that our approach is performing significantly better when detecting positive polarities compared to SentiWordNet's ability to detect negative polarities. However, it has to be taken into consideration that according to our emotion split there are more *negative* than *positive* emotions and that our approach can be split down further into more fine-grained categories.

V. Evaluation: Topic Modeling with Emotions

For our experiments we used LDA [14] to detect topics in our datasets. We conduct 3 different experiments to show *(a)* how tweet topics change when tweets are treated individually when compared to a pooled document; and *(b)* which emotions are dominant in each topic for each dataset.

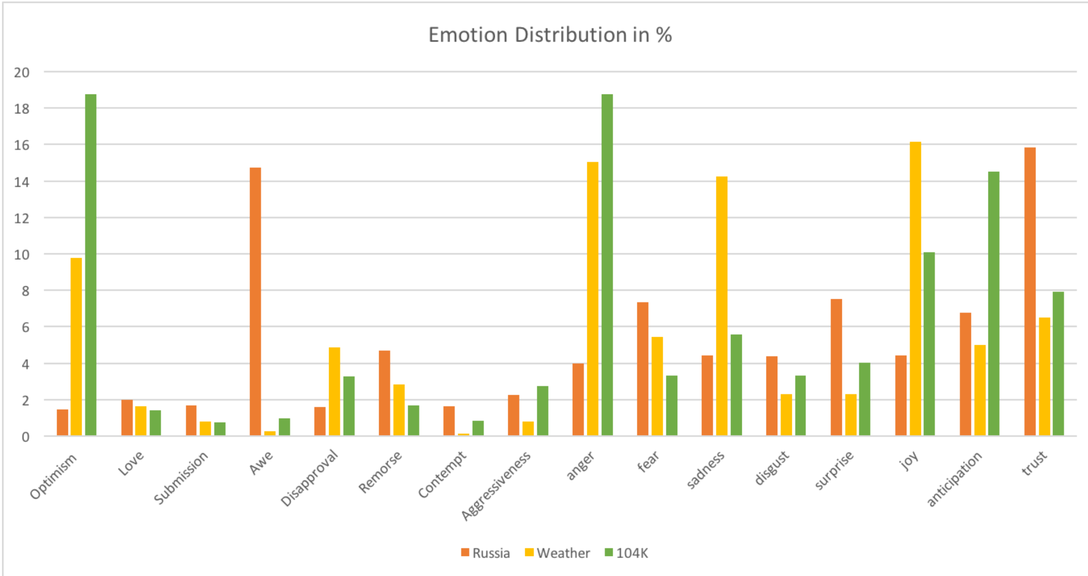

Fig. 2. Emotion Distribution in %

Polarity	Weather Data	104K Data
Positive (Annotation Scheme)	57.63	55.73
Negative (Annotation Scheme)	78.53	76.38
Positive (SentiWordNet)	65.53	79.12
Negative (SentiWordNet)	34.46	41.58

TABLE III
EMOTION ANNOTATION GROUND TRUTH IN %

- **Experiment 1** Each tweet is treated as an individual document to extract dominant topics
- **Experiment 2** All tweets belonging to one emotion annotation are combined into one document to extract dominant topics
- **Experiment 3** The Author-Topic Model is applied to see difference in Topics

The reason for pooling tweets into one document is to overcome the issue of sparsity in twitter data—this is due to the length of each tweet available is shorter than traditional documents used for training LDA models. We show that pooling tweets by emotion produces coherent topics for both sparse and less sparse data which can clearly be associated to emotions. Finally, we compare our pooling strategy to the *Author-Topic-Model*. We have chosen this model as parallels can be drawn from the *author* of a document to the dominant *emotion* of a tweet. We used topic coherence scores to evaluate the results of our models by using UCI coherence measure implemented in the Gensim library [38].

After obtaining all emotion annotations for each dataset we used all tweets that had a first basic emotions and dyads in each model. There are some tweets which have more than one basic emotion assigned to them, however in order to establish a baseline model we only consider tweets with one dominant emotion.

a) Experiment 1: Table IV shows the topics for each dataset, where the number of topics was based on the best coherence score for each dataset. Furthermore, each topic was named intuitively based on the words that occurred most commonly in these topics.

Topic	Weather	Russia	104K
1	storm/tornado	support America	watching Show
2	hot/sunshine	terrorism	thankfulness
3	weather/storm	attacks email	friends
4	rain/snow	guns	another year
5	sunny/may	black rights	well wishes
6	weather/link/user	trending news	weekend/music
7	mixed weather	police violence	game/luck
8	weather warning	vote white house	school
9	-	time	beautiful summer
10	-	presidential candidates	time/work

TABLE IV
RESULTS EXPERIMENT 1 - DATASET 1,2 AND 3

b) Experiment 2: After pooling each tweet by its emotion annotation we conducted further experiments to see which topic is most dominant for each emotion in Table V. We did not use minimum count of tokens and maximum frequency in our experiments with the Weather Data, because there is only a small amount of tweets and a total of 2721 tokens. For experiments with the Russia data we have adjusted the minimum number of words to 3 and the maximum frequency of occurrence to 0.5, where the maximum number of tokens was set to 7508. Experiments with the 104K dataset we used 100000 tokens and adjusted the maximum occurrence frequency to 0.5 and the minimum occurrence of a word to 3. We have observed that for results of Experiment 2 that there was an overlap of words used in topics of basic emotions and dyads that are close to each other on the *wheel of emotion*. This may indicate that (1) topics might be associated to similar emotions and (2) words used in certain topics may be specific

Emotion	Weather Data	Russia Data	104K Data
Joy	sunshine/beautiful	Trump/supporter	love/baby
Trust	snow/patrol	president/Trump l	school/follower
Fear	tornado/hot	war/police	change/homework
Surprise	weather/chance	Trump/Donald	Trump/Donald
Sadness	weather/rainy	black/refugee	lost/cry
Disgust	cold/weather	Trump/Clinton	boy/crap
Anger	morning/patchy	Hillary/Obama/hit	hot/hit
Anticipation	time/storm/tomorrow	time/Trump/start	time/tomorrow
Love	sunshine/weather	Trump/Supporter	wonderful/smile
Submission	link/user	Clinton/Trump	hospital/hurt
Awe	morning/patchy day	Hillary/ #politics	good/day
Disapproval	storm/weather	Clinton/black/vote	no good/don't like
Remorse	hope/sunny/forecast	Trump/ #tcot	sick/shame
Contempt	window/freezing	Clinton/criminal	horrible/damn
Aggressiveness	weather/ storm	against/Hillary	day/get
Optimism	sunny/day	Trump/young/gift	good/hope

TABLE V

RESULTS EXPERIMENT 2 - MOST COMMON WORDS FOR THE MOST DOMINANT TOPIC

to corresponding emotions. These notions also correspond to findings with the AT model in Experiment 3.

c) Experiment 3: For our experiments with the AT model we replaced the *authors* of a document with our emotion annotations, which means that each tweet was treated as a document. Our results show that it is possible to use the AT model for detecting emotions in topics, however there are some observations to be made w.r.t. experiments conducted using the AT model. Firstly, we have found that when visualizing our AT model results that in many cases emotions were overlapping with topics to the extend that multiple topics were associated to one emotion. This meant that it was not possible to extract coherent or dominant topics with the associated emotions. When comparing the results in Table VI to Table V, it can be seen that we get more detailed insight into topics for each emotion. One of the reasons for this could be due to the fact that the AT model was originally tailored for larger documents and therefore tailoring it further may yield more coherent topics in the future.

Topic	Title
1	storm/hot
2	user/link
3	storm/coming/ area
4	sunny weather
5	snow/patrol
6	rainy day
7	degrees/thunderstorm
8	tornado/warning

TABLE VI

NUMBER OF TOPICS FOR EACH DATASET USING THE AT MODEL

VI. DISCUSSION

Whilst the main purposes of this work was to develop a baseline system to detect fine-grained emotions with regards to topics on social media, it also demonstrates the knowledge one can extract—especially from the Russia dataset. Below we summaries a number of findings that we believe useful in better understanding such diverse emotions and provides an implicit measure of influence.

- Through our emotion pooling strategy, we found that positive emotions are associated with Donald Trump, whilst negative emotions often included references to Hillary Clinton or Barack Obama as per Table V. Similarly, we found that topics such as black rights, crime and war were associated with emotions such as fear, sadness and disapproval.
- We also found that there are few usernames that were more frequently referenced in some topic that we later correlated with emotions, such as *@realdonaldtrump* was frequently mentioned in tweets associated to topics that belonged to positive emotion topics which included terms such as *#makeamericagreatagain, GoP, Voter and Supporter*. Whilst *@HillaryClinton* is frequently mentioned with terms such as *#neverhillary, illegal, email and debate* and associated to emotions such as contempt, disapproval and aggressiveness.
- Last but not least, a black news agency was frequently mentioned with terms such as *freedom, history, Islam and Hillary Clinton*. Again these topics were closely associated to emotions such as sadness, fear and contempt.

Thus, it can be argued that using both sentiment analysis and topic modeling combined can give a detailed insight into large datasets and help detect trends.

VII. CONCLUSION AND FUTURE WORK

In this work we have shown that it is possible to use a lexicon to determine fine-grained emotions in tweets. Moreover, we have demonstrated that it is possible to achieve a ground truth accuracy of 78.53% when comparing our annotation algorithm against a human annotated dataset. We have shown that by pooling tweets by their main emotions allows us to see more coherent topics compared to treating each tweet as a document. Future work will aim to further investigate lexicons and other knowledge-based approaches to develop domain- independent fine-grained emotion annotation systems. Additionally, we will investigate how to best integrate additional meta-data into topic models, such as the time when a tweet was posted in order to see how emotions in topics change over time.

ACKNOWLEDGMENT

This work was supported by the STFC Hartree Centre's Innovation Return on Research programme, funded by the Department for Business, Energy & Industrial Strategy.

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

[1] A. Sheth, P. Anantharam, and C. Henson, "Physical-cyber-social computing: An early 21st century approach," *IEEE Intelligent Systems*, no. 1, pp. 78–82, 2013.

[2] P. Wang, R. Angarita, and I. Renna, "Is this the era of misinformation yet? combining social bots and fake news to deceive the masses," in *The 2018 Web Conference Companion*, 2018.

[3] E. Hunt, "Tay, microsoft's ai chatbot, gets a crash course in racism from twitter," *The Guardian*, vol. 24, 2016.

[4] J. Jouhki, E. Lauk, M. Penttinen, N. Sormanen, and T. Uskali, "Facebook's emotional contagion experiment as a challenge to research ethics," *Media and Communication*, vol. 4, 2016.

[5] S. M. Mohammad, "Sentiment analysis: Detecting valence, emotions, and other affectual states from text," in *Emotion measurement*. Elsevier, 2016, pp. 201–237.

[6] S. Rosenthal, P. Nakov, S. Kiritchenko, S. Mohammad, A. Ritter, and V. Stoyanov, "Semeval-2015 task 10: Sentiment analysis in twitter," in *Proceedings of the 9th international workshop on semantic evaluation (SemEval 2015)*, 2015, pp. 451–463.

[7] E. Kouloumpis, T. Wilson, and J. D. Moore, "Twitter sentiment analysis: The good the bad and the omg!" *Icwsm*, vol. 11, no. 538-541, p. 164, 2011.

[8] H. Saif, Y. He, and H. Alani, "Semantic sentiment analysis of twitter," in *International semantic web conference*. Springer, 2012, pp. 508–524.

[9] J. Bollen, H. Mao, and A. Pepe, "Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena." *Icwsm*, vol. 11, pp. 450–453, 2011.

[10] Y. Wang, E. Agichtein, and M. Benzi, "Tm-lda: efficient online modeling of latent topic transitions in social media," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2012, pp. 123–131.

[11] D. Alvarez-Melis and M. Saveski, "Topic modeling in twitter: Aggregating tweets by conversations." *ICWSM*, vol. 2016, pp. 519–522, 2016.

[12] S. Bao, S. Xu, L. Zhang, R. Yan, Z. Su, D. Han, and Y. Yu, "Joint emotion-topic modeling for social affective text mining," in *Data Mining, 2009. ICDM'09. Ninth IEEE International Conference on*. IEEE, 2009, pp. 699–704.

[13] S. Stieglitz and L. Dang-Xuan, "Social media and political communication: a social media analytics framework," *Social Network Analysis and Mining*, vol. 3, no. 4, pp. 1277–1291, 2013.

[14] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.

[15] M. Guerini, L. Gatti, and M. Turchi, "Sentiment analysis: How to derive prior polarities from sentiwordnet," *arXiv preprint arXiv:1309.5843*, 2013.

[16] M. Rosen-Zvi, T. Griffiths, M. Steyvers, and P. Smyth, "The author-topic model for authors and documents," in *Proceedings of the 20th conference on Uncertainty in artificial intelligence*. AUAI Press, 2004, pp. 487–494.

[17] P. Ekman, "An argument for basic emotions," *Cognition & emotion*, vol. 6, no. 3-4, pp. 169–200, 1992.

[18] R. Plutchik, "The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice," *American scientist*, vol. 89, no. 4, pp. 344–350, 2001.

[19] S. Kiritchenko, X. Zhu, and S. M. Mohammad, "Sentiment analysis of short informal texts," *Journal of Artificial Intelligence Research*, vol. 50, pp. 723–762, 2014.

[20] M. Abdul-Mageed and L. Ungar, "Emonet: Fine-grained emotion detection with gated recurrent neural networks," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, vol. 1, 2017, pp. 718–728.

[21] C. Strapparava, A. Valitutti *et al.*, "Wordnet affect: an affective extension of wordnet." in *Lrec*, vol. 4. Citeseer, 2004, pp. 1083–1086.

[22] S. Baccianella, A. Esuli, and F. Sebastiani, "Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining." in *Lrec*, no. 2010, 2010, pp. 2200–2204.

[23] M. Taboada, J. Brooke, M. Tofiloski, K. Voll, and M. Stede, "Lexicon-based methods for sentiment analysis," *Computational linguistics*, vol. 37, no. 2, pp. 267–307, 2011.

[24] C. O. Alm, D. Roth, and R. Sproat, "Emotions from text: machine learning for text-based emotion prediction," in *Proceedings of the conference on human language technology and empirical methods in natural language processing*. Association for Computational Linguistics, 2005, pp. 579–586.

[25] C. Lin and Y. He, "Joint sentiment/topic model for sentiment analysis," in *Proceedings of the 18th ACM conference on Information and knowledge management*. ACM, 2009, pp. 375–384.

[26] S. Kim, J. Bak, and A. H. Oh, "Do you feel what i feel? social aspects of emotions in twitter conversations." in *ICWSM*, 2012.

[27] R. Mehrotra, S. Sanner, W. Buntine, and L. Xie, "Improving lda topic models for microblogs via tweet pooling and automatic labeling," in *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2013, pp. 889–892.

[28] W.-H. Lin, E. Xing, and A. Hauptmann, "A joint topic and perspective model for ideological discourse," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2008, pp. 17–32.

[29] J. Skinner, "Social media and revolution: The arab spring and the occupy movement as seen through three information studies paradigms," *Working Papers on Information Systems*, vol. 11, no. 169, pp. 2–26, 2011.

[30] S. M. Mohammad and P. D. Turney, "Crowdsourcing a word–emotion association lexicon," *Computational Intelligence*, vol. 29, no. 3, pp. 436–465, 2013.

[31] S. M. Mohammad and S. Kiritchenko, "Using hashtags to capture fine emotion categories from tweets," *Computational Intelligence*, vol. 31, no. 2, pp. 301–326, 2015.

[32] X. Wang, F. Wei, X. Liu, M. Zhou, and M. Zhang, "Topic sentiment analysis in twitter: a graph-based hashtag sentiment classification approach," in *Proceedings of the 20th ACM international conference on Information and knowledge management*. ACM, 2011, pp. 1031–1040.

[33] R. Plutchik, "Emotions: A general psychoevolutionary theory," *Approaches to emotion*, vol. 1984, pp. 197–219, 1984.

[34] B. Popken. (2018) Twitter doesn't make it easy to track Russian propaganda efforts — this database can help. [Online]. Available: https://www.nbcnews.com/tech/social-media/now-available-more-200-000-deleted-russian-troll-tweets-n844731

[35] crowdflower. (2016) Weather sentiment. Accessed: 2018-06-30. [Online]. Available: https://data.world/crowdflower/weather-sentiment

[36] A. Go, R. Bhayani, and L. Huang, "Twitter sentiment classification using distant supervision," *CS224N Project Report, Stanford*, vol. 1, no. 12, 2009.

[37] E. Cambria, D. Das, S. Bandyopadhyay, and A. Feraco, *A practical guide to sentiment analysis*. Springer, 2017, vol. 5.

[38] R. Rehurek and P. Sojka, "Gensim–python framework for vector space modelling," *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic*, vol. 3, no. 2, 2011.