Capturing Signals of Enthusiasm and Support Towards Social Issues from Twitter

Shubhanshu Mishra
School of Information Sciences
University of Illinois at Urbana-Champaign
Champaign, Illinois
smishra8@illinois.edu

Jana Diesner
School of Information Sciences
University of Illinois at Urbana-Champaign
Champaign, Illinois
jdiesner@illinois.edu

ABSTRACT

Social media enables organizations to learn what users say about their products online, and to engage with their potential audiences. Social media has also been allowing individual users and the public to signal their enthusiasm, support, or lack thereof for a broad range of topics. In this paper, we analyze the robustness of a prior framework for tagging tweets across the dimensions of enthusiasm (labels: enthusiastic, passive) and support (labels: supportive, non-supportive). We investigate the quality of annotations in a collection of tweets about three topics, namely, cyberbullying, LGBT rights, and Chronic Traumatic Encephalopathy (CTE) in the National Football League. We train models that achieve >70% and 80% F1 score for classifying tweets for enthusiasm and support, respectively. We assess how text-based signals of enthusiasm and support vary depending on the different annotators. Finally, we propose and demonstrate a network analysis-based approach for combining the annotated tweets with account and hashtag mention networks. This step helps to identify top accounts and hashtags related to the considered categories (enthusiasm and support). Our work offers an alternative or supplemental classification schema and prediction model to standard sentiment analysis and stance detection.

CCS CONCEPTS

• Human-centered computing → Social networks; **• Computing methodologies → Information extraction**; **• Information systems →** *Computational advertising*;

KEYWORDS

Social media, Twitter, Enthusiasm, Support, Networks

ACM Reference Format:

Shubhanshu Mishra and Jana Diesner. 2019. Capturing Signals of Enthusiasm and Support Towards Social Issues from Twitter . In *5th International Workshop on Social Media World Sensors (SIdEWayS'19), September 17, 2019, Hof, Germany*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3345645.3351104

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 ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIdEWayS'19, September 17, 2019, Hof, Germany
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6903-9/19/09...$15.00
https://doi.org/10.1145/3345645.3351104

1 INTRODUCTION

On social media, users can express their thoughts and opinions. Organizations (corporate or not) are sometimes interested in knowing what users say about them and their products on social media. In the context of this paper, we refer to products in a broad sense, including commercial goods, services, campaigns, and initiatives related to topics of public interest, including social issues. Sometimes, organizations are also interested in identifying *influencers* [1], i.e., accounts whose content receives large audiences, or who might be likely to promote their products. Identifying influencers supports organizations in targeting their marketing and communication strategies towards specific (collections of) accounts that can propagate their message across intended audiences, which can be more efficient than individual audience acquisition. This situation has resulted in the creation of classification systems that help to identify relevant information in social media posts. The most common use case of content classification is identifying opinions expressed in a given tweet [4, 5] as *positive*, *negative*, or *neutral*; a process also known as sentiment analysis or opinion mining. Since opinions are often expressed with respect to a reference point, e.g., a product, topic, or event, research has also been done on stance detection (identifying if a tweet is pro or against a certain reference point) [10, 17], and target or aspect based sentiment analysis (where sentiment towards an entity or topic is identified) [6].

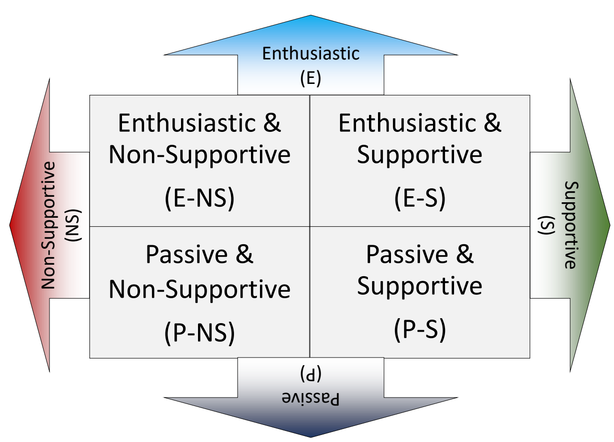

Figure 1: EPSNS classification schema with dimensions enthusiasm (labels: enthusiastic, passive) and support (labels: supportive, non-supportive).

In this paper, we combine the ideas of influencer identification and content classification to develop a framework for identifying enthusiastic and supportive tweets, which can be aggregated on the account level, by using a collection of tweets about three topics of public interest, namely, *cyberbullying*, *LGBT rights*, and *Chronic Traumatic Encephalopathy (CTE) in the National Football League (NFL)*. We reuse the tweet classification schema and data introduced in Mishra et al. [7] to categorize tweets across two dimensions, namely, **enthusiasm (labels: enthusiastic, passive** and **support (labels: supportive, non-supportive**. We validate the quality of the annotated data introduced in [7] through a series of inter-annotator and cross-topic evaluations. We develop machine learning classifiers for predicting if a tweet is enthusiastic or passive, and supportive or non-supportive, using text based features (as opposed to text and meta-data based features that we used in [7]). We use these models to identify features that signify enthusiasm and support towards the considered topics. Finally, we introduce an algorithm based on a weighted version of personalized PageRank [2, 11, 19], which combines the annotated tweets with account and hashtag mention networks. The algorithm supports the identification of top accounts across the dimensions of enthusiasm and support.

We find that the classifiers built using the datasets from Mishra et al. [7] achieve high F1 scores within and across the three topics considered in this paper. The set of top accounts and hashtags identified with our personalized PageRank approach includes accounts that were not found by using the general PageRank approach. We also contribute a unified classifier trained on the three datasets, along with an open source tool for classifying tweets and identifying top accounts and hashtags. Our approach supports the identification of Twitter accounts that enthusiastically support a topic or issue of public interest. We make the code for replicating our analysis publicly available[1].

2 SCHEMA FOR ENTHUSIASM AND SUPPORT CLASSIFICATION

We utilize the *enthusiastic, passive, supportive, and non-supportive (EPSNS)* orthogonal classification schema described in Mishra et al. [7] as the basis of tweet classification (see figure 1). This schema uses the following label definitions:

(1) **Enthusiastic**: Sender includes personal expression of emotion or call to action for others regarding a topic or issue.
(2) **Passive**: Lack of clear emotive content or call to action.
(3) **Supportive**: Actively showing favor for issue through use of outright statements/words of support. For the topics of cyberbullying and CTE, some supportive tweets use negative words (towards the issue), and by doing so, show support of the issue, i.e., being anti-bullying/CTE.
(4) **Non-supportive**: Being actively against the issue. These examples can be very similar in tone and content to supportive tweets in that they do show enthusiasm (exclamation, use of capitalization, etc.), but in blatant non-support of the issue. For the cases of cyberbullying and CTE, some tweets use positive words (towards the issue), and by doing so, are not supporting of the issue.

[1] https://github.com/napsternxg/TwitterEnthusiasmSupport

This classification system allows us to capture the level of **enthusiasm** and **support** towards a topic.

2.1 Comparison with existing classification tasks for tweets

The EPSNS classification schema used in Mishra et al. [7] is more suitable for identifying enthusiastic and supportive tweets compared to the positive, negative, and neutral classification schema commonly used in for sentiment analysis [12, 13, 18]. Figure 2 presents a few examples from the data, which illustrate a comparison of the EPSNS classification schema to the standard sentiment classification schema. This comparison shows that EPSNS captures different signals than sentiment classification. Overall, we consider the EPSNS schema as being is supplemental to, but not redundant with the sentiment analysis schema.

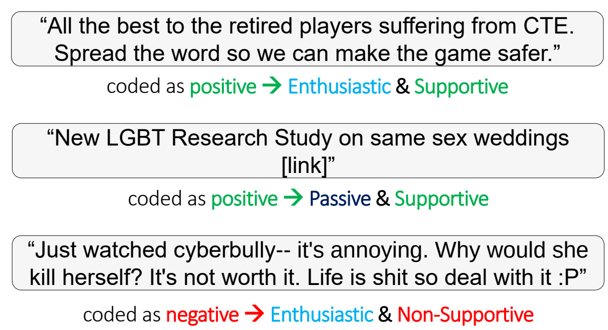

Figure 2: Example of using the enthusiasm and support dimensions instead of classic sentiment labels.

The EPSNS classification task is similar to the general stance detection task, where, given a topic, a model tries to identify the stance of a text towards a given topic. Since stance detection attempts to determine if a text is in favor of or against a topic [10], the stance dimension can be considered as similar to the support dimension in the EPSNS schema. However, the EPSNS schema is tailored to identifying reactions towards topics where support and enthusiasm are applicable, which is not the case for all topics. Another closely related task to EPSNS classification is affect identification [9], which aims to predict emotions and valence based on text data. However, affect identification is meant to generalize across topics, while EPSNS classification as presented herein is not as context-free and was only tested on the outlined topics.

2.2 Comparison with account labeling

Account labeling by aggregating text-level labels has been proposed in [16] for political ideology identification [15], and to identify machiavellianism, narcissism, and psychopathy. However, these classification schemas are often account-centric and do not incorporate additional context or knowledge that might be necessary to label a tweet. Furthermore, these schemas aim at classifying an account using a large sample of tweets, along with profile information. Our approach is based on classifying tweets only, and thereafter aggregating the level of enthusiasm and support expressed in tweets relevant to the topic in order to get a account level measure of enthusiasm and support towards a topic.

3 DATA

Mishra et al. [7, 8][2] introduced a dataset of tweets collected on the following topics: **cyberbullying (CB)**, **Chronic Traumatic Encephalopathy (concussions) in National Football League (CTE)**, and **Lesbian, Gay, Bisexual, and Transgender rights (LGBT)**. Each tweet in this dataset was annotated by two annotators, ensuring that annotators provide the labels based solely on the text of the tweet and in the absence of their own opinions.

Lable	CTE		CB		LGBT	
Anno.	1	2	1	2	1	2
NS	39	35	23	30	43	46
S	156	166	190	206	227	233
E	215	181	227	226	195	181
P	232	201	82	84	207	221

Table 1: Annotator (Anno.) Label Stats for each dataset. E: Enthusiastic, P: Passive, S: Supportive, NS: Non-supportive.

The distribution of the number of tweets per label in the dataset is shown in table 1. Although the dataset for a specific issue can be small (sometimes yielding less than 100 samples for certain classes), the resulting dataset has a high inter-annotator agreement identified using percentage agreement (% =) as well as Cohen's κ [3] (see table 2). However, a limitation of Cohen's κ is that it uses a baseline of chance agreement [14], which may hide disagreement. In this paper, we conduct several additional experiments to assess the similarity of annotator labels for training machine learning models. These experiments allow us to confirm if tweets with similar features have similar annotations.

	Enthusiasm			Support		
	%=	κ	N	%=	κ	N
CTE	0.96	0.91	379	0.98	0.92	165
CB	0.94	0.86	309	1.00	1.00	209
LGBT	0.93	0.87	395	0.97	0.89	257

Table 2: Inter annotator agreement between two annotators. % = is percentage agreement, and κ is Cohen's kappa.

4 EVALUATING DATA ROBUSTNESS FOR TRAINING CLASSIFICATION MODELS

We first describe the data from Mishra et al. [7], and then assess the robustness of the annotated data towards its suitability for training generalizable prediction models. The data were prepared for analysis as follows: The texts were tokenized using a Twitter tokenizer in NLTK[3]. Each term was lemmatized using the NLTK lemmatizer. A document is represented in terms of the TF-IDF score of its unigrams along with bigrams and trigrams (identified via pointwise mutual information in each data).

Our first analysis focuses on the top salient terms identified for each dataset, label, and annotator combination. The salient terms are identified using mean TF-IDF scores. Table 3 shows that for the majority of datasets and labels, the salient terms identified across the annotated tweets are highly similar across annotators. The similarity in top salient terms correlates with relatively high inter-annotator agreement scores presented in table 2. This finding also provides support for the claim that the data for each label are similar in their word distribution across the annotators. Table 3 reveals that salient terms for enthusiastic are more conversational compared to other labels, and express proclamations, e.g., *screw* (CTE), *emoticon* (CB), and *you's* (LGBT). The analysis of salient terms leads the way for our next experiments on assessing feature importance for training generalizable tweet classifiers.

The second analysis examines the quality of logistic regression models trained on data using labels from only one annotator. This allows us to assess the consistency of an annotator (similar to the concept of intra-coder reliability). We conduct three-fold cross-validation. All evaluation scores represent micro-F1 scores (unless specified otherwise). Table 4 shows that for each dataset, cross-validated models trained using labels from only one annotator result in $\sim$ 70% score for the dimension of enthusiasm and $\sim$ 83% for support. The scores are higher for identifying supportive versus non-supportive tweets compared to enthusiastic versus passive tweets. The high F1 scores are correlated with high inter-annotator agreement scores. Since this experiment relied on training and evaluation using a single annotator's labels, the F1 scores also provide information about the consistency in annotations by that annotator. Hence, it appears that it was easier for the annotators to label tweets for the support dimension than for the enthusiasm dimension. This can be due to codebook instructions or the nature of the texts.

In the third analysis, we identified if a model trained on two datasets (e.g., CTE+LGBT) and a single annotator's labels will transfer to the remaining dataset (e.g., CB) annotated by a different annotator. The results of this experiment are shown in table 5. We observe that a model trained on a given annotator's labels consistently gets comparable evaluation scores when tested on the remaining dataset and labels from either the same or the other annotator. It is important to note that the classifier performance for the support dimension is considerably worse ($\sim$ 60% compared to $\sim$ 83%) when the model trained on CTE+CB data is tested on LGBT data. However, the comparison between the drop in F1 scores is not fully compatible as the data in table 4 are the mean evaluation scores across three splits on one-third of the full data for a given topic (e.g., LGBT), whereas the scores in table 5 are based on the full dataset for that topic.

In the last analysis, we evaluated the transferability of a model trained on two of the datasets with combined annotator labels on the remaining dataset with combined annotator labels. Instead of taking the majority vote of the labels and discarding instances with conflicting labels, we created a dataset such that $(X, y) = \bigcup_i (X_i, y_i)$, here X is the feature matrix, and y is the label vector, y_i is the labels by annotator i. Table 6 shows that the model yields

[2]Data: https://doi.org/10.13012/B2IDB-2603648_V1
[3]https://www.nltk.org/

D	L	A	Word (mean TF-IDF)
CTE	E	1	screw (0.96), chilling (0.87), rtnew (0.77), prevalent (0.74), fund (0.70)
		2	screw (0.96), chilling (0.87), prevalent (0.74), fund (0.70), wow (0.69)
	P	1	explained (0.82), reminds (0.80), coverage (0.76), possible (0.75), jermaine (0.71)
		2	explained (0.82), coverage (0.76), difference (0.76), possible (0.75), jermaine (0.71)
	S	1	chilling (0.87), coverage (0.76), tragedy (0.69), terrible (0.67), reveal (0.67)
		2	coverage (0.76), tragedy (0.69), johnathan (0.68), terrible (0.67), mild (0.64)
	NS	1	isn't (0.53), CTE (0.51), pool (0.47), brandon (0.47), weeden (0.47)
		2	isn't (0.53), CTE (0.51), faced (0.50), sea (0.50), pool (0.47)
CB	E	1	): (0.98), convicted (0.94), truce (0.91), boot (0.90), bro (0.84)
		2	): (0.98), convicted (0.94), truce (0.91), boot (0.90), bro (0.84)
	P	1	ali (0.96), pledge (0.94), watching (0.86), actually (0.83), white_people (0.77)
		2	ali (0.96), pledge (0.94), actually (0.83), favorite (0.78), watching (0.78)
	S	1	ali (0.96), pledge (0.94), convicted (0.94), youre (0.90), bro (0.84)
		2	): (0.98), ali (0.96), pledge (0.94), convicted (0.94), bro (0.84)
	NS	1	gay (0.71), best (0.64), go (0.62), caleb (0.60), raver (0.60)
		2	actually (0.83), gay (0.71), best (0.64), live (0.62), caleb (0.60)
LGBT	E	1	opinion (0.95), intended (0.70), sexless (0.68), you's (0.68), uncomfortable (0.67)
		2	opinion (0.95), maybe (0.92), intended (0.70), sexless (0.68), you's (0.68)
	P	1	legalized (0.97), heart (0.84), outside (0.75), drink (0.73), 10_thing (0.71)
		2	dont (0.89), outside (0.75), drink (0.73), 10_thing (0.71), new_campaign (0.71)
	S	1	legalized (0.97), outside (0.75), drink (0.73), 10_thing (0.71), biblical (0.70)
		2	legalized (0.97), outside (0.75), drink (0.73), 10_thing (0.71), biblical (0.70)
	NS	1	larry (0.69), glb (0.69), passion (0.65), kill (0.62), ship (0.62)
		2	actually (0.70), larry (0.69), glb (0.69), kill (0.62), ship (0.62)

Table 3: Salient n-grams identified in annotations of each annotator (A), for each label (L), across datasets (D).

Data	Annotator Model	max 1	max 2	min 1	min 2	mean 1	mean 2	std. 1	std. 2
CTE	Enthusiasm	0.872	0.858	0.517	0.523	0.713	0.697	0.138	0.131
	Support	0.877	0.879	0.800	0.821	0.823	0.835	0.033	0.019
CB	Enthusiasm	0.806	0.796	0.625	0.538	0.740	0.719	0.053	0.070
	Support	0.929	0.910	0.875	0.861	0.899	0.881	0.018	0.017
LGBT	Enthusiasm	0.866	0.815	0.515	0.548	0.667	0.654	0.114	0.085
	Support	0.854	0.839	0.822	0.809	0.839	0.831	0.010	0.009

Table 4: Cross-validation micro-F1 scores for training the model with three-fold cross-validation using an individual dataset with labels from a single annotator.

Test Data	Annotator Model	Test → Train ↓	1	2
CTE	Enthusiasm	1	0.729	0.715
		2	0.772	0.743
	Support	1	0.800	0.826
		2	0.800	0.826
CB	Enthusiasm	1	0.761	0.758
		2	0.738	0.742
	Support	1	0.873	0.843
		2	0.883	0.864
LGBT	Enthusiasm	1	0.729	0.694
		2	0.634	0.604
	Support	1	0.604	0.602
		2	0.596	0.599

Table 5: Evaluation using micro-F1 scores for testing on one dataset from a single annotator and training on other data from the other annotator.

Test Data	Enthusiasm	Support
CTE	0.749	0.801
CB	0.763	0.855
LGBT	0.731	0.559

Table 6: Evaluation using micro-F1 scores for testing on one dataset and training on the others, combining annotations from all annotators.

evaluation scores comparable to those presented in table 5. The evaluation score for the support dimension of the LGBT dataset is again lower compared to the last experiment (table 5). The lower F1 score signifies least inter-annotator agreement as well as Cohen's κ for the LGBT data, as shown in table 2. Our analysis analysis also supports the hypothesis that the models trained using all annotator labels are often more accurate that those trained on single annotator label.

Finally, we created a combined model which was trained on all data from all annotators. This model was tuned using three-fold cross-validation. The model evaluation scores are summarized in table 7. The combined model achieves high mean F1 scores, suggesting the appropriateness of the extracted features that were used for classification. We also investigated the top features for each label identified by these combined models as shown in table 8. These features capture the presence of URL and account mentions, while some of the top features are also related to the respective datasets. Table 8 also shows that top features for enthusiasm contain emotive words like *agree*, *lol*, and *great*, as well as mentions of accounts. On the other hand, top features for the passive class include the presence of URLs, and mention of news outlets like *Reuters*. Similarly, top features for supportive include URL mentions, and explicit mentions of n-grams containing the word *support*, while top features for non-supportive include words like *hate*, *angry*, and *kill*.

The presence of dataset-specific features, i.e., specific words from a dataset, among top features can be smoothed out by the use of word embeddings, which can allow the model to learn more general features. We did not experiment with word embedding based models, owing to the small size of our training data and our need to use cross-validation to identify the variability in model evaluation scores.

Model	max	min	mean	std
Enthusiasm	0.943	0.544	0.798	0.174
Support	0.972	0.845	0.902	0.058

Table 7: Cross-validation micro-F1 scores on training the model with three-fold cross-validation using combined annotations from all annotators across all datasets.

5 IDENTIFYING ENTHUSIASTICALLY SUPPORTIVE ACCOUNTS AND HASHTAGS

In this section, we describe the construction of two types of networks for identifying the most enthusiastic and supportive accounts and hashtags in a corpus of tweets related to a given topic.

5.1 Network construction

In order to identify the top accounts per label and label combination, we created a network of account mentions as follows. If $account_1$ mentions $account_2$ in tweet t, we create a directed edge between $account_1$ and $account_2$. Additionally, we use the probability of t being predicted as either enthusiastic (E), passive (P), supportive (S), or non-supportive (NS) as edge attributes. Also, a weight $w = 1$ attribute is added to each edge, which represents edge frequency. Next, if the same edge occurs multiple times in a corpus, we sum the scores for each of E, P, S, NS, and w. Finally, for each directed edge between n_1 and n_2, we sum the above mentioned scores for each the outgoing edges to get a score for n_1.

A similar network based on account hashtag mentions was also constructed. Here, the edge is between $account$ and $hashtag$ instead of $account_1$ and $account_2$. This network allows for the identification of top hashtags as well as accounts in a corpus of tweets about a topic along the dimensions of enthusiasm and support.

5.2 Identification of top nodes in the network

We use the weighted personalized PageRank (PPR) algorithm [2] for identifying the top nodes of each type. The personalization weights are computed as $exp(\sum_j score_j^1 - \sum_j score_j^0)$. Here, $score_j^1$ and $score_j^0$ represent node scores for enthusiastic (or supportive) and passive (or non-supportive), respectively. The exp is used to ensure that the weights are positive, which is a requirement of the PPR algorithm.

Using the above algorithm, we identified the top accounts in the account mention graph as well as the top account/hashtags in the account-hashtag graphs. The top enthusiastic and supportive nodes are shown in table 9 (mention network) and table 10 (hashtag network). Individual account names were replaced with USR to protect privacy. The tables highlight that the PPR algorithm identifies accounts and hashtags that are different (almost no overlap across all topics) than the ones identified by the baseline page rank algorithm (All). For example, for CTE, one of the top nodes along the enthusiasm dimension in the mention as well as the hashtag network is *@Sports_Brain*. *@Sports_Brain* is the Twitter handle of a company that provides concussion management programs. Similarly, the NFL account is among the top supportive accounts for CTE. For CB, the top hashtag for enthusiasm as well as support is *#cyberbullying*, while the top account in the mention network is *USR2* who is a Youtuber. Finally, for LGBT, the top enthusiastic and supportive accounts are *@free_equal*, which is a United Nations initiative for LGBT equality, and *@USR_FilmExpert*. Further validation would be needed to assess if the accounts and hashtags identified by PPR have a higher relevance with respect to enthusiasm and support related to the selected topics than the accounts identified with the baseline algorithm.

6 CONCLUSIONS

In this work, we have evaluated the EPSNS classification schema for labeling tweets, and subsequent applications for identifying accounts and hashtags expressing enthusiasm and support towards topics of public interest. More specifically, we evaluated the robustness of the annotation of three datasets on different topics through a series of experiments. Our findings demonstrate the robustness of the annotations and the generalization capability of the models trained on these data. Furthermore, we utilized the tweet level classification scores in the personalized PageRank algorithm for identifying top accounts and hashtags that express enthusiasm and support towards the three considered topics.

Our approach is limited by the small dataset and the usage of simple linear models. Furthermore, a direct comparison with a stance classification, which could further establish the utility of the EPSNS classification task for account labeling, is not provided. However, since the goal of this work was to introduce the core idea of identifying accounts based on the measurement of enthusiasm and support expressed in tweets, the comparison with other tasks can be pursued in future studies. Our work can help in merging text classification and network analysis for account labeling on social media platforms.

ACKNOWLEDGMENTS

We thank Sneha Agarwal, Johna Picco, Kirstin Phelps, and Jinlong Guo for their help with developing the annotated data and discussing the classification schema with us.

REFERENCES

[1] Eytan Bakshy, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. 2011. Everyone's an influencer: quantifying influence on twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining - WSDM '11*. ACM Press, New York, New York, USA, 65. https://doi.org/10.1145/1935826.1935845

Model	Label	Feature scores
Enthusiasm	**E**	@account (-44.63), ! (-11.62), rt (-10.02), great (-6.70), read (-6.49), lol (-6.13), thechronicleher (-6.05), war (-5.96), agree (-5.85), ? (-5.79)
	P	head (5.00), scan (5.00), testing_company (5.47), love (5.68), reuters (5.89), supreme_court (5.92), actually (6.53), legalization (7.88), nhl_concussion_case (8.65), URL.COM (10.44)
Support	**NS**	kill (-5.66), go (-4.90), as (-4.85), hate (-4.21), mlk (-4.10), cyber_bully_white_people (-4.00), surveillance (-3.91), angry (-3.72), ? (-3.62), annie (-3.58)
	S	suicide (2.38), URL.COM (2.43), life (2.47), support_gay_right (2.54), absolutely (2.56), 100 (2.65), did (2.72), asshole (2.75), cyber_bully (3.12), sad (3.46)

Table 8: Top features for each class in the combined model trained using all datasets and annotator labels.

	CTE		CB		LGBT	
	Account	**PR**	**Account**	**PR**	**Account**	**PR**
E/P	USR1	0.191	USR2	0.050	free_equal	0.033
	Sports_Brain	0.191	USR4	0.050	UN_Women	0.030
	USR3	0.041	USR5	0.043	USR_FilmExpert	0.030
S/NS	USR6	0.186	USR2	0.062	free_equal	0.044
	USR12	0.068	USR4	0.062	HRC	0.033
	NFL	0.066	USR5	0.054	USR_FilmExpert	0.028
All	USR7	0.021	USR8	0.009	HRC	0.024
	NFL	0.015	USR9	0.008	Tedofficialpage	0.010
	frontlinepbs	0.009	USR10	0.008	USR11	0.010

Table 9: Top 3 nodes in the mention network based on different PageRank algorithms (PR=PageRank score). In the *All* row, ranking and scores are based on overall PageRank. Accounts of individuals were replaced with USR to protect privacy.

	CTE		CB		LGBT	
	Account	**PR**	**Account**	**PR**	**Account**	**PR**
E/P	Sports_Brain	0.264	USR5	0.479	USR1	0.234
	USR2	0.264	#cyberbullying	0.116	#lgbt	0.105
	#cte	0.137	#parenting	0.102	USR3	0.032
S/NS	#nfl	0.062	WestYorksPolice	0.357	USR4	0.427
	#cte	0.058	#cyberbullying	0.122	#lgbt	0.101
	Sports_Brain	0.051	#bullying	0.094	#gay	0.087
All	#nfl	0.048	#cyberbullying	0.048	#lgbt	0.063
	#cte	0.023	#cdnpoli	0.018	#gay	0.008
	#concussion	0.015	#cyberbully	0.015	#questionnier	0.006

Table 10: Top 3 nodes in the hashtag network based on different PageRank algorithms (PR=PageRank score). In the *All* row, ranking and scores are based on overall PageRank. Accounts of individuals were replaced with USR to protect privacy.

[2] Sergey Brin, Lawrence Page, Sergey Brin, and Lawrence Page. 1998. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems* 30, 1-7 (apr 1998), 107–117. https://doi.org/10.1016/S0169-7552(98)00110-X

[3] Jacob Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement* 20, 1 (apr 1960), 37–46. https://doi.org/10.1177/001316446002000104

[4] Svetlana Kiritchenko, Xiaodan Zhu, and Saif M. Mohammad. 2014. Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research* 50 (aug 2014), 723–762. https://doi.org/10.1613/jair.4272

[5] Efthymios Kouloumpis, Theresa Wilson, and Johanna Moore. 2011. Twitter Sentiment Analysis: The Good the Bad and the OMG!. In *Proceedings of the Fifth International AAAI Conference on Web and Social Media*.

[6] Bing Liu and Lei Zhang. 2012. A Survey of Opinion Mining and Sentiment Analysis. In *Mining Text Data*. Springer US, 415–463. https://doi.org/10.1007/978-1-4614-3223-4_13

[7] Shubhanshu Mishra, Sneha Agarwal, Jinlong Guo, Kirstin Phelps, Johna Picco, and Jana Diesner. 2014. Enthusiasm and Support: Alternative Sentiment Classification for Social Movements on Social Media. In *Proceedings of the 2014 ACM Conference on Web Science (WebSci '14)*. ACM, New York, NY, USA, 261–262. https://doi.org/10.1145/2615569.2615667

[8] Shubhanshu Mishra, Sneha Agarwal, Jinlong Guo, Kirstin Phelps, Johna Picco, and Jana Diesner. 2019. Tweet IDs annotated for enthusiasm and support towards social causes: CTE, cyberbullying, and LGBT. https://doi.org/10.13012/B2IDB-2603648_V1

[9] Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. SemEval-2018 Task 1: Affect in Tweets. In *Proceedings of The 12th International Workshop on Semantic Evaluation*. Association for Computational Linguistics, New Orleans, Louisiana, 1–17. https://doi.org/10.18653/v1/S18-1001

[10] Saif M. Mohammad, Parinaz Sobhani, and Svetlana Kiritchenko. 2017. Stance and Sentiment in Tweets. *ACM Transactions on Internet Technology* 17, 3 (jun 2017), 1–23. https://doi.org/10.1145/3003433

[11] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1998. The PageRank Citation Ranking: Bringing Order to the Web. , 17 pages. https://doi.org/10.1.1.31.1768 arXiv:1111.4503v1

[12] Bo Pang and Lillian Lee. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval* 2, 1âĂŞ2 (2008), 1–135. https://doi.org/10.1561/1500000011 arXiv:cs/0112017

[13] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing - EMNLP '02*, Vol. 10. Association for Computational Linguistics, Morristown, NJ, USA, 79–86. https://doi.org/10.3115/1118693.1118704 arXiv:arXiv:cs/0205070v1

[14] Robert Gilmore Pontius and Marco Millones. 2011. Death to Kappa: birth of quantity disagreement and allocation disagreement for accuracy assessment. *International Journal of Remote Sensing* 32, 15 (aug 2011), 4407–4429. https://doi.org/10.1080/01431161.2011.552923

[15] Daniel Preoţiuc-Pietro, Jordan Carpenter, Salvatore Giorgi, and Lyle Ungar. 2016. Studying the Dark Triad of Personality using Twitter Behavior. In *Proceedings of the 25th ACM Conference on Information and Knowledge Management (CIKM)*.

[16] Daniel Preoţiuc-Pietro, Ye Liu, Daniel J Hopkins, and Lyle Ungar. 2017. Beyond Binary Labels: Political Ideology Prediction of Twitter Users. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*.

[17] Parinaz Sobhani, Saif Mohammad, and Svetlana Kiritchenko. 2016. Detecting Stance in Tweets And Analyzing its Interaction with Sentiment. In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*. Association for Computational Linguistics, Stroudsburg, PA, USA, 159–169. https://doi.org/10.18653/v1/S16-2021

[18] T Wilson, J Wiebe, and P Hoffman. 2005. Recognizing contextual polarity in phrase level sentiment analysis. In *Acl*, Vol. 7. Association for Computational Linguistics, Morristown, NJ, USA, 12–21. https://doi.org/10.3115/1220575.1220619

[19] W. Xing and A. Ghorbani. 2004. Weighted PageRank algorithm. In *Proceedings. Second Annual Conference on Communication Networks and Services Research, 2004.* IEEE, 305–314. https://doi.org/10.1109/DNSR.2004.1344743