Extracting Stances on Pandemic Measures from Social Media Data

Erik Tjong Kim Sang
Social Sciences & Humanities
Netherlands eScience Center
Amsterdam, The Netherlands
e.tjongkimsang@esciencecenter.nl

Shihan Wang, Marijn Schraagen, Mehdi Dastani
Intelligent Software Systems,
University of Utrecht
Utrecht, The Netherlands
{s.wang2, m.p.schraagen, m.m.dastani}@uu.nl

***Abstract*—Support for national measures against the COVID-19 pandemic can be measured by collecting responses to questionnaires. In this paper we explore a less costly and less time-consuming method: by analyzing social media data. We compare stances regarding anti-pandemic measures extracted from tweets with the questionnaire results provided by the Dutch national health institute RIVM. We find similarities and differences and discuss the results.**

***Index Terms*—COVID-19 pandemic, social media analysis, opinion mining**

I. Introduction

The COVID-19 pandemic has forced governments to impose a variety of measures fighting the spread of the disease. For policy makers, it is important to have good insight into the impact that these measures have on the general public. In The Netherlands, the national health institute RIVM collects such data, by regularly asking tens of thousands of people to fill in questionnaires, answering questions about the impact of the pandemic and the anti-pandemic measures on their lives [1].

Collecting these useful data by questionnaires is a costly and time-consuming process. Therefore we want to explore if data derived from social media texts could complement the questionnaires. Each day, hundreds of thousands of Dutch-speaking users turn to Twitter express their thoughts. We will analyze their tweets, extract the ones regarding the COVID-19 pandemic and derive opinions from these. Then we will compare the opinions on six different topics over time with the opinions measured from questionnaires from the Dutch national health institute.

Several other authors have also studied pandemic social media messages. Müller et al. [2] applied a BERT model to several COVID-19 text classification tasks. Kumar et al. [3] studied stances of different groups involved in social media discussions. López et al. [4] tracked early responses to COVID-19 policies on Twitter for 15 languages.

After this introduction, we will describe our research methods and data in section 2. Section 3 contains an overview of the results of this study. In section 4 we discuss the results and conclude.

The study described in this paper was funded by the Netherlands eScience Center, under dossier number COVID10, as project number 27020S04. Our data (tweet ids) are available on twiqs.nl. Our software is available in github.com/puregome

II. Methods and data

Our data consists of Dutch tweets as collected by the service twiqs.nl [5]. The data sets consist of about 20 million tweets written in Dutch. We focus on the tweets from the start of the pandemic in The Netherlands (March 2020) until the date of this paper (May 2021). Five steps are performed for extracting opinions from the tweets, similar to Tjong Kim Sang et al. [6]:

1) extract relevant tweets for six topics with keyword search
2) manually annotate about 1,000 tweets for each topic
3) train a machine learning model for each topic, based on the annotations
4) label all tweets found with keyword search with the relevant model
5) perform a statistical analysis of the machine-labeled tweets

The six anti-pandemic measures we study are: *social distancing*, *testing*, *working from home*, *vaccination*, *wearing face masks* and *curfew*. For each of these topics, a keyword query was constructed manually, which was used to make the first selection of the tweets per topic. About 1,000 non-duplicate tweets per topic were selected randomly and labelled manually. Three labels could be assigned to each tweet: Supports, for tweets in agreement with the measure, Rejects, for tweets rejecting the measure and Irrelevant, for off-topic tweets and tweets for which no stance could be determined.

The text classifier FastText [7] was used for learning a model for labelling unseen tweets based on the annotations. In order to increase the coverage of the models, training was started from a language model developed from 230 million Dutch tweets from the year 2020. The default parameter settings of FastText were used, among others model is skipgram, learning rate is 0.05 and number of epochs is 5, except for word vector length, for which we used the value 300. The learning method was evaluated with 10-fold cross-validation, with the F_1 rate of the relevant labels (Supports and Rejects).

The six models were applied to all relevant Dutch tweets of the period March 2020 - May 2021. The comparison data of the Dutch national health institute RIVM is only available for twelve separate evaluation weeks in those 15 months. For each of these weeks, we computed the support for each measure by

TABLE I
DATA AND RESULTS

Topic	**Annotation analysis**				**Machine learning results**				
	Tweets	***Time frame***	***Relevant***	***Support***	***Tweets***	***Relevant***	***F***	***delta***	***r***
social distancing	1,186	Feb 2020 - Jul 2020	75%	72%	242,351	78%	0.685	9±5%	0.7
testing	1,181	Mar 2020 - Dec 2020	47%	62%	436,434	28%	0.496	9±4%	0.0
working from home	1,003	Mar 2020 - May 2021	34%	87%	34,610	26%	0.492	9±4%	-0.4
vaccination	1,340	Jan 2020 - May 2021	35%	29%	871,358	27%	0.354	46±8%	0.7
wearing face masks	1,011	Mar 2020 - Mar 2021	45%	47%	314,748	36%	0.472	54±3%	0.9
curfew	750	Mar 2020 - May 2021	27%	16%	63,309	24%	0.413	75±2%	1.0

dividing the number of found Supports tweets by the number of relevant tweets (Supports plus Rejects). In order to remove the bias effect of the machine learner, a constant was added to the numbers to make sure that the average of the predicted numbers for weeks in the manual annotation period was equal to the average support of the manually annotated tweets.

Two methods were used for evaluating the predicted support scores. First there was the delta score: the average of differences of all the available pairs of scores (from tweets and questionnaires) per evaluation week. The second comparison was the Pearson correlation coefficient (r), which compares the shapes of the graphs produced by a pair of models.

III. RESULTS

The results can be found in Table I and Figure 1. The topics are sorted by increasing delta score: the average difference between the tweet support scores and the scores derived from the questionnaires.

We assume that annotating 1,000 tweets is sufficient for this classification task. We tested this assumption by evaluating a larger number of annotated tweets for a single topic, *social distancing*. For 5,977 tweets we obtained a slightly higher F_1 rate (0.723 in comparison with 0.685) but the delta score (9) remained the same, as well as the shape of the graph.

IV. DISCUSSION

For three of the topics, opinions derived from tweets proved to be a good approximation of the opinions derived from the questionnaires: *working from home*, *testing* and *social distancing*. For the other three topics, this was not the case, as can be observed from the high delta scores in Table I for *vaccination*, *wearing face masks* and *curfew*.

The graphs for *working from home* and *testing* show one way in which the social media data could complement the questionnaires. Questions on these topics were only added to the questionnaires in the Summer of 2020. The social media data also provide insights into the support for the related measures for the period before the inclusion of the questions.

The most important cause for the differences measured for *vaccination*, *wearing face masks* and *curfew* is that the discussions on these topics on Twitter are quite politicized. They include strong opponents which repeat positions known from the English / American discussions, focusing on health risks (*vaccination*) and dangers for civil liberties (*wearing face masks* and *curfew*), often in connection with expressing support for one of the Dutch far right political parties (FVD).

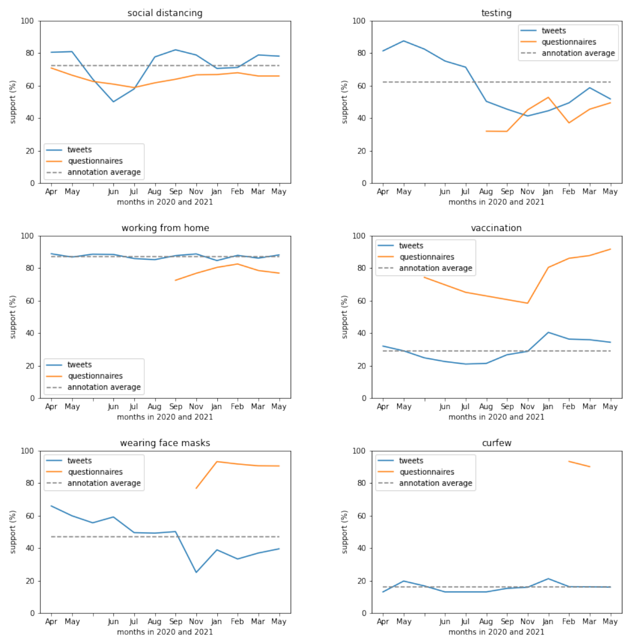

Fig. 1. Support for the six measures, based on tweets and questionnaires

REFERENCES

[1] RIVM, "Does the general public manage to follow the behavioral measures?" 2020, Visited on 15 June 2020. (in Dutch). [Online]. Available: https://www.rivm.nl/gedragsonderzoek/maatregelen-welbevinden/naleven-gedragsregels

[2] M. Müller, M. Salathé, and P. E. Kummervold, "COVID-Twitter-BERT: A Natural Language Processing Model to Analyse COVID-19 Content on Twitter," *CoRR*, 2020. [Online]. Available: https://arxiv.org/abs/2005.07503

[3] S. Kumar, R. A. V. Cox, M. Babcock, and K. M. Carley, "A Weakly Supervised Approach for Classifying Stance in Twitter Replies," *CoRR*, 2021. [Online]. Available: https://arxiv.org/abs/2103.07098

[4] C. E. López, M. Vasu, and C. Gallemore, "Understanding the perception of COVID-19 policies by mining a multilanguage twitter dataset," *CoRR*, 2020. [Online]. Available: https://arxiv.org/abs/2003.10359

[5] E. Tjong Kim Sang and A. van den Bosch, "Dealing with big data: The case of Twitter," *Computational Linguistics in the Netherlands Journal*, vol. 3, pp. 121–134, 2013.

[6] E. Tjong Kim Sang, M. Schraagen, S. Wang, and M. Dastani, "Transfer Learning for Stance Analysis in COVID-19 Tweets," in *CLIN31: Computational Linguistics in The Netherlands*, 2021.

[7] A. Joulin, E. Grave, P. Bojanowski, and T. Mikolov, "Bag of Tricks for Efficient Text Classification," in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Association for Computational Linguistics, April 2017, pp. 427–431.