Coronavirus on Social Media: Analyzing Misinformation in Twitter Conversations

Karishma Sharma, Sungyong Seo, Chuizheng Meng, Sirisha Rambhatla, Aastha Dua, Yan Liu
Department of Computer Science
University of Southern California
Los Angeles, USA.
{krsharma,sungyons,chuizhem,sirishar,aastha,yanliu.cs}@usc.edu

March 30, 2020

ABSTRACT

The ongoing Coronavirus Disease (COVID-19) pandemic highlights the interconnected-ness of our present-day globalized world. With social distancing policies in place, virtual communication has become an important source of (mis)information. As increasing number of people rely on social media platforms for news, identifying misinformation has emerged as a critical task in these unprecedented times. In addition to being malicious, the spread of such information poses a serious public health risk. To this end, we design a dashboard to track misinformation on popular social media news sharing platform – Twitter. Our dashboard allows visibility into the social media discussions around Coronavirus and the quality of information shared on the platform as the situation evolves. We collect streaming data using the Twitter API from March 1, 2020 to date and provide analysis of topic clusters and social sentiments related to important emerging policies such as "#socialdistancing" and "#workfromhome". We track emerging hashtags over time, and provide location and time sensitive analysis of sentiments. In addition, we study the challenging problem of misinformation on social media, and provide a detection method to identify false, misleading and clickbait contents from Twitter information cascades. The dashboard maintains an evolving list of detected misinformation cascades with the corresponding detection scores, accessible online at https://ksharmar.github.io/index.html.

1 Introduction

As social media becomes the primary source of information for people around the world, it has become increasingly critical and challenging to curb misinformation. According to Pew Research report, social media outpaced print news paper in 2018 (Shearer, 2018 (accessed March 20, 2020, Mitchell, 2018), as the share of Americans who get their news online continues to increase (Geiger, 2019 (accessed March 20, 2020).

The misinformation surrounding COVID-19 pandemic is especially damaging, since any mis-steps can poses a serious public health risk by leading to exponential spread of the disease and accidental death due to self-medication (Vigdor, 2020 (accessed March 24, 2020). The risk of misinformation surrounding the pandemic has motivated the World Health Organization (WHO) to launch a "Mythbuster" page WHO (2020 (accessed March 20, 2020), however, these counter measures face challenges with the fast-paced evolution and spread of news on social media. As a result, it is extremely important to identify and potentially curb the spread of misinformation as close as possible to its point of origin. To this end, we present a dashboard to provide insights about the nature of information that is currently shared through social media regarding the Coronavirus pandemic. The dashboard provides an analysis of topics, sentiments, and trends, assessed from Twitter posts; along with posts identified as spreading false, misleading and clickbait information on Coronavirus. Here, we focus our analysis on Twitter since it has the highest number of news focused users (Hughes and Wojcik, 2019 (accessed March 20, 2020) and provides access to public Tweets data.

Figure 1: Coronavirus Misinformation Dashboard: Panel shows the landing page of the developed dashboard to analyze sentiments, trends, and misinformation on Twitter.

2 Dataset Collection

We collect social media posts on Twitter using the streaming API service starting from March 1, 2020 to date. We use keywords related to Coronavirus to filter the Twitter stream and obtain relevant tweets about the pandemic. The dataset from March 1, 2020 to March 17, 2020 contains 8.1M tweets from 182 countries. The subset of English tweets equals 4.7M. We obtain geolocation information at the country-level based on Dredze et al. (2013). The data collection is ongoing and will be used to update the analysis on the dashboard. The the dashboard is available at `https://ksharmar.github.io/index.html`.

3 Sentiment Analysis

We analyze the evolving country-wise sentiments surrounding the Coronavirus pandemic. The public perceptions constitute an important factor for gauging the reactions to policy decisions and preparedness efforts. In addition, they also reflect the nature of news coverage and potential misinformation. We extract sentiments from social media posts at the country-level and over time, to study the evolving public perceptions towards the pandemic. Using lexical sentiment extraction based on (Hutto and Gilbert, 2014), we obtain the valence (positive or negative) along with its intensity for each tweet based on its textual information. The sentiment is aggregated over tweets to estimate the overall sentiment distribution. The distribution of sentiments was found to vary over time and country.

In addition, we analyze the public perception of emerging policies such as social distancing and remote work. These disease mitigation strategies also provide unprecedented glimpse into the effect of remote work and isolation on mental health. Although the option to work remotely is limited to the white collar workforce, nevertheless absence of child and dependent-care has emerged as an important challenge. Furthermore, this forced remote work will impact workdays of white collar workers beyond the pandemic. In order to understand public sentiment and opinion about different social issues, we extract hashtag information from the collected tweets, and filter based on keywords "#workfromhome, #wfm, #workfromhome, #workingfromhome, #wfhlife' and "#socialdistance, #socialdistancing". The filtered tweets are analyzed to obtain positive and negative sentiments and ranked and visualized based on valence and intensity.

4 Topic and Trend Analysis

We analyze Twitter conversations to identify topics and trends in the Twitter data on Coronavirus. We use topic modeling based on character embeddings extracted from social media posts. We identified 20 different topics corresponding to Coronavirus from the English tweets in the data. We found that the prominent topics of discussions during early March were centered around global outbreaks (Wuhan, Italy, Iran), travel restrictions, prevention measures such as hand washing and masks, hoarding, symptoms and infections, immunization, event cancellations, testing kits and centers, government response and emergency funding. We visualize the most relevant (representative) tweets in each topic cluster based on relevance to the word distribution in the estimated topics. The representative tweets and word distribution of each of the 20 topics is provided on the dashboard. The label to each cluster of tweets was assigned by manual inspection of the word distribution and representative tweets of the cluster.

The emerging trends on Twitter highlight changes in perception or importance of topics as the pandemic situation changes. We extract hashtags from the tweet text for all tweets in the dataset for mid-march. The hashtags with emerging popularity are estimated based on fitted linear regression curves on the usage counts of hashtags over the period. The Top-30 emerging hashtags from March 9 - 17 are visualized on the dashboard and regularly updated. On the dashboard, we also provide the distribution of tweets based on the collected geolocation information over time to visualize the proportion of tweets with available geolocation in each country and its trend over time.

5 Misinformation Analysis

Misinformation forms an important aspect of our online world. Due to social distancing measures, the reliance on information available online has become critical. To address this challenge, we identify misinformation and provide the list of identified posts propagating such information on the dashboard.

The task of distinguishing legitimate vs. false, misleading and clickbait content is challenging from both, a human and machine perspective. In fact, this increases the importance of eliminating such information from social media platforms, because the general public is easily manipulated into believing false information, which in this case can be detrimental to public health and have dire consequences Vigdor (2020 (accessed March 24, 2020). In order to identify misinformation, we first extract information cascades from the collected dataset, i.e. retweet trees starting from a source post. To determine the veracity of each cascade, we build a detection model based on the tweet text in the cascades Sharma et al. (2019). We use externally compiled fact-checking sources to label a subset of the cascades as either false, misleading, or clickbait vs. legitimate, following the procedure in Qian et al. (2018); and trained a classifier based on the cascade tweet texts using a neural network with character-level embeddings for classification Joulin et al. (2016).

6 Future Work

There are several critical directions of future work to address this large-scale "infodemic" surrounding COVID-19. The proportion of Twitter users in the United states is higher than in other countries like China with alternate social media platforms. Since the situation is at a global, social media analysis for other platforms and languages beyond English is critical towards curbing misinformation. The second important factor is that the annotation for misinformation is a challenging task and requires expert verification. Therefore, research in unsupervised or distant supervision are important towards alternate forms of labeling, to improve classification and handle the imbalance in the distribution of misinformation vs. legitimate information. We also plan to include social context information towards improving misinformation detection. In addition, we plan to analyze sentiments about other emerging topics.

References

Mark Dredze, Michael J Paul, Shane Bergsma, and Hieu Tran. Carmen: A twitter geolocation system with applications to public health. In *Workshops at the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.

A. W. Geiger. *Key findings about the online news landscape in America*, 2019 (accessed March 20, 2020). URL https://www.pewresearch.org/fact-tank/2019/09/11/key-findings-about-the-online-news-landscape-in-america/.

Adam Hughes and Stefan Wojcik. *10 facts about Americans and Twitter*, 2019 (accessed March 20, 2020). URL https://www.pewresearch.org/fact-tank/2019/08/02/10-facts-about-americans-and-twitter/.

Clayton J Hutto and Eric Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth international AAAI conference on weblogs and social media*, 2014.

Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Hérve Jégou, and Tomas Mikolov. Fasttext.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*, 2016.

Amy Mitchell. *Americans Still Prefer Watching to Reading the News - and Mostly Still Through Television*, 2018. URL https://www.journalism.org/2018/12/03/americans-still-prefer-watching-to-reading-the-news-and-mostly-still-through-television/.

Feng Qian, Chengyue Gong, Karishma Sharma, and Yan Liu. Neural user response generator: Fake news detection with collective user intelligence. In *IJCAI*, volume 3834, page 3840, 2018.

Karishma Sharma, Feng Qian, He Jiang, Natali Ruchansky, Ming Zhang, and Yan Liu. Combating fake news: A survey on identification and mitigation techniques. *ACM Transcations on Intelligent Systems and TEchnology*, 2019.

Lisa Shearer. *Social media outpaces print newspapers in the U.S. as a news source*, 2018 (accessed March 20, 2020). URL https://www.pewresearch.org/fact-tank/2018/12/10/social-media-outpaces-print-newspapers-in-the-u-s-as-a-news-source/.

N. Vigdor. *Man Fatally Poisons Himself While Self-Medicating for Coronavirus, Doctor Says*, 2020 (accessed March 24, 2020). URL https://www.nytimes.com/2020/03/24/us/chloroquine-poisoning-coronavirus.html.

WHO. *Coronavirus disease (COVID-19) advice for the public: Myth busters*, 2020 (accessed March 20, 2020). URL https://www.who.int/emergencies/diseases/novel-coronavirus-2019/advice-for-public/myth-busters.