An Analysis of Users Engagement on Twitter During the COVID-19 Pandemic: Topical Trends and Sentiments

Sultan Alshamrani¹, Ahmed Abusnaina¹, Mohammed Abuhamad², Anho Lee³, DaeHun Nyang⁴, and David Mohaisen¹

¹University of Central Florida, Orlando, FL 32816, USA
²Loyola University Chicago, Chicago, IL 60660, USA
³Korea University, Seoul, South Korea
⁴Ewha Womans University, Incheon, South Korea

Abstract. The outbreak of COVID-19 pandemic raised health and economic concerns. With social distancing along with other measures that are enforced in an attempt to limit the spread of the virus, our life has dramatically changed. During this period, the web and social media platforms have become the main medium for communication, expression, and entertainment. Such platforms are a rich source of information, enabling researchers to better understand how the pandemic affected the users' everyday life, including interaction with and perception of different topics. In this study, we focus on understanding the shift in the behavior of Twitter users, a major social media platform used by millions daily to share thoughts and discussions. In particular, we collected 26 million tweets for a period of seven months, three months before the pandemic outbreak, and four months after. Using topic modeling and state-of-the-art deep learning techniques, the trending topics within the tweets on monthly-bases, including their sentiment and user's perception, were analyzed. This study highlights the change of the public behavior and concerns during the pandemic. Users expressed their concerns on health services, with an increase of 59.24% in engagement, and economical effects of the pandemic (34.43% increase). Topics such as online shopping have had a remarkable increase in popularity, perhaps due to the social distancing, while crime and sports topics witnessed a decrease. Overall, various topics related to COVID-19 have witnessed an improved sentiment, alluding to users adoption to the pandemic and associated topics of the public discourse.

Keywords: Coronavirus; COVID-19; Sentiment Analysis; NLP; Topic Modeling.

1 Introduction

The coronavirus disease 2019 (COVID-19) is the largest pandemic in the information age. Caused by SARS-CoV-2, a highly transmissible respiratory virus, COVID-19 is the biggest public health concern of this century, declared as a global pandemic by the World Health Organization on March 11, 2020 [19]. As of August 2020, there are close to 21 million confirmed COVID-19 cases, with close to 6.4 million active cases, 750 thousand deaths, and 14 million recovered cases [20]. The outbreak of COVID-19 has changed people's life and behavior, including their interaction and communication. Social distancing and other measures are taken in an attempt to limit the spread of

2 Authors Suppressed Due to Excessive Length

COVID-19, making the Web, and social media, in particular, the main medium for communication, expression, and entertainment.

In such a period, data science and mining play a central role in understanding the effect of the pandemic on users, and their behavior and perception. Several studies focused on the users' behavior on social media, including Twitter as a major platform [1, 2, 8, 16]. Understanding the trends, and people's perception, using sentiment analysis allows a better understanding of users' behavior, and their reaction toward a particular topic. In this work, we study the effect of COVID-19 on the behavior of users on Twitter, and how the trends and topics have shifted, including the user sentiment and perceptions. In particular, we collected 26 million tweets from four English-speaking countries for the period of October 2019 to April 2020, three months before the COVID-19 outbreak, and four months after. We investigate the change of the users' behaviors through the duration, in both topics discussed, and their sentiments and thoughts toward each topic. Our findings highlight increased concerns about the economic effects of the pandemic, the quality of the provided health services. In addition, several topics gained popularity during the pandemic, such as online shopping and social mediarelated tweets, with an increase of more than 30% in popularity.

Contributions. This work studies the shift in the Twitter user's behavior coinciding with COVID-19 pandemic. In particular, we make the following contributions:

- We collected a large dataset of 26 million tweets from four English-speaking countries, namely, the United States, Canada, England, and Australia, and provided an in-depth analysis of the collected tweets, including their sentiments and the topics they discuss.
- We used state-of-the-art deep learning and natural language processing techniques to extract and track topics discussed in the social media platform along with the trends of user behavior towards these topics. For topic modeling, we used Latent Dirichlet Allocation, and for per-topic sentiment analysis, we used BERT to track behavioral changes within topics during the studied period. Revealing 23 distinct topics, and the increase of popularity in certain topics during the pandemic, such as health services, economy, and online shopping. Highlighting the public concerns and general trend during the lock-down.

Organization. This remaining of the paper is organized as follows. In section 2, we review the related work in the field of users' perceptions and behavior analysis during the COVID-19 pandemic. Section 3 provides a description of the data collection approach and measurements. In section 4, we perform topic modeling to uncover a variety of topics and their trends during the pandemic from our data. Section 5 provides deep learning-based analysis on users' perceptions and emotions towards different topics raised on the social media platform. Finally, we conclude our work in section 6.

2 Related Work

The huge impact of COVID-19, potential and actual, has led many researchers to study and investigate users' perceptions and behavior coping with and conducting activities during the pandemic. This interest of studying users' perception is motivated by the need for providing well-informed measures to address the spread of misinformation about the pandemic. Several studies have examined the spread of COVID-19 misinformation on Twitter, as it is considered one of the most utilized social media platforms.

Ordun *et al.* [13] conducted an exploratory analysis using different topics, terms, and features on Twitter, to investigate the speed at which the information about COVID-19 spreads in comparison to Chinese social media. The study also examined the network behavior during the COVID-19 pandemic using various techniques, such as machine learning, topic modeling, uniform manifold approximation and projection, and statistical analysis. The findings of this study showed that the median retweeting time of COVID-19-related posts is roughly 50 minutes faster than the re-postings on Chinese social media about the Avian influenza A H7 (H7N9) outbreak in March 2013. Another study focusing on the spread of COVID-19 misinformation on Twitter was presented by Kouzy *et al.* [9]. The authors studied the amount of COVID-19 misinformation on Twitter by collecting and analyzing tweets from 14 different hashtags and words related to COVID-19. The study employed statistical analysis, by comparing terms and hashtags, to identify certain tweets and account characteristics. Using a dataset of 673 tweets, the authors reported 153 tweets to have some degree of misinformation, and 107 to be posted from unverifiable sources.

Schild *et al.* [14] analyzed data from Twitter and 4Chan for a period of five months to investigate the presence of Sinophobia, an issue that was raised at the beginning of the pandemic outbreak, across these two platforms. The study showed a strong correlation between the spread of Sinophobia-related content and COVID-19, and such a correlation was less obvious on mainstream platforms such as Twitter than other platforms such as 4Chan. In addition to analyzing users' behavior during the pandemic and the spread of misinformation regarding COVID-19, recent studies focused on providing techniques to help researchers and scholars mitigating the risk of spreading the virus. For example, Latif *et al.* [11] surveyed several papers, repositories, and datasets related to the pandemic with the intent of making them available to the public to help to track and mitigate the spread of the virus. Their work contributes to highlighting challenges and strategies, as well as creating a live repository for the latest data in this space.

In this study, we provide insights into how the pandemic has impacted people's behavior towards topics discussed on Twitter. Moreover, we analyze trends and shifts of emotions and feelings observed when discussing these topics before and after the pandemic. We utilize state-of-the-art techniques to detect, model, and track different topics from tweets collected in time-frame covering periods before and after the COVID-19 outbreak. After observing the major topics discussed during the data collection period, we adopted a deep learning-based sentiment analysis to study the people's perceptions of these topics before and after the pandemic.

3 Data Collection and Measurements

3.1 Dataset

In this study, we aim to understand the change in the users' behavior and trends before and after COVID-19 outbreak towards certain topics. As such, we collected 26 million tweets from four English-speaking countries, namely, the United States, England, Canada, and Australia. In particular, we collected tweets from 14 major cities with the



Fig. 1: The distribution of collected tweets over countries. 67% of the tweets are collected from the United States.

Fig. 2: The distribution of collected tweets over different cities. Eight cities (57%) are within the United States.



Fig. 3: The number of collected tweets per month within the studied duration. Note that the number of tweets is evenly distributed over the months, with an average of 3.7 million tweets per month.

highest number of daily tweets from a seven-month period, starting from the beginning of October 2019 until the end of April 2020. The data collection period consists of three months before the outbreak and four months after. We used the GetOldTweets3 API [12] to collect tweets from the social media platform, Twitter. In addition to our collected tweets, we obtained Sentiment140 dataset [7], a collection of 1.6 million tweets, including 800,000 tweets labeled as positive, and 800,000 tweets labeled as negative. The latter dataset is used as a ground-truth benchmark for the sentiment analysis task.

Data Collection: Statistics and Measurements. The collected tweets are distributed over a span of seven months across four countries, as shown in Figure 1. Since the United States has by far the largest number of users on Twitter (62.55 million users), this is reflected in the dataset, where 67% of the tweets are from the United States [6]. Since both Australia and Canada have fewer users, we limited the number of cities from these countries to one city from each country, with the highest number of tweets originated from Toronto (Canada) and Sydney (Australia) as shown in Figure 2. The 26 million tweets are distributed over seven months, with an average of 3.5 million tweets per month, with a peak of 4 million tweets in March 2020, which represents the duration of the peak of the first wave of the outbreak [17].

3.2 Data Preprocessing

In order to understand users' behavior on twitter before and after COVID-19 outbreak, two tasks are required: topic modeling and sentiment analysis. The topic modeling task aims to detect topics that were discussed during the collection period and while keeping track of the evolution of trends related to these topics. The sentiment analysis task aims to explore people's perceptions on topics and whether the behaviors towards these topics are impacted by the pandemic (in a correlation-based analysis).

Each task requires different data preprocessing steps. For topic modeling, we cleaned the collected tweets by removing any special characters, hashtags, URLs, as well as non-English characters. Then, we removed the stopwords and short phrases with less than three characters. This preprocessing step is important considering the targeted task, *i.e.*, topic modeling, since topics are informed by keywords that occur across a large number of documents. Cleaning the collected tweets from unrelated and/or infrequent words can positively impact the outcome of the topic modeling task.

To handle data for sentiment analysis, we kept the original tweets and applied Word-Piece tokenization [15] on the collected tweets. Skipping the other preprocessing steps, such as the eliminations of stopwords and non-English characters representing emojis, improves the performance without impacting the accuracy of the analysis as we employ Bidirectional Encoder Representations from Transformers (BERT), which utilizes WordPiece tokenization anyway. WordPiece tokenization is a word segmentation algorithm that forms a new subword of a token using a pre-trained likelihood probability model, *e.g.*, the word "working" is divided into two subwords "work" and "ing". Such an approach is beneficial when handling the out-of-vocabulary problem. We describe our use of WordPiece in the following section.

3.3 Data Representation

Each of the two tasks we conducted (topic modeling, and sentiment) requires certain numerical representation. For that, we used bag-of-words representation for topic modeling, and WordPiece tokenizer for sentiment analysis.

Bag-of-Words Representations. For the topic modeling task, we represent tweets with bag-of-words. Representing the tweets as clusters of different topics, based on the similarity score of their context, goal, or interest, requires considering frequent keywords from the content of tweets.

In order to utilize different state-of-the-art topic modeling techniques, tweets are first represented using a bag-of-words method during the data representation phase, in which a tweet is transformed into a vector of values that represent the presence or the weight of all unique words in the corpus in relation to the tweet. In the bag-of-words scheme, a tweet is commonly represented with a hot-encoding vector that highlights the existing words of the tweet against all terms in the corpus. The collection of all unique terms in a corpus is often referred to as a dictionary. Considering the fact that most of the employed datasets are large in size, the dictionary can be very large, producing sparse and high-dimensional bag-of-words representations. To optimize the bag-of-words representation of tweets, we adopted several steps, including: (1) removing frequent and rare words, and (2) feature selection. Common words that appear in more than, for example, 50% of all tweets are maybe general terms with less discriminative power than



Fig. 4: The general flow of the sentiment analysis pipeline. Sentiment140 dataset is used to fine-tune BERT-based model for sentiment classification task.

other less frequent words. Rare words, that appear less than 1,000 times in the entire corpus, for example, are also eliminated. We note that we created a pre-defined list of words related to COVID-19 in the final feature set, to have enough features to produce topics related to COVID-19 outbreak. This list of terms include terms such as: *coronavirus*, *virus*, *corona*, *covid19*, *covid-19*, and *covid*.

To reduce the dimensionality of the bag-of-words representation and after the preprocessing step, we selected the most frequent 10,000 words to be the features for representing the tweets. In a preliminary experiment, choosing the 10,000 words produced the best trade-off between performance and accuracy for the topic modeling task.

WordPiece-based Representations. Due to the difference in the two tasks, we utilized different methods in the data representation. For the sentiment analysis task, we use a WordPiece-based representation method to represent tweets to the BERT model [5]. The WordPiece technique creates a vocabulary of a fixed number of words, sub-words and characters, and solves the out-of-vocabulary problem by splitting unrecognized words into sub-words, if no sub-word match in the pre-defined dictionary, it is then split further into characters, and then mapped to the corresponding embedding. This technique has been proven efficient as opposed to other embedding mechanisms that map all of the out-of-vocabulary words to one token such as 'UNK' [15,21].

Tweets are represented as a matrix with rows fixed as the number of words in the representation and columns as the embedding of each word. For our implementation, we used 70 words as the length of tweets, and 768 as the length of the word embedding vector, *i.e.*, each tweet is represented as 70×768 matrix representation. We use 70 words to represent a tweet because the majority of tweets (99.8%) in our dataset were within this length, and the 768 is the pre-defined size of the word embeddings of BERT.

4 LDA-based Topic Modeling and Tracking

This study aims to explore the user's perceptions and behavior in response to a variety of topics and subjects during the data collection period. Moreover, we aim to investigate the trends of these topics and the impact of the pandemic on users' reactions to these topics (in a correlation analysis), requiring powerful tools and techniques such as topic modeling. Topic modeling is an unsupervised machine learning technique, that is typically used to extract a set of topics in a group of documents. Topic modeling processes a set of documents and detects the repeated patterns of word and phrase across documents, to cluster the documents based on their similarities [3, 10]. This study utilizes the MALLET's Latent Dirichlet Allocation (LDA)-based topic modeling technique, a state-of-the-art topic modeling approach that maps each document in the corpus to a set

6

7



Fig. 5: The distribution of tweets over LDA generated topics. Most tweets are associated with "Thoughts/Emotions" and "Family/Social Life" related topics.

of topics. Each document, *i.e.*, each tweet, is assigned to a topic with a probability score, allowing the tweet to be recognized in different topics. However, we assign the tweet to the topic with the highest score. Using the topic modeling, topics are represented with a cluster of similar and related words [4]. This enables the detection and tracking of topics through the data collection period. Accurate detection of topics allows real-time analysis and observation of trends, *e.g.*, users' reactions and behavior towards topics.

4.1 LDA-based Topic Modeling: Configuration

After the collection and processing of data, tweets are represented with vector representations of a bag-of-words. Extracting the bag-of-words representation is described in section 3. Receiving input vectors of 10,000 bag-of-word representation, the LDA model assigns topics for each tweet. The abstract pipeline of the process is presented in Figure 7. Establishing the topic model requires a training phase in which a number of topics are investigated in terms of the coherence score of topics. The coherence score is a score calculated for each topic, by measuring the semantic similarity between words that have the highest score within the given topic, the word score is calculated based on the frequency of the word within the topic, and its inverse frequency with other topics.



Fig. 6: The over-time distribution of positive and negative tweets per generated topic.



Fig. 7: The general flow of the topic modeling task. Each tweet is represented as bagof-words generated vector for LDA topic modeling training.

This method provides distinguishable measurements between topics that are semantically similar. The higher the coherence value the better the quality of the clustering, indicating better topic modeling and assignment. We examined the effect of changing the number of the extracted topics on the modeling task. We explored extracting 15 to 50 topics with an increase of 5 topics, each iteration, while observing the coherence score achieved in each iteration. The LDA-model achieves the best performance when the number of extracted topics is 40 with a coherence score of 0.55.

4.2 LDA-based Topic Modeling: Outcome

Using the best-performing LDA model, we manually inspected the topics through the frequent keywords generated for each topic, and assigned names and descriptions to those topics. Since some topics share similar keywords, we assigned the same topic name for multiple clusters (*i.e.*, LDA-extracted topics). This manual inspection of topics produced 23 unique topics.

Topic-Tweet Distribution. Figure 5 shows the distribution of topics across the collected tweets in the seven months. The topics related to "Family/Social Life" and "Thoughts/Emotions" represent 35.38% of the collected tweets. Moreover, during this period, 0.5 million tweets were COVID-19/Pandemic-related. In general, except for the aforementioned topics, the tweets are distributed in the range of 0.4-2 million tweets per topic. **Topic Temporal Tracking.** Using the best-performing model, we tracked topics through time to observe trends and topic evolution. Figure 6 shows the temporal distribution of tweets through different months of the collection period.

An obvious increase in the tweets related to "COVID-19/Pandemic" is observed in March as the World Health Organization declared COVID-19 as a global pandemic. This increase is shown in Figure 6d with a 472.15% rise in the volume of the tweets from December to March. A similar trend is observed for topics related to "Health Services" with a 76.7% increase from December to March, as shown in Figure 6l.

This pandemic has also affected the trends of topics related to "Economy", as governments and organizations started enforcing lock-downs and social distancing. Topics related to the economy have enjoyed a 35.64% increase in tweets between December and March, as shown in Figure 6h. The "Online Shopping/Tickets" topic has shown a similar increase, of around 34%, from October to March, as shown in Figure 6p.

The negative impact of the pandemic on the volume of tweets is observed on topics related to entertainments, *e.g.*, "Game/Sport" with a decrease of 57.95% from October to April as most sports events were suspended around the world. We also observed a

decrease of tweets related to "Crime/Violence" with more than 21% drop in volume from October to April as shown in Figure 6e, although unclear if that is due to a decrease in crime rates or the engagement of users with crime-related contents. The full illustrations of the trends of the topics are shown in Figure 6.

5 Topic-derived Sentiment Analysis

The second part is to investigate the trends and shift of perceptions of certain topics discussed on Twitter during the data collection period. This part is done by examining the sentiments observed on different topics before and after the pandemic. To establish a baseline sentiment analysis model, we used a ground-truth benchmark dataset of 1.6 million annotated tweets for sentiment analysis task [7]. At the beginning of our research, we conducted several preliminary experiments to define the best model architecture to perform the task, including deep learning-based techniques such as Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), Long Short Term Memory (LSTM), and BERT. Using the ground-truth dataset, the achieved F-1 scores were 80.53%, 82.03%, 81.65%, and 87.06% for DNN, CNN, LSTM, and BERT, respectively. Table 1 shows the achieved results of different models performing the sentiment analysis in terms of true positive rate, true negative rate, precision, and the F-1 score. Note that BERT outperforms all others, achieving a precision accuracy of 87.35%. Therefore, we selected the BERT model to perform the sentiment analysis in our study. An illustration of the pipeline of the workflow to conduct the sentiment analysis using BERT is shown in Figure 4, including the preprocessing and data representation stages.

5.1 BERT-based Sentiment Analysis

In essence, Bidirectional Encoder Representations from Transformers (BERT) [5] is a language model that benefits from the attention mechanism used in the transformer architecture [18]. This attention mechanism has two six-layers of encoders that have the ability to learn contextual relation between words in a given text, as well as six layers of decoders that generate the needed output for a given task. As opposed to traditional NLP models that read textual data sequentially from right-to-left or leftto-right, the transformer encoder is considered bidirectional since it reads the entire given text at once, allowing the model to capture the context of each word based on its surroundings. Since BERT is a language model, it uses only the encoder part of the transformer. By adding a new layer to the core model. BERT fits a wide variety of NLP language tasks, such as classification, question answering, as well as named entity recognition. In our implementation of BERT, we used the same structure and model configuration of the original work. For more details, we refer the reader to the original research paper in [5]. In this study, tweets are separated per sentence, with two special tokens indicating the start of the tweet and end of each sentence. Then, each tweet is fed into the trained BERT model, providing the embedding of each word in the tweet considering its surroundings. The output of the BERT model is a one-neuron output layer with a sigmoid activation function for binary classification signaling the polarity of the tweet, *i.e.*, either positive or negative sentiment.

Model	True Positive Rate	True Negative Rate	Precision	F-1 score
DNN	80.15%	80.97%	80.91%	80.53%
CNN	81.51%	82.67%	82.55%	82.03%
LSTM	81.60%	81.62%	81.71%	81.65%
BERT	86.77%	87.45%	87.35%	87.06%

Table 1: The evaluation of deep learning models on sentiment analysis task. BERTbased model outperform its counterparts, therefore, used as the baseline in our analysis.

Topics-derived Sentiment. Figure 6 shows the observed positive and negative tweets for different topics. Generally, the percentage of reported sentiments are approximately similar throughout the collection period for all topics. However, the tweets related to "COVID-19/Pandmeic" have shown a significant decrease in negative tweets from 41.48% in January to 34.46% in March. Such negativity drop is also observed on tweets related to "Weather/Travel" with 56.37% of negative tweets in October to 46.12% in March. For the tweets related to "Health Services", the results show an increase of the positivity in the tweets, as the percentage of positive tweets increased from 42.65% in October, to 47.49% in March. Similar trends in the sentiment analysis of the topic related to "Online Shopping/Tickets" with an increase in the positive tweets from 30.64% in October to 37.08% in March. While it is impossible to accurately pinpoint the root cause of the "positivity", it seems as though that the overall pandemic and associated measures are accepted by some as a reality, reducing the negative reaction.

6 Conclusion

In this study, we aim to better understand the effect of the pandemic on users' interaction on social media, including their public perception. Using a large-scale dataset of 26 million English tweets from four countries, and 14 major cities, we conducted topic modeling across 23 topics, and performed a temporal and semantic analysis on monthlybases. Our analysis highlights the increasing concern in the public discourse on the provided health services, shifting the public discussing from topics such as sports and politics, into online shopping and economical effects of the pandemic on the society. More interesting, over time, at the aggregate-level, users have become more positive in their expression, as measured by the sentiment in their tweets on various topics.

Acknowledgement. This work was supported by NRF grant 2016K1A1A2912757 (Global Research Lab) and a gift from NVIDIA. S. Alshamrani was supported by a scholarship from the Saudi Arabian Cultural Mission.

References

- Alshamrani, S., Abuhamad, M., Abusnaina, A., Mohaisen, D.: Investigating online toxicity in users interactions with the mainstream media channels on youtube. In: The 5th International Workshop on Mining Actionable Insights from Social Networks. pp. 1–6 (2020)
- Alshamrani, S., Abusnaina, A., Mohaisen, D.: Hiding in plain sight: A measurement and analysis of kids' exposure to malicious urls on youtube. In: Third ACM/IEEE Workshop on Hot Topics on Web of Things. pp. 1–6 (2020)

12 Authors Suppressed Due to Excessive Length

- 3. Blei, D.M., Lafferty, J.D.: Topic models. In: Text mining, pp. 101-124 (2009)
- Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. J. Mach. Learn. Res. 3, 993– 1022 (2003)
- Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Burstein, J., Doran, C., Solorio, T. (eds.) Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers). pp. 4171–4186 (2019)
- J. Clement: Leading countries based on number of twitter users as of july 2020 (2020), https://bit.ly/3gDVbrt
- 7. KazAnova: Sentiment140 dataset with 1.6 million tweets (2020), https: //www.kaggle.com/kazanova/sentiment140
- Kouloumpis, E., Wilson, T., Moore, J.: Twitter sentiment analysis: The good the bad and the omg! In: Fifth International AAAI conference on weblogs and social media (2011)
- Kouzy, R., Abi Jaoude, J., Kraitem, A., El Alam, M.B., Karam, B., Adib, E., Zarka, J., Traboulsi, C., Akl, E.W., Baddour, K.: Coronavirus goes viral: quantifying the covid-19 misinformation epidemic on twitter. Cureus 12(3) (2020)
- Landauer, T.K., Foltz, P.W., Laham, D.: An introduction to latent semantic analysis. Discourse processes 25(2-3), 259–284 (1998)
- Latif, S., Usman, M., Manzoor, S., Iqbal, W., Qadir, J., Tyson, G., Castro, I., Razi, A., Boulos, M.N.K., Weller, A., et al.: Leveraging data science to combat covid-19: A comprehensive review (2020)
- 12. Mottl: GetOldTweets3 Twitter Scrapping API (2020), https://github.com/Mottl/ GetOldTweets3
- Ordun, C., Purushotham, S., Raff, E.: Exploratory analysis of covid-19 tweets using topic modeling, umap, and digraphs. CoRR abs/2005.03082 (2020)
- Schild, L., Ling, C., Blackburn, J., Stringhini, G., Zhang, Y., Zannettou, S.: "go eat a bat, chang!": An early look on the emergence of sinophobic behavior on web communities in the face of COVID-19. CoRR abs/2004.04046 (2020)
- Schuster, M., Nakajima, K.: Japanese and korean voice search. In: 2012 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2012, Kyoto, Japan, March 25-30, 2012. pp. 5149–5152. IEEE (2012)
- Severyn, A., Moschitti, A.: Twitter sentiment analysis with deep convolutional neural networks. In: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 959–962 (2015)
- 17. The Atlantic: The public deserves the most complete data available about covid-19 in the us. no official source is providing it, so we are. (2020), https://covidtracking.com/
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Guyon, I., von Luxburg, U., Bengio, S., Wallach, H.M., Fergus, R., Vishwanathan, S.V.N., Garnett, R. (eds.) NeuRIPS (2017)
- 19. World Health Organization: WHO Director-General's opening remarks at the media briefing on COVID-19 11 March 2020 (2020), https://tinyurl.com/vyvm6ob
- 20. WorldOMeter: COVID-19 Coronavirus Pandemic (2020), https://www.worldometers.info/coronavirus/
- 21. Wu, Y., Schuster, M., Chen, Z., Le, Q.V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., Dean, J.: Google's neural machine translation system: Bridging the gap between human and machine translation. CoRR abs/1609.08144 (2016)