

“A Virus Has No Religion”: Analyzing Islamophobia on Twitter During the COVID-19 Outbreak

Mohit Chandra*
International Institute of Information
Technology, Hyderabad
Hyderabad, India
mohit.chandra@research.iiit.ac.in

Manvith Reddy*
International Institute of Information
Technology, Hyderabad
Hyderabad, India
manvith.reddy@students.iiit.ac.in

Shradha Sehgal*
International Institute of Information
Technology, Hyderabad
Hyderabad, India
shradha.sehgal@students.iiit.ac.in

Saurabh Gupta
Indraprastha Institute of Information
Technology Delhi
New Delhi, India
saurabhg@iiitd.ac.in

Arun Balaji Buduru
Indraprastha Institute of Information
Technology Delhi
New Delhi, India
arunb@iiitd.ac.in

Ponnuram Kumaraguru[†]
International Institute of Information
Technology, Hyderabad
Hyderabad, India
pk.guru@iiit.ac.in

ABSTRACT

The COVID-19 pandemic has disrupted people’s lives driving them to act in fear, anxiety, and anger, leading to worldwide racist events in the physical world and online social networks. Though there are works focusing on Sinophobia during the COVID-19 pandemic, less attention has been given to the recent surge in Islamophobia. A large number of positive cases arising out of the religious *Tablighi Jamaat* gathering has driven people towards forming anti-Muslim communities around hashtags like #coronajihad, #tablighijamaatvirus on Twitter. In addition to the online spaces, the rise in Islamophobia has also resulted in increased hate crimes in the real world. Hence, an investigation is required to create interventions. To the best of our knowledge, we present the first large-scale quantitative study linking Islamophobia with COVID-19.

In this paper, we present *CoronaBias* dataset which focuses on anti-Muslim hate spanning four months, with over 410,990 tweets from 244,229 unique users. We use this dataset to perform longitudinal analysis. We find the relation between the trend on Twitter with the offline events that happened over time, measure the qualitative changes in the context associated with the *Muslim* community, and perform macro and micro topic analysis to find prevalent topics. We also explore the nature of the content, focusing on the toxicity of the URLs shared within the tweets present in the *CoronaBias* dataset. Apart from the content-based analysis, we focus on user analysis, revealing that the portrayal of religion as a symbol of patriotism played a crucial role in deciding how the Muslim community was perceived during the pandemic. Through these experiments, we reveal the existence of anti-Muslim rhetoric around COVID-19 in the Indian sub-continent.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in collaborative and social computing**; **Social media**; • **Information systems** → **Social networks**.

*Authors contributed equally to this research.

[†]Major part of this work was done while Ponnuram Kumaraguru was a faculty at IIIT-Delhi

KEYWORDS

social network analysis; data mining; web mining; social media

1 INTRODUCTION

The global Coronavirus (COVID-19) outbreak has impacted people’s personal, social, and economic lives. The pandemonium caused due to misinformation, and insufficient preparation to handle such outbreaks has resulted in increased levels of fear, anxiety, and outbursts of hateful emotions among the population [1, 30, 47]. Consequently, specific communities are being targeted with acts of microaggression, physical and verbal abuse. Social media is flooded with such hateful posts causing online harassment of particular communities [4, 49]. Chinese communities and Asians at large have been subjected to racism in the physical world as well as in online spaces due to the COVID-19 origin theories [13, 24, 35, 49].

Sinophobia has been one of the prime issues during this pandemic, but the surging level of Islamophobia has not received due attention from the research community. During this pandemic, the Islamic missionary movement *Tablighi Jamaat*’s events that happened in multiple countries of South-East Asia have been in the news due to a large number of COVID-19 positive cases being traced back to the attendees of these congregations [6, 23, 32]. Unfortunately, this has resulted in a discriminatory movement where the people from the *Muslim* community are being blamed for the virus spread and portrayed as “human bombs” and “corona jihadis” [21, 37] leading to trends like #coronajihad, #tablighijamatvirus on Twitter.

Islamophobia has been studied broadly in the context of terrorism in the past. Previous works have shown that the Muslim community is often linked to terror and violence [12]. However, there are other ways in which Islamophobia manifests itself, as has been the case with the ongoing pandemic. This study brings to light the communalism of the COVID-19 pandemic in India. Through multiple quantitative and qualitative experiments, we reveal the existence of anti-Muslim rhetoric around COVID-19 in the Indian sub-continent. In previous studies, it has been shown that real-world crimes can be related to incidents in online spaces [46]. A similar phenomenon has been observed with the hate crimes against the *Muslim* community after the surge COVID-19 cases [34]. In this

paper, we propose the *CoronaBias* dataset (data available here ¹), hoping that it helps take forward the hate speech detection research. Our work has multiple stakeholders, including legislators, social media companies/moderators, and their users.

In this paper, we focus on analyzing the spread of Islamophobia on the social media platform Twitter, especially in the context of the Indian sub-continent. To the best of our knowledge, this is the first quantitative study at the intersection of Islamophobia and COVID-19. Through this paper, we focus on the following research questions:

- **RQ1** : How did the offline events that happened during the course of this study statistically affect online behaviour? How did the context associated with the *Muslim* community change over the period of our study?
- **RQ2** : Which topics were prevalent in the tweets concerned with the *Muslim* community. Were some topics more prevalent than others during a particular window of time?
- **RQ3** : What were the differentiating characteristics of users who were indulged in spreading Islamophobia in contrast to the other users?
- **RQ4** : What was the role played by external sources, especially the news media outlets? What was the nature of the content that was referenced in the tweets through URLs?

Our contributions: In this paper, we present the first of its kind “*CoronaBias*” dataset with over 410,990 tweets from 244,229 unique users, which we will publicly release on the acceptance of this paper. Additionally, we present a long-term longitudinal analysis. Our key contributions are summarised as follows: i) Drawing a correlation between real-world events and corresponding changes on Twitter in a statistically sound manner, with the help of Temporal Analysis using the PELT algorithm. We also find a growing association of the *Muslim* community with the COVID-19 pandemic through the semantic similarity experiment; ii) Extracting the prevalent topics of discourse in broad and focused windows of time through *Macro* and *Micro* topic modelling experiment. Qualitatively, we find a blend of mixed sentiments over the topics; iii) We surface the differentiating characteristics of users sharing Islamophobic content from other users. We specifically analyze user network, user activity, and descriptions of user bios to reveal the differentiating characteristics; iv) Comparing the nature of external content that is referenced on Twitter via URLs. In addition to presenting the content referenced through URLs from YouTube, BBC, and OpIndia (top shared domains), we present a comparative study of the toxicity of the referenced content.

2 RELATED WORK

The unprecedented nature of the ongoing pandemic has attracted much attention from the research community. This section presents some of the past works on research related to COVID-19 and hate speech.

COVID-19 and Online Social Media: The presence of a massive amount of data related to COVID-19 on various social media platforms has attracted researchers to collect it for different studies [10, 35, 40]. Chen et al. [10] proposed a multilingual coronavirus

Twitter dataset with over 123 million tweets. Schild et al. [35] presented one of the first studies related to the Sinophobic behavior during the COVID-19 pandemic. The researchers collected over 222 million tweets and performed experiments revealing the rise in Sinophobia. Works done in the past were related to the characterization of online discussion on different issues. Previous works like Schild et al. [35], Ziems et al. [49] have focused on a quantitative longitudinal analysis of Sinophobia that emerged during this pandemic. On a different front, Kouzy et al. [26] have presented a study to analyze the magnitude of misinformation related to COVID-19 on Twitter. Along similar lines, Ferrara [17] have emphasized the role of bots in spreading conspiracy theories related to the pandemic on Twitter. The topic of privacy & policy in the health sector has been a central focus during this pandemic, primarily due to the importance of contact tracing [11, 20]. Although there have been quite a few works exploring different aspects of the ongoing pandemic, to the best of our knowledge there is no quantitative analysis based work exploring Islamophobia during the COVID-19 pandemic. We fill this gap in this paper.

Research related to hate speech and Islamophobia: Hate speech has been a popular topic in the research community since a long time. While there have been works that have focused on the general hate speech detection [8, 14, 16, 19] there are others that have focused on specific forms of hate speech like racism & sexism [43] and cyberbullying [9]. In another line of work, researchers have used deep learning for hate speech detection [5, 31].

Unlike other forms of hate speech, Islamophobia has not been explored in depth. Soral et al. [36] presented that social media users were exposed to a higher level of islamophobic content than the users who consumed news from traditional forms of media. In another work, Vidgen and Yasseri [41] proposed an automated software tool that distinguishes between non-Islamophobic, weak Islamophobic, and strong Islamophobic content.

3 ETHICAL CONCERNS

We have used two sources of data in this work: i) COVID-HATE from Twitter (Chen et al. [10]), and ii) News articles scraped from BBC & OpIndia, and video titles and descriptions from YouTube. The data from Twitter was collected from a publicly available source and we hydrated the tweet IDs using the official Twitter API. Information about YouTube videos was collected with the official YouTube Data API ². For the news articles, we provided a user agent string that made our intentions clear and provided a way for the administrators to contact us with questions or concerns. We requested data at a reasonable rate and strived never to be confused for a DDoS attack. We saved only the data we needed from the page. The use of data in this research warranted a justification of the ethical implications and the decisions made during the work. The methods followed during the study are influenced by the works around the internet research ethics and, more specifically, Twitter research ethics [18, 42, 50]. We handled data and made decisions in data collection in an ethical manner with no intent to harm any individual’s privacy in any way [27, 29, 39].

¹<https://github.com/mohit3011/Analyzing-Islamophobia-on-Twitter-During-theCOVID-19-Outbreak>

²<https://developers.google.com/youtube/v3>

4 CORONABIAS: AN ANTI-MUSLIM HATE DATASET

For our study, we used the data collected from Twitter, Chen et al. [10] have presented a large corpus using specific keywords related to COVID-19 and have published the tweet IDs on an online repository.³ We filtered these only to consider tweets written in Roman script (English) from February 1, 2020, till May 31, 2020. Since our work focuses on the *Muslim* community, we further filtered out tweets based on a curated set of keywords. We used positive, negative, and neutral terms so as to avoid bias in our dataset. Some examples of keywords of each type include *tablighiheroes*, *muslim-saviours* for positive keywords, *muslimvirus*, *coronajihad* as negative keywords and *islam*, *muslim* for neutral keywords. (Link to the entire keyword list can be found here ⁴).

The filtered dataset (referred as “*CoronaBias* dataset” from now on) contains 410,990 tweets with a month-wise distribution of number of tweets being: February: 105,974, March: 70,682, April: 161,780, May: 72,554. 244,229 unique users in the dataset have tweeted at least once during the course of the study, out of which 2,107 are verified accounts. The average word length for the tweets in the dataset is 34.

To answer **RQ4**, we collected data from 764 videos on YouTube, 231 unique articles on OpIndia, and 82 unique articles on BBC, which were referenced as URLs in the tweets present in our dataset (ethical concerns related to the data collection have been discussed in Section 3).

4.1 Dataset Annotation

In addition to creating the *CoronaBias* dataset, we assigned binary labels to the tweets based on the stance of Islamophobia. In this labelling task, we assigned each tweet one of the two labels - ‘Hateful’ or ‘Non-Hateful’. The sheer number of tweets made it impossible to annotate the entire dataset manually; therefore, we approached the problem of semi-automatic annotation of tweets in three ways – 1) Using Linguistic Inquiry and Word Count (LIWC) [38] 2) Using Bidirectional Encoder Representations from Transformers (BERT) based method, 3) SVM based method.

Method	Accuracy	Recall	Precision	F-1
BERT	0.854± 0.088	0.854± 0.087	0.854± 0.086	0.853± 0.088
SVM	0.799± 0.024	0.803± 0.024	0.806± 0.023	0.799± 0.024

Table 1: Results for the task of tweet classification into ‘Hateful’ and ‘Non-Hateful’ category using BERT and SVM (Using 5-fold Cross-Validation). We report the macro-averaged values for each metric.

Out of the three approaches, LIWC based method performed the worst. Alternatively, we used SVM and a BERT based model for this task due to its recent success in various downstream NLP tasks [15, 33, 44, 45]. We used a 2-layer MLP with dropout following

³<https://github.com/echen102/COVID-19-TweetIDs>

⁴<https://github.com/mohit3011/Analyzing-Islamophobia-on-Twitter-During-theCOVID-19-Outbreak/blob/main/keywords-major.txt>

the BERT module to classify the data, keeping the learning rate at 10^{-6} and dropout value at 0.2. We fine-tuned the BERT based model on our dataset for maximum of 20 epochs (code ⁵). We also experimented with SVM where we used a linear kernel along with regularization parameter= 0.1 and maximum iteration= 3500.

To create the dataset used for training our models, we filtered out 2000 tweets from the *CoronaBias* dataset using a curated list of keywords that contained positive, negative, and neutral *Muslim*-related terms, to ensure a balanced training data. We used a modified version of the guidelines provided in [41] for the annotation process. The filtered tweets were annotated on binary labels – 1) Non-Hateful and 2) Hateful according to the stance of tweets concerning the *Muslim* community, by three annotators. The Fleiss’ Kappa score for the annotation was 0.812, which translates to a near-perfect agreement among the annotators.

After the annotation process, the distribution of the number of tweets came out as: Hateful = 1,051, Non-Hateful = 949. We used this labelled dataset to train the BERT based model and SVM classifier. Table 1 presents the results for the classifiers. We chose BERT based approach due to its better performance. After training our BERT-based classifier, we classified all the tweets present in the *CoronaBias* dataset as ‘Hateful’ or ‘Non-Hateful’. We use the classified tweets in the *CoronaBias* dataset for further experiments.

5 TEMPORAL ANALYSIS

To answer the first part of **RQ1** concerning the relation of offline and online events, we carried out the temporal analysis of the annotated tweets present in the *CoronaBias* dataset. In particular, we focused on the correlation between real-world events and the behaviour on Twitter during the course of this study along with the difference in the dynamics between ‘Hateful’ and ‘Non-Hateful’ labelled tweets. We used a methodology similar to the one proposed by Zannettou et al. [48]. We used change-point analysis to rank each change in the time-series graph based on mean and variance for both the tweet categories. For the change-point analysis, we used the PELT algorithm as described in [25].

PELT is an exact algorithm for change-point analysis that determines the points in time at which the mean and variance changes by maximizing the likelihood of the distribution given the data. We ran the PELT algorithm on the tweets labelled as ‘Hateful’; additionally, we used a penalty with the algorithm to keep the count of change-points limited. We ran the algorithm with different values of penalty constant (1 – 10) in decreasing order and kept track of the largest penalty amplitude at which each change-point first appeared. This provided us a ranking of the change-points in order of their significance (lower the rank, higher is the significance of the event). For correlating the change points to the real-world events, we considered a window of 10 days around the date of change-point and collected the information of the concerned events through multiple media outlets and Wikipedia. Table 2 comprises event descriptions along with the change-point significance rank and the date of event occurrence.

Figure 1 presents the frequency of tweets present in the ‘Hateful’ and ‘Non-Hateful’ categories from February till May 2020. As

⁵https://github.com/mohit3011/Analyzing-Islamophobia-on-Twitter-During-theCOVID-19-Outbreak/tree/main/bert_files

Change Point Event Rank	Date (MM-DD-YYYY)	Event Description
1	03-31-2020 & 04-01-2020	On 31st March, reports linked the <i>Tablighi Jamaat</i> congregation event that happened in Delhi to the sudden spread of COVID-19 across the country.
1	04-26-2020	COVID Explosion in Maharashtra and record new cases. 24th April marked the start of Ramadan.
2	04-12-2020	On 14 April, India extended the nationwide lockdown till 3 May. On 18 April, the Health ministry announced that 4,291 cases were directly linked to the <i>Tablighi</i> event.
3	04-06-2020	As of 4 April, about 22,000 people who came in contact with the <i>Tablighi Jamaat</i> missionaries had to be quarantined.
3	03-17-2020	Malaysia reported its first two deaths. By 17 March, the Sri Petaling event had resulted in the biggest increase in COVID-19 cases in Malaysia, with almost two thirds of the 673 confirmed cases in Malaysia linked to this event.
4	02-26-2020	On the 25th of February, the Iranian government first told citizens that the U.S. had "hyped COVID-19 to suppress turnout" during elections, and that it would "punish anyone" spreading rumors about a serious epidemic.
4	05-26-2020	24th May was end of Ramadan.

Table 2: Table linking the ranked change points to the set of offline/real-world events. The ranked change-points are computed using the penalty factor in the PELT algorithm (i.e, higher the penalty, higher is the significance of the change-point and lower is the event rank).

expected, the number of ‘Non-Hateful’ tweets was higher than the number of ‘Hateful’ tweets in general because the prior category contains both *pro-Muslim* and *Muslim-neutral* tweets.

As observed in Table 2 and Figure 1, several change-points coincide with the increased number of ‘Hateful’ tweets on Twitter. Table 2 includes the significant offline events that took place in the real-world around the time of each of the online change points of our dataset. After looking at two of the highest significant events (event rank 1) in Table 2, we observed a local maxima in the number of ‘Hateful’ labelled tweets (in fact the number attained a global maxima on 31st March / 1st April 2020). Interestingly, the rank one event on 31st March/1st April is one of the few occurrences in the entire study where the number of ‘Hateful’ tweets was more than the number of ‘Non-Hateful’ tweets. In contrast to that, another event rank 1 (26-04-2020) in Table 2 before which we observed a drop in the number of ‘Hateful’ tweets marked the start of *Ramadan*.

We observed a clear correlation between the offline events and the corresponding behavioural changes on Twitter through the temporal analysis. Additionally, using the PELT algorithm, we surfaced the changes in the statistical properties (mean & variance) of the trend on Twitter following the real world in a ranked fashion. We observed that offline events with greater significance impacted the online trends more.

6 SEMANTIC SIMILARITY EXPERIMENT

In the second part of **RQ1**, we investigated the change in context associated with the *Muslim* community during the pandemic. For this, we performed the semantic similarity experiment. Through this experiment we also wanted to validate our hypothesis that with the due course of time, the *Muslim* community was blamed for the mishappenings during the pandemic. We modified the methodology presented in [35] and trained Word2Vec Continuous-Bag-of-Words (CBOW) model [28].

We chose two general words, ‘*Muslim*’ and ‘*Virus*’ as these can be used in many contexts (referred to as reference words in the rest of this section). We trained a separate Word2Vec model for each month from February till May 2020; more formally, we defined the time period of the study as $\mathcal{T} = \{Feb, March, April, May\}$. For

each of the Word2Vec model $\mathcal{W}_i, i \in \mathcal{T}$, we computed the most similar words to the two reference words (‘*Muslim*’ and ‘*Virus*’) using the cosine similarity criterion. Table 3 presents the most similar words for each of the reference words across all the months. Observing the trend for the word *Muslim*, we saw that the list of similar words in February was not related to COVID-19, with a few words being racist slurs. However, we observed the appearance of words like ‘China’, ‘Virus’, ‘COVID19’ in the subsequent months, suggesting that the context around the word ‘*Muslim*’ was used in a context similar to the aforementioned words in the dataset. Similar was the case with the reference word ‘*Virus*’ where we could observe the word being related to various disease-related terminologies and symptoms in the initial two months. However, in the month of March and onwards, we observed the words ‘Allah’ and ‘Muslims’, indicating that people started associating COVID-19 with the *Muslim* community.

Additionally, in Table 3 we observed that for the word ‘*Muslim*’, the cosine similarity score of words like ‘China’, ‘Virus’ increased from March to May. Similar was the case with the word ‘*Virus*’ where we observed the cosine similarity score for ‘Muslims’ steadily rise for the same time period, indicating the context around the pandemic inclined more towards Islamophobia due course of time. We also noticed the association of the word ‘*Muslim*’ with less familiar words like ‘Dungan’, ‘Indigenous’, and ‘Reality’. A closer look at tweets containing these words revealed the reason behind this association. The word ‘Dungan’ is a term used in the former Soviet Union territories to refer to a group of Muslim people. The word ‘Reality’ was used in many tweets to describe the reality of Islamophobia on the ground. The word ‘Indigenous’ was often used to portray the association of the *Muslim* community being indigenous to India but still facing religious discrimination.

Through this experiment, we showed that the context around the *Muslim* community changed through the course of the pandemic. As time progressed, we observed growing association between the words like ‘Virus’, ‘COVID19’ with the reference word *Muslim* which is indicative of a rising Islamophobic sentiment.

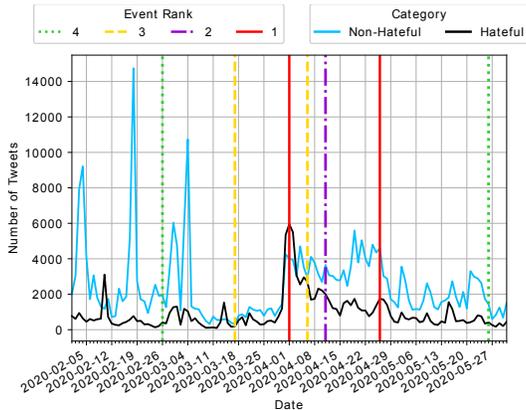


Figure 1: Distribution of ‘Hateful’ and ‘Non-Hateful’ tweets made by the users. The vertical lines represent the event ranks referenced in Table 2.

7 TOPIC MODELLING

To answer RQ2, we delve deep into the content analysis of the collected data through topic modelling. We divided this experiment into two levels of granularity – 1) *Macro Level Topic Modelling*, 2) *Micro Level Topic Modelling*. At the macro level, we performed the task of topic modelling on the *CoronaBias* dataset to reveal the prevalent topics across the period of our study. At the micro-level, we chose two significant events during our study and performed the task of topic modelling over a window of ten days around these events.

We experimented with three different approaches – 1) Latent Dirichlet Allocation (LDA), 2) Top2Vec Angelov [3], 3) Non-Negative Matrix Factorization (NMF), out of which NMF based method provided the best results. Due to the limited space, we only discuss the results obtained using the NMF based technique.

For the NMF based method, we pre-processed the tweets (removal of punctuation, non-ASCII characters, stopwords, and whitespaces, followed by lowercasing). To prevent any anomaly in the method, we removed all the topics that occurred in more than ~ 85% or less than ~ 3% of the tweets. For finding the number of topics for both *Macro level* and *Micro level* topic modelling, which could best represent the data, we used the *Coherence score*. We iterated through the number of topics from 5 to 50 with a step size of 5. We then used TF-IDF based vectorization to create tweet vectors which were used in the NMF based method (code here ⁶). Furthermore, we classified the topics as *Hateful* and *Non-Hateful* on the stance of portrayal of *Muslims* and the Islamic religion related to COVID-19.

⁶https://github.com/mohit3011/Analyzing-Islamophobia-on-Twitter-During-theCOVID-19-Outbreak/tree/main/topic_modelling_files

Muslim				Virus			
February	March	April	May	February	March	April	May
Systematic (0.318)	Islam (0.434)	Islam (0.444)	People (0.457)	Disease (0.447)	Coronavirus (0.446)	Coronavirus (0.491)	Coronavirus (0.540)
Travelers (0.305)	Bully (0.288)	People (0.389)	Islam (0.432)	Disease (0.415)	Disease (0.411)	Covid19 (0.427)	Covid19 (0.481)
Bully (0.294)	China (0.283)	China (0.346)	China (0.402)	Creator (0.361)	Disease (0.380)	Muslims (0.406)	Disease (0.471)
Dungan (0.282)	Moslem (0.281)	Virus (0.329)	Covid19 (0.386)	Wrath (0.350)	Muslims (0.369)	Corona (0.382)	Muslims (0.457)
Reality (0.280)	Virus (0.275)	Covid19 (0.314)	coronavirus (0.361)	Coughing (0.343)	Allah (0.342)	Disease (0.363)	Corona (0.434)
Moslem (0.273)	Seeing (0.233)	US (0.293)	India (0.344)	Diseases (0.332)	China (0.2817)	People (0.349)	Muslim (0.379)
Indigenous (0.268)	People (0.224)	Americas (0.292)	One (0.329)	Source (0.296)	Muslim (0.275)	India (0.334)	People (0.379)

Table 3: Most similar words to *Muslim* and *Virus* in decreasing order of the value of cosine similarity value (in parenthesis). We observed an increase in cosine similarity score between words like *Muslim*, *Virus* and *China* indicating that the context around the pandemic moved towards Islamophobia with the due course of time.

Table 4 presents the prevalent topics along with the tokens and the number of tweets belonging to the particular topic for the *Macro* topic modelling. Among the topics which were categorized as ‘Non-Hateful’, we observed a high number of tweets for topic 1, which talks about the COVID-19 pandemic distracting people from the issue of torture on Uyghur *Muslims*. Topic 2 talks about the financial crisis and market collapse during the pandemic, while Topic 3 represents tweets related to general precautions & measures against COVID-19. Topic 4, interestingly, talks about members of *Tablighi Jamaat* donating plasma for plasma therapy and helping others, tweets belonging to this topic were very prevalent during the month of *Ramadan*.

In contrast, when we looked at the ‘Hateful’ topics, we found a lot of them related to the different aspects of the *Tablighi Jamaat* event and its members. Topic 5 represents anti-Abrahamic religious sentiments that arose due to various conspiracy theories related to the origin of COVID-19. Topic 6 is a discriminatory topic that blames the entire *Muslim* community for the stone-pelting incident on the health workers. Topic 7 & 8 are related to blaming the *Tablighi Jamaat* event for the rise in COVID-19 cases and representing Islamophobic sentiments. An interesting observation comes from the number of tweets belonging to each of the topics - we found the hateful topics to be more representative of the data. This can be confirmed with the number of tweets belonging in each category - while topics in the non-hateful category have less than 10k tweets, almost 2-2.5 times more tweets belong to the last three hateful topics (6,7,8).

For the *Micro* topic modelling we chose two significant events – 1) The release of reports related to the *Tablighi Jamaat* event on 31st March/1st April 2020 ; 2) The occasion of *Ramadan* on 25th April. The choice of these two events came from the fact that we

	S.No.	Topic	Tokens	No. of Tweets
Non-Hateful	1	The Coronavirus pandemic is distracting people from the Uyghur <i>Muslim</i> crisis.	coronavirus distraction, muslims tortured, imprisoned raped, concentration killed, thousands chinese	5,368
	2	The Coronavirus Pandemic leading to a global market collapse.	markets collapse, global demand, corona hits, find difficult, attention needed	1,975
	3	General measures to prevent Covid	washing hands, covering face, five times, shaking hands, covering	4,470
	4	<i>Tablighi Jamaat</i> members donating plasma	donate plasma, tablighijamaat members, members recovered, covid donate, patients	7,555
Hateful	5	Anti-Abrahamic religious sentiments attached with COVID-19.	kill muslims, china kill, hate muslims, hate christians, hate jews	6,089
	6	<i>Muslims</i> attacking Covid-19 health workers	indian muslims, fight, pandemic, doctors, poor, violence, govt	28,747
	7	<i>Tablighi Jamaat</i> religious congregation	tablighi jamaat, maulana, jamaat attendees, police, jamaat members, pakistan, markaz	24,718
	8	Assigning blame to <i>Muslims</i> and Allah for the spread of Covid-19	corona jihad, spreading corona, allah, muslims, corona virus, people	19,561

Table 4: Some prevalent topics in Macro topic modeling experiment for the entire corpus. The rows in green (1-4) and red (5-8) represent the non-hateful and hateful topics respectively.

S.No.	Tablighi Jamaat Event (April 1st)		Occasion of Ramadan (April 25th)	
	Topic	Tokens	Topic	Tokens
1	India's efforts to contain virus foiled by <i>Tablighis</i>	every tablighi, indias efforts, severe setback, single cluster, efforts containing, keep counting	Criticising media for spreading Islamophobia	hey islamophobics, islamophobics medicine, indian media, tablighijamaat, social media
2	Islamic preacher praying to divert Covid-19 to Non- <i>Muslim</i> nations	allah divert, nonmuslim nations, infections nonmuslim, divert covid19, preacher praying	Discriminatory behaviour towards <i>Muslims</i>	muslims denied, pandemic differentiate, failure indian, denied food, racist rhetoric
3	Islamophobia tainting India as doctors refuse to treat <i>Muslim</i> patients	islamophobia taints, hospital refuses, muslim patient, crisis morality, news channels	<i>Tablighi jamaat</i> willing to donate plasma	plasma patients, members recovered, covid donate, cure others, tj members

Table 5: Some prevalent topics in Micro topic modelling experiment for the entire corpus.

wanted to look at the prevalent topics in two contrasting events on the stance towards the *Muslim* community. Table 5 presents the relevant topics for each of the event. For the *Tablighi Jamaat* event, we see that most of the relevant topics are negatively portraying the *Muslim* community and blaming them for the spread of COVID-19. Though there are topics (topic 3) where people are condemning Islamophobia and negligence in treating *Muslim* patients after the *Tablighi Jamaat* incident, a majority of the tweets relate to how the *Muslim* community caused a setback to India's efforts to curb the virus.

On the contrary, topics associated with the occasion of Ramadan include many positive topics which talk about peace, humanity, and the good work that people have been doing during the pandemic, along with condemning the discriminatory behaviour against *Muslims*.

Through this experiment, we analyzed the topics that were prevalent throughout the period of our study. While we observed topics in support of the discrimination against the *Muslim* community, some others included condemnation of this discriminatory behaviour. Through this experiment, we were able to identify the different

aspects of hate associated with *Muslims* after going through a variety of topics. Additionally, the *Micro* topic modelling helped us co-relate the prevalence of topics and events happening in the real world.

8 USER CHARACTERISTICS AND COMMUNITIES

The discourse on social media platforms is a reflection of how the users react to the offline events. Hence, it becomes essential to analyze the differentiating characteristics among the users. For this analysis, we categorized the users present in the *CoronaBias* dataset in three categories based on the percentage of 'Hateful' tweets among the total number of tweets tweeted by the user. To remove any instance of aberration, we only considered users who had tweeted at least five times during our study. Through this filtration process, we finally conducted this experiment on 12,328 unique users. We categorized the users into the following classes: 1) **Class 1:** The percentage of hateful tweets is < 25% of the total number of tweets by the user (6,021 users); 2) **Class 2:** The percentage of hateful tweets is $\geq 25\%$ and < 50% of the total number of tweets by



Figure 2: Word cloud generated from the bio of the users present in Class 1 (0 – 25% ‘Hateful’ tweets). Added the color version to show you (will add the grayscale version in the main draft)



Figure 3: Word cloud generated from the bio of the users present in Class 2 (25 – 50% ‘Hateful’ tweets). Added the color version to show you (will add the grayscale version in the main draft)



Figure 4: Word cloud generated from the bio of the users present in Class 3 (50 – 100% ‘Hateful’ tweets). Added the color version to show you (will add the grayscale version in the main draft)

the user (3, 638 users); 3) **Class 3:** The percentage of hateful tweets is $\geq 50\%$ of the total number of tweets by the user (2, 669 users).

To answer **RQ3** relating to user characteristics and communities, we performed the following set of experiments:

- We created word clouds from the user bios for each category of users. This experiment helped us analyze the ideologies for each set of users who actively engaged with the platform.
- We also created a user network graph to reveal the interactions between the users from different classes. Furthermore, we analyzed the pattern in the online activity of the users and co-related it to the offline events.

8.1 Analysis of User Bio

In this experiment, we wanted to analyze the differences in ideologies among the users and understand how they identified or presented themselves online. Hence, we extracted the user bios for the users in each category, pre-processed them, and generated word clouds for each class.

Figure 2, Figure 3 and Figure 4 presents the word clouds for Class 1, 2 and 3 respectively. As observed, when we move from Class 1 to Class 3, we see the growing significance of terms related to Nationalism and *Hinduism*. A detailed look at the word cloud for Class 1 users revealed the usage of non-hateful/general terms like *Proud Indian*, *Indian Muslim*, *Humanity*, *Muslim*, *Love* etc. This set of users had the least amount of hateful tweets against Muslims ($< 25\%$), and their user bios also indicated that many claimed to be Muslim and support ideologies of humanity and love. Out of the 6, 021, 4.9% of user had the word *Indian* written in their bio, 4.4% of users used the word *love* whereas 3.89% of users used the word *Muslim* in their bio.

Class 2 users’ word cloud has many terms related to Nationalism/patriotism (*India*, *India First*, *Proud Indian*) but at the same time also has religious terms (both *Pro-Muslim* and *Pro-Hinduism*). 7.4% of the 3, 638 users used the word *Indian*, 6.9% of users had the word *Proud* whereas 5.2% of users used the word *Hindu* in their bio. As this is a transition class with medium level of hateful tweets against Muslims, the users’ bios also indicated the same. In contrast to the prior two word clouds, Class 3 users’ word cloud shows the sign

of blatant *Hinduism* which is portrayed as patriotism through the usage of Nationalist terminologies (*Hindu Nationalist*, *Nation First*, *Proud Hindu*). The word cloud also contains terms related to Islamophobia (*Chinese Jihadi*, *Ex Secular*, *Jihadi Products*). Moreover, we also observed a sharp rise in the percentage of users belonging to this class using these terms. 8.6% of the 2, 669 users used *Hindu*, 8.5% used *Proud* and 5.4% used *Nationalist* in their bio.

This analysis revealed that religion was used as a symbol of patriotism by the far-right community present on Twitter. Moreover, the *Muslim* community was portrayed as ‘anti-nationalist’ due to the *Tablighi Jamaat* incident. Overall, the user bio descriptions gave us insights into what ideologies users from each class believed in and we found a correlation between how they presented themselves and the nature of content they posted.

8.2 User Network Graph and Activity Analysis

As the first part of this experiment, we generated the following network among the 12, 328 users. For creating the network graph, we considered each user as a node and added an edge between the nodes if one of the users (node A) followed the other user (node B). We removed isolated nodes i.e. nodes without any outgoing or incoming edges, to obtain 10, 366 nodes. We used the ForceAtlas2 layout algorithm [22] on Gephi for the network spatialization with a scaling factor of 2.0 and gravity of 1.0.

In Figure 5 we found that users who posted a majority of hateful tweets (Class 3 coloured in red: users with $\geq 50\%$ hateful tweets) are clustered closely together. The predominance of red colour is restricted to this one community implying that users that spread hate closely followed each other. This cluster does have some yellow nodes (Class 2: users with $\geq 25\%$ and $< 50\%$ hateful tweets) but very few green nodes (Class 1: users with $< 25\%$ hateful tweets), furthering our observation that this community of users spread more hateful content. We observed two other clusters of users that posted less hateful tweets. One cluster is predominantly green, and the other has majorly green nodes with some yellow nodes interspersed. This shows that there are also communities that posted majorly non-hateful content. To further validate the results obtained from our experiment, we ran Louvain community detection algorithm [?]

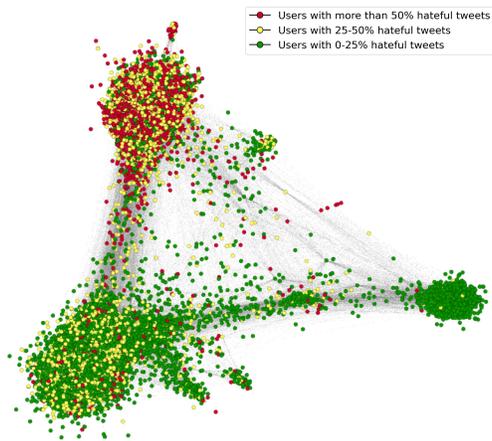


Figure 5: Network Graph for the case where the users are classified in one of three categories based on the percentage of ‘Hateful’ tweets posted. Green nodes represent the users in Class 1 (0 – 25%), Pink nodes represent the users in Class 2 (25 – 50%) and Orange nodes represent the users in Class 3 (50 – 100%). Added the color version to show you (will add the grayscale version in the main draft)

] on our dataset. We obtained a total of 14 separate communities based on the modularity class, out of which three communities constituted 99.7% of the graph. On further investigation we found these three communities to be congruous with the 3 major clusters we observed in Figure 5, where user demarcation was based on the percentage of hateful tweets posted.

Overall, we conclude that the follower-following relations between users are based on the kind of content they post and that the users in a cluster/community have a similar nature of the amount of hateful tweets, generating an echo-chamber effect overall. Some users of class 2 are scattered among the two major clusters, showing the equivocal behaviour of some users.

Apart from the network graph, we also focused on observing the patterns in user activity during our study. We performed a temporal analysis of the number of unique users in each category posting tweets for each day to further reveal the differences in the behaviour. Figure 6 presents the user activity graph for the three-class classification taxonomy based on % of ‘Hateful’ tweets. We observe a relatively high number of Class 1 unique users (users with < 25% hateful tweets) active during the first week of March; we believe this was due to people posting *Pro-Muslim* tweets on the issues of *Uyghur Muslims* and rising COVID-19 cases in Iran. Furthermore, a subset of Class 1 users also formed a separate cluster in the network graph that focused on *Muslims* in contexts other than the *Tablighi Jamaat* event in India. From the first week of March until 30th March, the user activity remained at normalcy.

31st March 2020 was when the reports about *Tablighi Jamaat* congregation event were published. This day marked a sharp rise in the active users belonging to Class 2 and Class 3, suggesting the

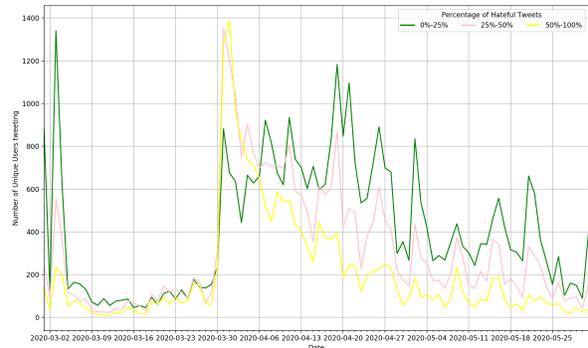


Figure 6: User Activity Graph for the case where the users are classified in one of three categories based on the percentage of ‘Hateful’ tweets posted. Class 1 (0 – 25%) in Green, Class 2 (25 – 50%) in Pink and Class 3 (50 – 100%) in Yellow.

rise in Islamophobic content on Twitter. Interestingly, this incident also marked the start of behaviour where the users started posting more about the *Muslim* community (both in a positive and negative context). We again observed a rise in activity for Class 1 users around the time of *Ramadan* (last week of April). During this time, people started posting tweets with positive messages for the *Muslim* community with hashtags like *#muslimsaviours*, lauding Muslims for their contribution to aiding with plasma therapy and distribution of food packages.

This experiment showed how particular events trigger a specific section of users to engage more with the platform. We also observed a pattern of increased discussion revolving around the *Muslim* community during the ongoing pandemic.

9 EXTERNAL URL EXPERIMENT

Interestingly, the *CoronaBias* dataset contains ~ 68.5% of tweets with one or more URLs referencing external content. Hence, we focused on the analysis of URLs to answer **RQ4** relating to nature of external sources.

In this experiment, we extracted every URL from each of the unique tweets (excluding RTs). We then expanded each of the extracted URLs using the Requests⁷ library and counted the frequency for each domain. Table 6 presents the most popular external URL domains sorted by the number of tweets those links occur in. As expected, all but one domain (YouTube) belong to news media companies, as the users referenced these sources for the latest news headlines related to the disease spread, its prevention/cure, and associated lockdown policies. Among the popular social media

⁷<https://requests.readthedocs.io/en/master/>

S.No.	External URL Domain	Frequency	S.No.	External URL Domain	Frequency
1	opindia.com	9483	6	indiatimes.com	2865
2	aljazeera.com	4032	7	altnews.in	2554
3	bbc.com	3347	8	thewire.in	2250
4	jihadwatch.org	3061	9	dw.com	2080
5	youtube.com	2924	10	washingtonpost.com	2051

Table 6: Top 10 most frequently referenced domains through URLs present in the tweets in the CoronaBias dataset. The frequency denotes the total number of URLs that belong to these domains.

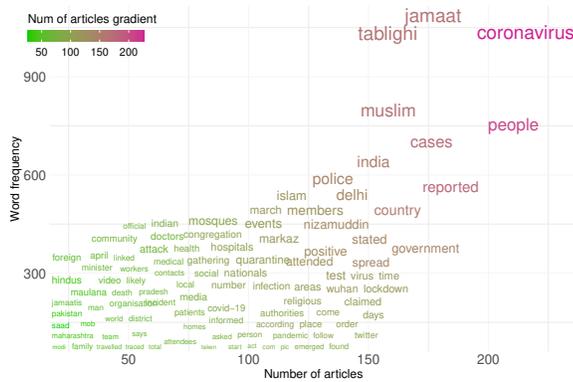


Figure 7: Chatter plot for article titles and content collected from OpIndia.

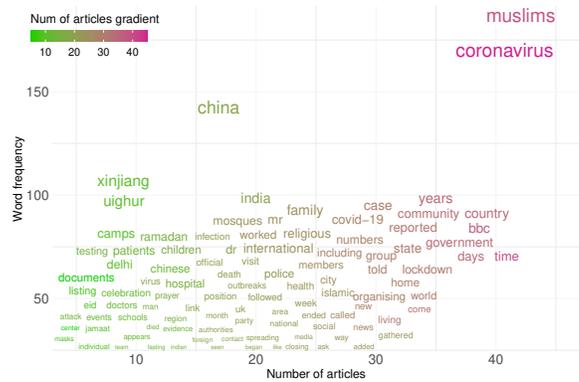


Figure 8: Chatter plot for article titles and content collected from BBC.

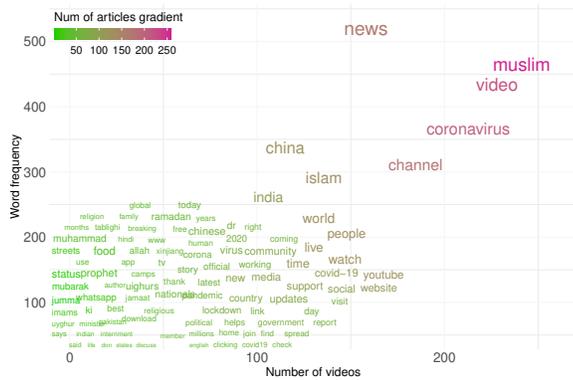


Figure 9: Chatter plot for the video titles and descriptions collected from Youtube.

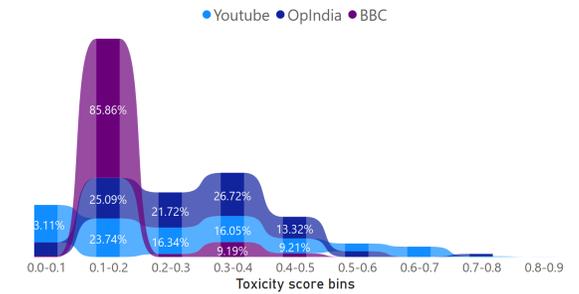


Figure 10: Ribbon plot for percentage of videos from Youtube, articles from OpIndia, and BBC in different Toxicity score bins obtained using Perspective API.

platforms, we found the order as Facebook (579) followed by Instagram (322), Tumblr (155), and LinkedIn (62). Surprisingly, we observed a large number of URLs from a few websites that have been known to disseminate anti-Muslim news ⁸ (OpIndia, Jihad-watch, and Swarajyamag). The frequency for URLs from OpIndia is maximum and double that of Aljazeera, which shows the scale of probable Islamophobic content shared through the URLs on Twitter.

In the second step, we dug deeper into the content referred through the URLs. For this experiment we chose three domains –

1) OpIndia⁹, 2) BBC¹⁰, 3) Youtube¹¹. While we collected article title and article text data from OpIndia and BBC, we focused on title along with description for YouTube videos (ethical concerns related to data collection have been discussed in Section 3). Fig 7, Fig 8 and Fig 9 present the chatter plots obtained for the data collected from OpIndia, BBC and YouTube respectively. As observed, the content referred from OpIndia talks about *Muslims* and the *Tablighi Jamaat* incident in great detail, with words like *Tablighi*, *Jamaat* occurring more than 900 times across ~ 180 articles. The content presented

⁸<https://theprint.in/opinion/india-anti-muslim-fake-news-factories-anti-semitic-playbook/430332/>

⁹<https://www.opindia.com/>

¹⁰<https://www.bbc.com/>

¹¹<https://www.youtube.com/>

by BBC varied broadly while focusing on topics of Uighur *Muslims*, general measures to prevent COVID-19, and lockdown due to the pandemic. The content referred from Youtube lay in between that of OpIndia and BBC.

Additionally, we analyzed the toxicity of external content. For this experiment, we used Perspective API ¹² to get the toxicity score for the data collected from the three sources. Perspective API defines toxicity as "rude, disrespectful, or unreasonable content". Figure 10 presents the toxicity ribbon plot for each of the content from OpIndia, BBC, and Youtube. We divided the toxicity scores into bins 0-0.1, 0.1-0.2 etc, and plotted the percentage of articles from that domain belonging to a given toxicity bin. We observed that ~ 86% of the tweets containing a link from BBC lie in the bin 0.1-0.2 representing minor to no use of harsh words (here harsh words does not mean toxic words). In contrast to this, the number of articles from OpIndia were more evenly distributed in the range of 0.1-0.5. Moreover, the presence of a significant number of tweets in the range of 0.5-0.8 indicated the toxic nature of content present on OpIndia. Furthermore, occurrence of majority of articles from BBC in the 0.1-0.2 bin shows that any form of bias arising from the Perspective API due to religious terms like *Muslims* etc. does not have a confounding effect on our claim of presence of toxic content for articles lying in the bin 0.5-0.8. For further validation of the results, we manually analyzed 50 most frequently occurring articles referred from BBC and OpIndia in our dataset. We found that ~ 66% of the 50 most frequently occurring articles from OpIndia portrayed Islamophobic behaviour, whereas we did not find any Islamophobic article among the 50 most frequently occurring articles from BBC.

The experiments performed in this section raise the concern that while the text/images presented in a tweet may not explicitly be hateful, tweets can still spread toxicity and hate from the URLs referenced in them. The widespread presence of media sources like OpIndia in our dataset, that frequently publish anti-Muslim content, shows that people used external sources to further Islamophobic views. The high percentage of tweets containing news media sources also highlights that people reference ongoing news in crucial times like a pandemic, where information becomes key.

10 CONCLUSION AND DISCUSSION

Through this work, we aim to provide a data-driven overlook of rising Islamophobia around the world. We curate a dataset of tweets related to the *Muslim* community, known as "*CoronaBias*" to study how anti-*Muslim* discrimination starts and spreads during the pandemic caused by the Coronavirus. The dataset consists of 410,990 tweets from 224,229 unique users having an average of 34 words in each tweet. Using this dataset, we perform temporal analysis with the PELT experiment to draw a correlation between the offline events and the corresponding changes on Twitter. We find a growing association of the *Muslim* community with the COVID-19 pandemic through the semantic similarity experiment. We perform *Macro* and *Micro* topic modelling to understand the topics prevalent during the course of our study in detail and during two focused windows. Qualitatively, we find a blend of mixed sentiments over the topics. Apart from the content-based analysis, we also perform a user-based experiment to reveal the differences in characteristics of

various classes of users through their user bios and online activity. Finally, we conduct an external URL experiment to study the nature of the content referred to outside of Twitter.

Islamophobia has been studied broadly in the light of terrorism in the past. Researchers have shown that the Westerners often link Muslims to terror and violence [12], their policies result in an increase of Islamophobia [2]. In addition, [7] discuss that media, government, and community discourses converge to promote Islam as dangerous and deviant. Through this work, we create an understanding of anti-Muslim sentiments that are not directly coherent with terrorism but can harm the community in a dire manner. The spread of negative sentiments towards any community is unacceptable; therefore, we wish to build on this work to explore its mitigation.

REFERENCES

- [1] Daniel Kwasi Ahorsu, Chung-Ying Lin, Vida Imani, Mohsen Saffari, Mark D Griffiths, and Amir H Pakpour. 2020. The fear of COVID-19 scale: development and initial validation. *International journal of mental health and addiction* (2020).
- [2] Yunis Alam and Charles Husband. 2013. Islamophobia, community cohesion and counter-terrorism policies in Britain. *Patterns of Prejudice* 47, 3 (2013), 235–252.
- [3] Dimo Angelov. 2020. Top2Vec: Distributed Representations of Topics. (2020). arXiv:2008.09470 [cs.CL]
- [4] Matthew PJ Ashby. 2020. Initial evidence on the relationship between the coronavirus pandemic and crime in the United States. *Crime Science* 9 (2020), 1–16.
- [5] Pinkesh Badjatiya, Shashank Gupta, Manish Gupta, and Vasudeva Varma. 2017. Deep learning for hate speech detection in tweets. In *Proceedings of the 26th International Conference on World Wide Web Companion*. 759–760.
- [6] Jayshree Bajoria. 2020. Coronajihad is Only the Latest Manifestation: Islamophobia in India has Been Years in the Making. <https://www.hrw.org/news/2020/05/01/coronajihad-only-latest-manifestation-islamophobia-india-has-been-years-making> [Online; accessed 11-November-2020].
- [7] Linda Briskman et al. 2015. The creeping blight of Islamophobia in Australia. *International Journal for Crime, Justice and Social Democracy* 4, 3 (2015), 112–121.
- [8] Mohit Chandra, Ashwin Pathak, Eesha Dutta, Paryul Jain, Manish Gupta, Manish Shrivastava, and Ponnurangam Kumaraguru. 2020. AbuseAnalyzer: Abuse Detection, Severity and Target Prediction for Gab Posts. In *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, Barcelona, Spain (Online), 6277–6283. <https://doi.org/10.18653/v1/2020.coling-main.552>
- [9] Despoina Chatzakou, Nicolas Kourtellis, Jeremy Blackburn, Emiliano De Cristofaro, Gianluca Stringhini, and Athena Vakali. 2017. Measuring# gamergate: A tale of hate, sexism, and bullying. In *Proceedings of the 26th international conference on world wide web companion*. 1285–1290.
- [10] Emily Chen, Kristina Lerman, and Emilio Ferrara. 2020. Tracking Social Media Discourse About the COVID-19 Pandemic: Development of a Public Coronavirus Twitter Data Set. *JMIR Public Health Surveill* 6, 2 (29 May 2020), e19273. <https://doi.org/10.2196/19273>
- [11] Hyunghoon Cho, Daphne Ippolito, and Yun William Yu. 2020. Contact tracing mobile apps for COVID-19: Privacy considerations and related trade-offs. *arXiv preprint arXiv:2003.11511* (2020).
- [12] Sabri Ciftci. 2012. Islamophobia and threat perceptions: Explaining anti-Muslim sentiment in the West. *Journal of Muslim Minority Affairs* 32, 3 (2012), 293–309.
- [13] Melanie Coates. 2020. Covid-19 and the rise of racism. *bmj* 369 (2020).
- [14] Thomas Davidson, Dana Warmsley, Michael Macy, and Ingmar Weber. 2017. Automated Hate Speech Detection and the Problem of Offensive Language. arXiv:1703.04009 [cs.CL]
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [16] Mai ElSherief, Vivek Kulkarni, Dana Nguyen, William Yang Wang, and Elizabeth Belding. 2018. Hate lingo: A target-based linguistic analysis of hate speech in social media. *arXiv preprint arXiv:1804.04257* (2018).
- [17] Emilio Ferrara. 2020. # covid-19 on twitter: Bots, conspiracies, and social media activism. *arXiv preprint arXiv:2004.09531* (2020).
- [18] Casey Fiesler and Nicholas Proferes. 2018. "Participant" Perceptions of Twitter Research Ethics. *Social Media+ Society* 4, 1 (2018), 2056305118763366.
- [19] Antigoni Maria Founta, Despoina Chatzakou, Nicolas Kourtellis, Jeremy Blackburn, Athena Vakali, and Ilias Leontiadis. 2019. A unified deep learning architecture for abuse detection. In *Proceedings of the 10th ACM Conference on Web Science*. 105–114.

¹²<https://www.perspectiveapi.com/>

- [20] Scott D Halpern, Robert D Truog, and Franklin G Miller. 2020. Cognitive bias and public health policy during the COVID-19 pandemic. *Jama* 324, 4 (2020), 337–338.
- [21] Shaikh Azizur Rahman Hannah Ellis-Petersen. 2020. Coronavirus conspiracy theories targeting Muslims spread in India. <https://www.theguardian.com/world/2020/apr/13/coronavirus-conspiracy-theories-targeting-muslims-spread-in-india> [Online; accessed 11-November-2020].
- [22] Mathieu Jacomy, Tommaso Venturini, Sebastien Heymann, and Mathieu Bastian. 2014. ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLOS ONE* 9, 6 (06 2014), 1–12. <https://doi.org/10.1371/journal.pone.0098679>
- [23] Suhasini Raj Jeffrey Gettleman, Kai Schultz. 2020. In India, Coronavirus Fans Religious Hatred. <https://www.nytimes.com/2020/04/12/world/asia/india-coronavirus-muslims-bigotry.html> [Online; accessed 11-November-2020].
- [24] Alexa Alice Joubin. 2020. Anti-Asian Racism during COVID-19 Pandemic, *GW Today*, April 20, 2020. *GW Today* (2020).
- [25] Rebecca Killick, Paul Fearnhead, and Idris A Eckley. 2012. Optimal detection of changepoints with a linear computational cost. *J. Amer. Statist. Assoc.* 107, 500 (2012), 1590–1598.
- [26] Ramez Kouzy, Joseph Abi Jaoude, Afif Kraitem, Molly B El Alam, Basil Karam, Elio Adib, Jabra Zarka, Cindy Traboulsi, Elie W Akl, and Khalil Baddour. 2020. Coronavirus goes viral: quantifying the COVID-19 misinformation epidemic on Twitter. *Cureus* 12, 3 (2020).
- [27] Annette Markham and Elizabeth Buchanan. 2012. Ethical decision-making and internet research: Version 2.0. recommendations from the AoIR ethics working committee. Available online: aoir.org/reports/ethics2.pdf (2012).
- [28] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [29] Alan Mislove and Christo Wilson. 2018. A practitioner’s guide to ethical web data collection. In *The Oxford Handbook of Networked Communication*. Oxford University Press.
- [30] Nicola Montemurro. 2020. The emotional impact of COVID-19: From medical staff to common people. *Brain, behavior, and immunity* (2020).
- [31] Ji Ho Park and Pascale Fung. 2017. One-step and two-step classification for abusive language detection on twitter. *arXiv preprint arXiv:1706.01206* (2017).
- [32] Billy Perrigo. 2020. It Was Already Dangerous to Be Muslim in India. Then Came the Coronavirus. <https://time.com/5815264/coronavirus-india-islamophobia-coronajihad> [Online; accessed 11-November-2020].
- [33] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. *arXiv:1908.10084* [cs.CL]
- [34] Arunabh Saikia. 2020. The other virus: Hate crimes against India’s Muslims are spreading with Covid-19. <https://scroll.in/article/958543/the-other-virus-hate-crimes-against-indias-muslims-are-spreading-with-covid-19>
- [35] Leonard Schild, Chen Ling, Jeremy Blackburn, Gianluca Stringhini, Yang Zhang, and Savvas Zannettou. 2020. " go eat a bat, chang!": An early look on the emergence of sinophobic behavior on web communities in the face of covid-19. *arXiv preprint arXiv:2004.04046* (2020).
- [36] Wiktor Soral, James Liu, and Michał Bilewicz. 2020. Media of contempt: Social media consumption predicts normative acceptance of anti-Muslim hate speech and islamoprejudice. *International Journal of Conflict and Violence (IJCV)* 14, 1 (2020), 1–13.
- [37] Bharath Syal. 2020. CoronaJihad: Stigmatization of Indian Muslims in the COVID-19 Pandemic. <http://southasiajournal.net/coronajihad-stigmatization-of-indian-muslims-in-the-covid-19-pandemic/> [Online; accessed 11-November-2020].
- [38] Yla R Tausczik and James W Pennebaker. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of language and social psychology* 29, 1 (2010), 24–54.
- [39] Leanne Townsend and Claire Wallace. 2016. Social media research: A guide to ethics. *Aberdeen: University of Aberdeen* (2016).
- [40] Bertie Vidgen, Austin Botelho, David Broniatowski, Ella Guest, Matthew Hall, Helen Margetts, Rebekah Tromble, Zeerak Waseem, and Scott Hale. 2020. Detecting East Asian Prejudice on Social Media. *arXiv preprint arXiv:2005.03909* (2020).
- [41] Bertie Vidgen and Taha Yasseri. 2020. Detecting weak and strong Islamophobic hate speech on social media. *Journal of Information Technology & Politics* 17, 1 (2020), 66–78.
- [42] Jessica Vitak, Katie Shilton, and Zahra Ashktorab. 2016. Beyond the Belmont principles: Ethical challenges, practices, and beliefs in the online data research community. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. ACM, 941–953.
- [43] Zeerak Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of the NAACL student research workshop*. 88–93.
- [44] Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2019. BERT Post-Training for Review Reading Comprehension and Aspect-based Sentiment Analysis. *arXiv:1904.02232* [cs.CL]
- [45] Jiacheng Xu, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Discourse-Aware Neural Extractive Text Summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 5021–5031. <https://doi.org/10.18653/v1/2020.acl-main.451>
- [46] Dingqi Yang, Terence Heaney, Alberto Tonon, Leye Wang, and Philippe Cudré-Mauroux. 2018. CrimeTelescope: Crime Hotspot Prediction Based on Urban and Social Media Data Fusion. *World Wide Web* 21, 5 (Sept. 2018), 1323–1347.
- [47] Atefeh Zandifar and Rahim Badrfam. 2020. Iranian mental health during the COVID-19 epidemic. *Asian journal of psychiatry* 51 (2020).
- [48] Savvas Zannettou, Joel Finkelstein, Barry Bradlyn, and Jeremy Blackburn. 2020. A Quantitative Approach to Understanding Online Antisemitism. *Proceedings of the International AAAI Conference on Web and Social Media* 14, 1 (May 2020), 786–797. <https://www.aaai.org/ojs/index.php/ICWSM/article/view/7343>
- [49] Caleb Ziems, Bing He, Sandeep Soni, and Srijan Kumar. 2020. Racism is a Virus: Anti-Asian Hate and Counterhate in Social Media during the COVID-19 Crisis. *arXiv preprint arXiv:2005.12423* (2020).
- [50] Michael Zimmer and Katharina Kinder-Kurlanda. 2017. *Internet research ethics for the social age: New challenges, cases, and contexts*. Peter Lang International Academic Publishers.