"A Virus has No Religion": Analyzing Islamophobia on Twitter During the COVID-19 Outbreak

Mohit Chandra*, Manvith Reddy*, Shradha Sehgal*, Saurabh Gupta, Arun Balaji Buduru, Ponnurangam Kumaraguru

International Institute of Information Technology - Hyderabad Indraprastha Institute of Information Technology - Delhi

* Authors contributed Equally to this Research









Motivation

- The global Coronavirus (COVID-19) outbreak has resulted in increased levels of fear, anxiety, and outbursts of hateful emotions among the population.
- Islamic missionary movement *Tablighi Jamaat*'s events, that happened in multiple countries of South-East Asia has caused a rising tide of Islamophobia that has not received due attention from the research community.
- In this study, we aim to study how Islamophobia has manifested itself across much of south-east asia with a special focus on India, where the bulk of it was found.



Tablighi Jamaat Ha congregation Mus Superspreader event #Congregation

Superspreader event (March '20) Hate around Muslims online

#CoronaJihad #MuslimVirus



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?



CoronaBias Dataset

- Filtered data collected by **Chen et al.**, by only considering tweets written in Roman script (English) and restricting the timeframe from February 1, 2020, till May 31, 2020. We further filtered this dataset using 64 Muslim related keywords that consisted of positive, negative and neutral terms.
- Classified the tweets using two labels 'Hateful' or 'Non-Hateful' by manually annotating 2000 randomly picked tweets and running it through a BERT Model.

410,990 Tweets

244, 229 unique users

0.812 Fleiss' Kappa score

Method	Accuracy	Recall	Precision	F-1
BERT	0.854 ± 0.088	0.854 ± 0.087	0.854 ± 0.086	0.853 ± 0.088



RQ1 : Temporal Analysis



- Conducted Temporal analysis to investigate the relation between offline and online events.
- Assigned ranks to events by using the **PELT** change-point detection algorithm. Significance of change points increases as the rank decreases, with 1 being the most crucial change points.

Event Rank	Date (MMDD-YYYY)	Event Description
1	03-31-2020	On 31st March, reports linked the Tablighi Jamaat 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.



RQ1 : Semantic Similarity

- To investigate the change in context around the muslim community, we trained Word2Vec Continuous-Bag-of-Words (CBOW) models for each month for both reference words(Muslim, Virus) and observed the top 5 most semantically similar words in the dataset.
- There is a clear pattern in both tables that shows the growing association between the coronavirus (COVID-19) and the Muslim Community.
- Through this experiment, we showed that the context around the Muslim community changed through the course of the pandemic.

March	April	May
Coronavirus	Coronavirus	Coronavirus
(0.446)	(0.491)	(0.540)
Disease	Covid19	Covid19
(0.411)	(0.427)	(0.481)
Desease	Muslims	Disease
(0.380)	(0.406)	(0.471)
Muslims	Corona	Muslims
(0.369)	(0.382)	(0.457)
Allah	Disease	Corona
(0.342)	(0.363)	(0.434)
China	People	Muslim
(0.2817)	(0.349)	(0.379)

Virus

Muslim

March	April	May
Islam	Islam	People
(0.434)	(0.444)	(0.457)
Bully	People	Islam
(0.288)	(0.389)	(0.432)
China	China	China
(0.283)	(0.346)	(0.402)
Moslem	Virus	Covid19
(0.281)	(0.329)	(0.386)
Virus	Covid19	coronavirus
(0.275)	(0.314)	(0.361)
Seeing	US (0.293)	India
(0.233)		(0.344)

RQ2 : Topic Modelling

Macro

Through this experiment, we analyzed the topics and aimed to establish the topics that were prevalent throughout the period of our study. We performed topic modelling at 2 levels : Macro and Micro topic modelling.

		S.No.	Topic	Tokens	No. of
					Tweets
ng	Non-Hateful	1	The Coronavirus pandemic is dis-	coronavirus distraction, muslims tortured, imprisoned raped,	5, 368
			tracting people from the Uyghur Muslim crisis.	concentration killed, thousands chinese	
		2	The Coronavirus Pandemic lead- ing to a global market collapse.	markets collapse, global demand, corona hits, find difficult, at- tention 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 donat- ing plasma	donate plasma, tablighijamaat members, members recovered, covid donate, patients	7, 555
		5	Anti-Abrahamic religious senti- ments attached with COVID-19.	kill muslims, china kill, hate muslims, hate christians, hate jews	6,089
	lul	6	Muslims attacking Covid-19 health workers	indian muslims, fight, pandemic, doctors, poor, violence, govt	28, 747
	Hate	7	Tablighi Jamaat religious congre- gation	tablighi jamaat, maulana, jamaat attendees, police, jamaat mem- bers, 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

Many hateful topics focused on different aspects of the Tablighi Jamaat event and its members. Large number of tweets were clustered into few key topics that assigned blame and expressed hate towards the muslim community.

Non-Hateful topics ranged from neutral topics referring to the global market collapse (Topic 1) as well as pro muslim topics describing plasma donation of Tablighi Jamaat members (Topic 4).

RQ2 : Topic Modelling

Micro

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 and considered a 10 day window around them.

Tabhlighi Jamaat Event (April 1st)			
Topic	Tokens		
India's efforts to contain virus foiled	every tablighi, indias efforts, se- vere setback, single cluster, efforts		
by Tablighis	containing, keep counting		
Islamic preacher praying to divert Covid-19 to Non- <i>Muslim</i> nations	allah divert, nonmuslim nations, infections nonmuslim, divert covid19, preacher praying		
Islamophobia taint- ing India as doctors refuse to treat <i>Mus-</i> <i>lim</i> patients	islamophobia taints, hospital refuses, muslim patient, crisis morality, news channels		

Occasion of Ramadan (April 25th)			
Topic	Tokens		
Criticising me-	hey islamophobics, islamophobics		
dia for spreading	medicine, indian media, tablighi-		
Discriminatory	Jamaat, social media		
behaviour towards	ferentiate, failure indian, denied		
Muslims	food, racist rhetoric		
Tablighi jamaat	plasma patients, members recov-		
willing to donate	ered, covid donate, cure others, tj		
plasma	members		

We observe a clear shift in sentiment of topics between both events. Around April 1st, a large number of tweets blamed the muslim community for the spread. By April 25th, more positive topics recognizing the issue of islamophobia were being talked about.



RQ3 : Analysis of User Bio's







0-25 % Hateful Tweets

25-50 % Hateful Tweets

50-100 % Hateful Tweets

- Class 1 users revealed the usage of non-hateful/general terms like Proud Indian, Indian Muslim. Many claimed to be Muslim and support ideologies of humanity and love.
- Class 2 users' word cloud has many terms related to Nationalism/patriotism but at the same time also has religious terms (both Pro-Muslim and Pro-Hinduism).
- Class 3 users' word cloud shows the sign of blatant Hinduism which is portrayed as patriotism through the usage of Nationalist terminologies and terms related to Islamophobia



RQ3 : User Network and Activity Analysis



- Users who posted a majority of 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.
- We observed two other clusters of users that were predominantly green. This shows that there are also communities that posted majorly non-hateful content.
- 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**.



RQ4 : External URL Experiment

- 68.5% of tweets in the CoronaBias Dataset contained an external URL.
- We observed a large number of URLs from a few websites that have been known to disseminate anti-Muslim news (OpIndia, Jihadwatch, and Swarajyamag).
- Among the popular social media platforms, we found the order as Facebook (579) followed by Instagram (322), Tumblr (155), and Linkedin (62).

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



RQ4 : External URL Experiment

- The content from OpIndia refers to Muslims in great detail and often in a negative light.
- The content presented by BBC varied and focused on global topics.
- The content referred from Youtube lay in between that of OpIndia and BBC.







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

- Perspective API used to calculate toxicity score
- 86% of the tweets containing a BBC link lied in the 0.1-0.2
- Tweets referring OpIndia were more evenly distributed in the range of 0.1-0.5. We also observe tweets in higher toxicity brackets 0.5 0.8.





Contributions

- We present the first of its kind *CoronaBias* dataset with over 410, 990 tweets from 244, 229 unique users, which we have publicly released.
- 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.
- Find a growing association of the Muslim community with the COVID-19 pandemic through the semantic similarity experiment
- Extracting the prevalent topics of discourse in broad and focused windows of time through Macro and Micro topic modelling experiment.
- Specifically analyse user network, user activity, and descriptions of user bios to reveal the differentiating characteristics of *Hateful* and *Non-Hateful* users.
- Examine the nature and toxicity of content spread through external URL's referenced in the tweets of the *CoronaBias* dataset.



Thank You