



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Pei, X;Mehta, D

Title:

Multidimensional racism classification during COVID-19: stigmatization, offensiveness, blame, and exclusion

Date:

2022-12-01

Citation:

Pei, X. & Mehta, D. (2022). Multidimensional racism classification during COVID-19: stigmatization, offensiveness, blame, and exclusion. *Social Network Analysis and Mining*, 12 (1), <https://doi.org/10.1007/s13278-022-00967-9>.

Persistent Link:

<https://hdl.handle.net/11343/320280>

License:

[CC BY](#)



Multidimensional racism classification during COVID-19: stigmatization, offensiveness, blame, and exclusion

Xin Pei¹ · Deval Mehta²

Received: 27 February 2022 / Revised: 17 August 2022 / Accepted: 26 August 2022
© The Author(s) 2022

Abstract

Transcending the binary categorization of racist texts, our study takes cues from social science theories to develop a multi-dimensional model for racism detection, namely stigmatization, offensiveness, blame, and exclusion. With the aid of BERT and topic modelling, this categorical detection enables insights into the underlying subtlety of racist discussion on digital platforms during COVID-19. Our study contributes to enriching the scholarly discussion on deviant racist behaviours on social media. First, a stage-wise analysis is applied to capture the dynamics of the topic changes across the early stages of COVID-19 which transformed from a domestic epidemic to an international public health emergency and later to a global pandemic. Furthermore, mapping this trend enables a more accurate prediction of public opinion evolution concerning racism in the offline world, and meanwhile, the enactment of specified intervention strategies to combat the upsurge of racism during the global public health crisis like COVID-19. In addition, this interdisciplinary research also points out a direction for future studies on social network analysis and mining. Integration of social science perspectives into the development of computational methods provides insights into more accurate data detection and analytics.

Keywords Racism · COVID-19 · Social media · Deviant behaviours

1 Introduction

The global outbreak of COVID-19 has been accompanied by the worldwide upsurge of racism. An increasing research stream has illuminated the more infectious nature of racist reactions than coronavirus, which is leading towards a more harmful social consequence (Kapilashrami and Bhui 2020; Wang et al. 2021). BBC has reported that the United Nations raised racially motivated violence and other hate incidents against Asian Americans to “an alarming level” in 2020 (BBC 2021). And hate crimes occurred in the New York City in 2020 experienced a ninefold increase from the

previous year. Therefore, it has become urgent to comprehend the racist discourse so as to enact effective intervention strategies to prevent the escalation of deviant behaviours such as hate crimes and social exclusion during COVID-19.

Against this backdrop, many studies have drawn attention to social media platforms which provide critical avenues for pandemic-related public discussion. Scholars have widely adopted highly advanced computational methods and state-of-the-art language models for big social data analytics on these platforms, with the purpose of achieving a better understanding of racist reactions from the public. Unsupervised machine learning techniques such as topic modelling, keyword clustering have been widely employed in studies (e.g. Tahmasbi et al. 2021) for analysing Twitter and Reddit data during COVID-19. Scholars (He et al. 2021; Lu and Sheng 2020) have also adopted supervised learning methods such as support-vector machines (SVMs) and Transformers for hate and racist speech detection.

Despite the contribution in technical advancement, the extant literature shows the tendency of neglecting the theoretical foundation for data detection and analysis—that is how to define racism in the first place. To specify, the existing computational techniques and models tend to apply a

Xin Pei and Deval Mehta have contributed equally to this work.

✉ Xin Pei
xin.pei@unimelb.edu.au
Deval Mehta
Deval.mehta1092@gmail.com

¹ School of Culture and Communication, The University of Melbourne, Parkville, Melbourne, VIC 3010, Australia

² Monash eResearch Center, Monash University, Clayton, Melbourne, VIC 3800, Australia

binary definition that primarily categorises the linguistic features of texts into either the racist or non-racist ones (e.g. He et al. 2021). It is important to note that some studies mentioned different dimensions for racism identification. For instance, study by Davidson and colleagues (2017) utilized offensive language for automatic detection of hate speech. In the same vein, some other studies (Fan et al. 2020; Liu et al. 2022) particularly focused on stigmatization. However, they are not only restricted in numbers but also lack of a comprehensive model based on a summary of relevant indicators from social science studies. Given the dynamic nature of racist behaviours (Richeson 2018), a comprehensive classification capturing more nuances of racist discourse will allow for more insights into the behavioural change across different stages of COVID-19.

To fill this research gap, our study transcends the binary of (non)racism by introducing a model that classifies racist behaviours into four categories—stigmatization, offensiveness, blame, and exclusion. It is important to note that this model is built upon a combination of social science theories and computational methods. To specify, while the categorization is generated from prior scholarly discussion on racism across the domains of sociology, psychology, and social psychology, the application of the model involves deep learning techniques - BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al. 2018) and topic modelling (Blei et al. 2003).

Our study makes unique contribution that enriches the scholarly discussion on deviant racist behaviours on social media. First, applying this model on a stage-wise analysis, our study captures the dynamic evolvement of racist behaviours across the early development of COVID-19 - how the four racist categories competed with one another at different stages, and how the themes of each category shifted across time. Furthermore, mapping this trend will enable a more accurate prediction of public opinion evolvement concerning racism in the offline world, and meanwhile, the enactment of specified intervention strategies to combat the upsurge of racism during the global public health crisis like COVID-19. In addition, this interdisciplinary research also enhances a direction for future studies on social network analysis and mining. Integration of social science perspectives into the development of computational methods can provide a new route for a more accurate data detection and analytics.

2 Literature review

2.1 Racism, social media, and COVID-19

Many have argued that social media platforms are providing critical avenues for racist opinion expression. Especially, the widely advocated speech freedom on social media platforms

is paving the way for toxic and provocative languages replete with trolls, often with the target at a particular race/ethnicity, nation, or (im)migrant community (Lim 2017). Anonymity further enables hate speech and biased opinions to avoid detection (Keum and Miller 2018). Meanwhile, boundless connectivity on social media platforms allows racist opinions to travel at a fast speed and to reach a broad scope of audiences (He et al. 2021). Moreover, such connectivity also permits people with similar racial ideologies to cluster and collectively build up the racist discourse to increase its visibility and influence online (Kapilashrami and Bhui 2020). Therefore, Matamoros-Fernández (2017) coined the term “platformed” racism to refer to people’s usage of affordances on different social media platforms to duplicate and extend the offline social inequalities. (Oboler 2016) indicated the emergence of ‘Hate 2.0’, under which the repetitive occurrence of hate speech on social media keeps on justifying the racist discourse as a normalized collective behaviour.

It is important to note the rise of racism on social media during COVID-19. As one of the most severe global pandemics since the turn of the new millennium, COVID-19 has caused more than forty million confirmed cases and almost five million deaths across the globe till the submission date of this manuscript. Suspected to be originated from Wuhan, China, the global outbreak of COVID-19 has widely raised social exclusion against China, which has been evolved into discrimination, bias, and even hatred against Chinese and even Asians at large. This phenomenon has been unofficially coined as sinophobia.

It is worth noting that racism has been largely extended to the online world under the pandemic. On 16 March 2020, a post from the official Twitter account of Donald Trump, the former president of the USA, referred to COVID-19 as Chinese virus. Ironically, this overtly racist and xenophobic label immediately became an emergent popular hashtag - #chinesevirus, which was massively disseminated and circulated on Twitter and other social media platforms. Besides #chinesevirus, social media platforms have witnessed the proliferation of many other offensive hashtags centring on a particular race and nation in the pandemic context, such as #kungflu embodying the conflation of coronavirus with racial/ethnic cultural identities, and #boycottchina manifesting social exclusion. Hate speech and discriminative opinions are massively circulated and disseminated through these hashtags. Consequently, many scholars have initiated the investigation into racist reactions on social media during COVID-19. The following section will elaborate on the contribution and drawbacks of the research stream using computational methods.

2.2 Bridging social science theory and computational methods

Many studies have adopted computational methods for big data mining and social network analytics to better understand the dynamics of the racist deviant behaviours on social media platform from a macro-level. For example, (Garland et al. 2020) used the supervised machine learning models such as random forests and support-vector machines (SVMs) for classification of hate speech. Leveraging the large textual data corpus, studies have also used unsupervised techniques such as word2vec (Mikolov et al. 2013), Glove embeddings (Pennington et al. 2014) for topic clustering and keywords analysis. More recently, with the advent of deep learning and availability of training data, models such as long short-term memory networks (LSTMs) (Hochreiter and Schmidhuber 1997), RNNs, and much recently Transformers such as BERT (Devlin et al. 2018) have been employed in most studies (He et al. 2021; Lu and Sheng 2020) for classification of textual data on social media platforms.

Regardless of the contribution made to advancing the tools and techniques, prior studies tend to ignore the most fundamental issue that shall be addressed in the first place - that is how to define racism. Especially, many studies tend to adopt a binary classification of linguistic features that categorizes the texts into either racist or non-racist ones. Although some studies have committed the efforts to enriching the linguistic features of hateful and offensive speech (Abderrouaf and Oussalah 2019; Fahim and Gokhale 2021), the extant classification still tends to largely underestimate the complexity of racist behaviors, thereby leading to an oversimplified mechanism for racist data detection and analysis. This tends to prevent the discovery of the nuances of the themes embodied in the racist opinion expression, and the dynamics of the themes that are very likely to evolve alongside the development of a public event (Pei and Mehta 2020).

To fill this research gap, our study proposes a multidimensional model to detect and classify racism which is built upon the conceptualization of racism in prior social science research. To specify, transcending the binary category, our model specifies racist behaviours into stigmatization, offensiveness, blame, and exclusion. Taking references from the study by (Miller and Kaiser 2001) centring on theorizing stigma, our model defines stigma as confirming negative stereotypes for conveying a devalued social identity within a particular context. Similarly, built upon the research by (Jeshion 2013) surrounding the expression of offensive slurs, we refer to offensiveness as attacking a particular social group through aggressive and abusive language. The study by (Coombs and Schmidt 2000) in the context of Texaco's racism crisis points a direction for framing blame as attributing the responsibility for the negative consequences of

the crisis to one social group. The dimension of exclusion stems from the study by (Bailey and Harindranath 2005) that noted exclusion as a critical step of racializing others which embodies the process of othering to draw a clear boundary between in-group and out-group members. Please refer Table 1 which includes the definition of the four dimensions accompanied by the corresponding examples from the dataset.

This multidimensional classification model is applied to a stage-wise analysis, with the purpose of mapping the dynamics - how these four racist themes were competing with one another alongside the development of COVID-19. This will provide a more nuanced idea about the trend regarding the possibly shifting focus of public opinion concerning racism. Especially, we focus on the most turbulent early phase of COVID-19 (Jan to Apr 2020) where the unexpected and constant global expansion of the virus kept on changing people's perception of this public health crisis and how it is related to race and nationality. To specify, this research divides the early phase into three stages based on the changing definitions of COVID-19 made by the World Health Organization (WHO) - (1) 1 to 31 Jan 2020 as a domestic epidemic referred to as stage 1 (S1); (2) 1 Feb to 11 Mar 2020 as an International Public Health Emergency (after the announcement made by WHO on 1 Feb) referred to as stage 2 (S2); (3) 12 Mar to 30 Apr 2020 as a global pandemic (based on the new definition given by WHO on 11 Mar) referred to as stage 3 (S3). We select Twitter, the most influential platform for political online discussion, as the field for data mining and analysis.

3 Data and methods

This section deals with five parts - first, it outlines method used to scrape the data; second, it defines the four dimensions of racism; third, it describes the process of annotation; fourth, it explains the method employed for category-based racism and xenophobia detection; and last, it details the process of topic modelling employed for extracting topics from the categorized data.

Dataset of this research is comprised of 247,153 tweets extracted through Tweepy API.¹ We built a custom python-based wrapper utilizing the Tweepy API functionalities to continuously scrape the data starting from the 1 January until the 30 April 2020, which falls within our interest period of early covid-19 including the three durations of a domestic epidemic, an International Public Health Emergency, and eventually a global pandemic as highlighted in the paper.

¹ <https://www.tweepy.org/>.

For selecting the hashtags to scrape the Twitter data we first used the most common and topmost hashtags—#chinesevirus and #chinavirus, which were also used by the other popular studies (Tahmasbi et al. 2021; Ziems et al. 2020) of analysing covid-19 data on Twitter. For selecting all the other hashtags in our data scraping process, we followed a dynamic hashtag selection process which is a similar approach to the study (Srikanth et al. 2019) relevant for “rapidly-evolving online datasets”. Our strategy of dynamic hashtag scraping involves the following steps:

- After scraping a sample of 500 tweets from the topmost hashtags, we collect the most frequent top five hashtags occurring in those samples of 500 tweets.
- The collected top five most frequent hashtags (excluding the ones which were used to collect them) are then used to scrape new tweet samples containing them.
- We then repeat the first step for the new sample of the scraped tweets and collect the new most frequent top five hashtags having at least 50 occurrences.
- In the above step, we noticed that most of the new hashtags do not have a high frequency of repetition, so we stop the recursive scraping involving the first step at this stage.

The above process is adopted for the first two weeks at the beginning of a new stage of the data collection in our duration of scraping, i.e. this strategy is first employed on the 1 January 2020 (beginning of stage 1); then repeated on the 1 February 2020 (beginning of stage 2) and on 12 March 2020 (beginning of stage 3).

With this strategy, we developed the following list of hashtags which were then used to mine the data - #chinavirus, #chinesevirus, #boycottchina, #ccpvirus, #chinaflu, #china_is_terrorist, #chinaliedandpeopledied, #chinalied-peopledied, #chinalies, #chinamustpay, #chinapneumonia, #chinazi, #chinesebioterrorism, #chinesepeumonia, #chinesevirus19, #chinesewuhanvirus, #viruschina, and #wufu. The extracted tweets from the above hashtags are further divided into three stages that define the early development of Covid-19 as mentioned earlier. We show the number of tweets extracted for each day using this method in Fig 1.

3.1 Method

3.1.1 Category-based racism and xenophobia detection

Beyond a binary categorization of racism and xenophobia, this research applies the perspective of social science to categorizing racism and xenophobia into four dimensions as demonstrated in Table 1. This basically translates into a problem of five class classification of text data, where four

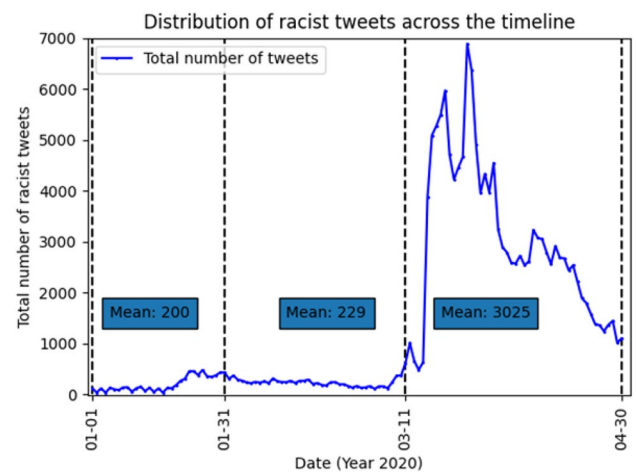


Fig. 1 Analysis of the number of tweets returned daily by our custom wrapper (based on Tweepy API) from 1 Jan to 30 Apr 2020

classes represent the four types of racism, and the fifth class refers to the category of non-racist and nonxenophobic.

3.1.2 Annotated dataset

To train machine learning and deep learning classifiers for this task, we aimed to build a reasonable size dataset (not so large that it becomes difficult to annotate, and not so small enough to compromise on proper training and evaluation of our methods). Thus, we selected 6,000 as the number of representative tweets that can be utilized for training the model. To evenly represent the opinions from the three stages, we divided the selection of 6000 tweets into 2000 tweets from each stage (S1, S2, S3). These 2000 tweets were then randomly selected from each development stage.

We hired four research assistants who were initially trained under the supervision of one co-author on a pilot data of 200 tweets to categorize them into four different categories based on their definitions. The annotation followed a coding method with 0 representing stigmatization, 1 for offensiveness, 2 for blame, and 3 for exclusion in alignment with the linguistic features of the tweets. The non-marked tweets were regarded as non-racist and non-xenophobic and represented class category 4. We limited the annotation for each tweet to only one label which aligned with the strongest category. All four research assistants had Asian ethnicity (Chinese).

After completing their initial training and review from the co-author, they were provided feedback if there was a dispute in labelling. The four research assistants reached overall inter-coder reliability above 70%, which is a moderate to high threshold selected by prior studies (Guntuku et al. 2019; Jaidka et al. 2019) for reliability of data annotation. Post this pilot data training of the four research assistants,

Table 1 Definition and example of categorization of racist and xenophobic behaviours

Category	Definition	Example
Stigmatization	Confirming negative stereotypes for conveying a devalued social identity within a particular context (Miller and Kaiser 2001)	For all the #ChinaVirus jumped from a bat at the wet market
Offensiveness	Attacking a particular social group through aggressive and abusive language (Jeshion 2013)	Real misogyny in communist China. #chinazi #China_is_terrorist #China_is_terrorists #FuckTheCCP
Blame	Attributing the responsibility for the negative consequences of the crisis to one social group (Coombs and Schmidt 2000)	These Chinese are absolutely disgusting. They spread the #ChineseVirus. Their lies created a pandemic #ChinaMustPay
Exclusion	The process of othering to draw a clear boundary between in-group and out-group members (Bailey and Harindranath 2005)	China deserves to be isolated by all means forever. SARS was also initiated in China, 2003 by eating anything & everything #BoycottChina

they were then given 500 tweets from each stage to categorize, which enabled us to make our dataset of 6000 tweets across the three stages. The distribution of 6000 tweets amongst the five classes is as follows—1318 stigmatization, 1172 offensive, 1045 blame, 1136 exclusion, and 1329 non-racist and non-xenophobic.

We view the task of classification of the above-mentioned categories as a supervised learning problem and target developing machine learning and deep learning techniques for the same. We firstly pre-process the input data text by removing punctuation and URLs from a text sample and converting it to lower case before providing it to train our models. We split the data into random train and test splits with 90:10 ratio for training and evaluating the performance of our models respectively by using the standard fivefold cross-validation.

3.2 BERT

Recently, word language models such as Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al. 2018) have become extremely popular due to their state-of-the-art performance on natural language processing tasks. Due to the nature of bidirectional training of BERT, it can learn the word representations from unlabelled text data powerfully and enables it to have a better performance compared to the other machine learning and deep learning techniques (Devlin et al. 2018). The common approach for adopting BERT for a specific task on a smaller dataset is to fine-tune a pre-trained BERT model which has already learnt the deep context-dependent representations. We select the “bert-base-uncased” model which comprises of 12 layers, 12 self-attention heads, a hidden size of 768 totalling 110 M parameters. We fine-tune the BERT model with a categorical cross-entropy loss for the five categories. The various hyperparameters used for fine-tuning the BERT model are selected as recommended from the paper (Devlin et al. 2018). We use the AdamW optimizer with the standard learning rate of

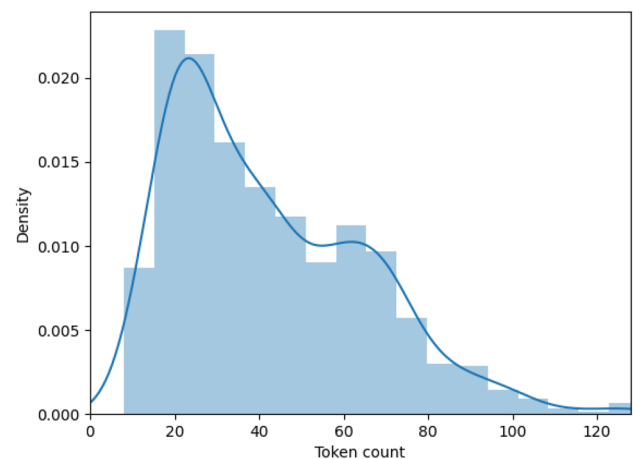


Fig. 2 Density distribution of token lengths of the tweets in our dataset

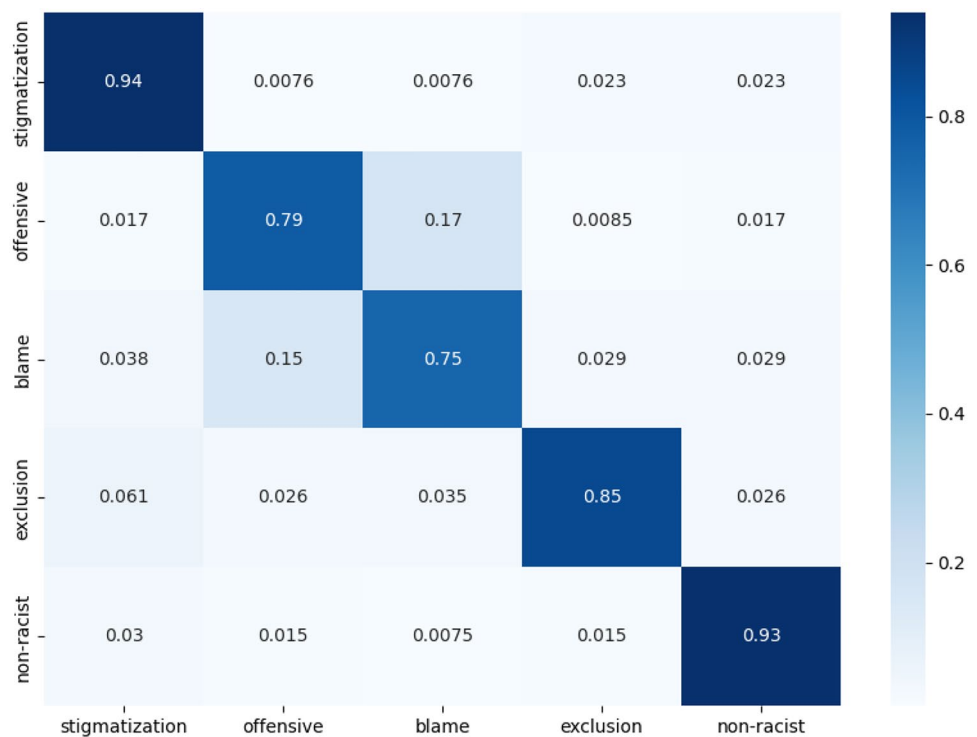
$2e-5$, a batch size of 16, and train it for 5 epochs. For selecting the maximum length of the sequences, we tokenize the whole dataset using Bert tokenizer and check the distribution of the token lengths. We notice that the minimum value of token length is 8, maximum is 130, median is 37 and mean is 42. Based on the density distribution shown in Fig. 2, we experiment with two values of sequence length - 64 and 128 and find that the sequence length of 64 provides a better performance.

As additional baselines, we also train two more techniques. Long short-term memory networks (LSTMs) (Hochreiter and Schmidhuber 1997) have been very popular with text data as they can learn the dependencies of various words in the context of a text. Also, machine learning algorithms such as support-vector machine (SVMs) (Hearst et al. 1998) have been used previously by researchers for text classification tasks. Moreover, the use of various feature extraction techniques such as term frequency inverse document frequency (TF-IDF), word2vec and Bag-of-Words (BoW) has proven to improve the performance of the classifiers (Li

Table 2 Performance of different models on the manually annotated test dataset.

Technique	Accuracy(%)	F1-score
SVM	69.04	0.66
SVM + TF-IDF features	75.8	0.74
SVM + BOW features	71.6	0.70
SVM + Word2Vec features	71.4	0.70
LSTM	74.01	0.72
BERT	86.06	0.81

Mean accuracy and f1-score for the fivefolds

Fig. 3 Confusion matrix (averaged for the fivefolds of validation data) of our trained BERT model for racism classification

et al. 2019; Zhang et al. 2010). The work in (Gebre et al. 2013) explored the use of TF-IDF on machine learning classifiers such as SVM and it was found that it helped to improve the classification performance significantly compared to the original baselines. It was found that TF-IDF and BoW perform the best with uni-gram or bi-gram collection of word features. We operate the bi-gram BoW, bi-gram TF-IDF and Word2Vec feature engineering techniques for training our SVM model. We adopt the same data pre-processing and implementation technique as mentioned earlier and train the SVM with grid search, a 5-layer LSTM (using the pre-trained Glove (Pennington et al. 2014) embeddings) and BERT model for the category detection of the racist and xenophobic tweets.

For evaluating the machine learning and deep learning approaches on our test dataset, we use the metrics of average

accuracy and weighted f1-score for the five categories. The performance of the models is shown in Table 2. It can be seen from Table 2 that the fine-tuned BERT model performs the best compared to SVM and LSTM in terms of both accuracy and f1 score. Although adding engineered features from TF-IDF improves the performance of the SVM classifier significantly, it cannot surpass the performance of the BERT model. Thus, we employ this fine-tuned BERT model for categorizing all the tweets from the remaining dataset. Having employed BERT on the remaining dataset, we get a refined dataset of the four categories of tweets spreaded across the three stages as shown in Table 2.

We also calculate the confusion matrix for our best performing model BERT as shown in Fig 3. As can be seen from the confusion matrix, we obtain an excellent classification performance (>0.90) for stigmatization and non-racism classification categories, a higher performance (>0.85 while <0.90) for the exclusion category, and moderately higher performance (>0.75 while <0.85) for the other two remaining categories of offensiveness and blame.

3.3 Topic modelling

Topic modelling is one of the most extensively used methods in natural language processing for finding relationships across text documents, topic discovery and clustering, and extracting semantic meaning from a corpus of unstructured

data (Jelodar et al. 2019). Many techniques have been developed by researchers such as Latent Semantic Analysis (LSA) (Deerwester et al. 1990), Probabilistic Latent Semantic Analysis (pLSA) (Hofmann 1999) for extracting semantic topic clusters from the corpus of data. In the last decade, Latent Dirichlet Allocation (LDA) (Blei et al. 2003) has become a successful and standard technique for inferring topic clusters from texts for various applications such as opinion mining (Zhai et al. 2011), social medial analysis (Cohen and Ruths 2013), event detection (Lin et al. 2010) and consequently there have also been various developed variants of LDA (Blei and McAuliffe 2010; Blei et al. 2003).

For our research, we adopt the baseline LDA model with Variational Bayes sampling from Gensim² and the LDA Mallet model (McCallum 2002) with Gibbs sampling for extracting the topic clusters from the text data. Before passing the corpus of data to the LDA models, we perform data pre-processing and cleaning which include the following steps. Firstly, we remove any new line characters, punctuations, URLs, mentions and hashtags. Later we tokenize the texts in the corpus and also remove any stopwords using the Gensim utility of pre-processing and stopwords defined in the NLTK³ corpus. Finally, we make bigrams and lemmatize the words in the text.

After employing the above pre-processing for our corpus, we employ topic modelling using LDA from Gensim and LDA Mallet. We perform experiments by varying the number of topics from 5 to 25 at an interval of 5 and checking the corresponding coherence score of the model (Fang et al. 2016). We train the models for 1000 iterations with varying number of topics, optimizing the hyperparameters every 10 passes after each 100 pass period. We set the values of α , β which control the distribution of topics and the vocabulary words amongst the topics to the default settings of 1 divided by the number of topics. We notice from our experiments that LDA Mallet has a higher coherence score (0.60–0.65) compared to the LDA model from Gensim (0.49–0.55) and thus we select LDA Mallet model for the task of topic modelling on our corpus of data.

The above strategy is employed for each racist and xenophobic category and for every stage individually. We find the highest coherence score corresponding to a specific number of topics for each category and stage. To analyse the results, we reduce the number of topics to five by clustering closely related topics using Eq. 1.

where N refers to the number of topics to be clustered, M represents the number of keywords in each topic, p_j corresponds to the probability of the word x_i in the topic, and T_c is the resultant topic containing the average probabilities

Table 3 Distribution of tweets amongst the four categories across the three stages

Category	Total	S1	S2	S3
Stigmatization	116,584	3723	5687	107,174
Offensiveness	10,503	1722	1808	6973
Blame	39,765	31	777	38,957
Exclusion	10,293	872	1341	8080

of all the words from the N topics. We then represent the top ten highest probability words in the resultant topic for every category and stage as is shown in Tables 4 to 7.

$$T_c = \frac{\left(\sum_{i=1}^N \sum_{j=1}^M p_j x_j \right)}{N} \quad (1)$$

4 Findings

Table 3 illustrates the distribution of racist tweets of the four categories across the three stages. Tables 4, 5, 6, and 7 demonstrate the ten most salient terms related to the generated five topics for each stage (S1, S2, and S3) of four categories. Each topic was summarized through the correlation between the ten terms. We put a question mark for topics from which no pattern can be generated. The below section provides a detailed analysis of the dynamics of the four categories.

4.1 Stigmatization

According to Table 3, stigmatization continuously acted as the dominant racist theme across the three early stages of COVID-19 (S1:3723; S2:5687; S3:107174). Based on the result generated from topic modelling (see Table 4), the main topics related to stigmatization at the first stage included “virus”, “China/Chinese”, “infection”, “outbreak” and “travel”. In general, stigmatization of this stage focused on the infectious nature of the virus, the association between the virus and China/Chinese, and the outbreak of the virus. Especially, the sub-topics under “China/Chinese” included “mask”, “animal”, and “eat”, which echoed the heated debates around China and Chinese at that time - to specify, whether wearing a mask that was advocated by China government would be helpful; and whether the origin of this virus was associated with the eating habits of Chinese people.

At the second stage, the leading topics changed to “emergency”, “globe”, “infection”, “China”, and “Chinese”. First, we noticed that due to the global outbreak of COVID-19 at this stage, expression of stigmatization started to pay

² <https://pypi.org/project/gensim/>.

³ <https://pypi.org/project/nltk/>.

Table 4 Extracted topics and their corresponding keywords for the category of stigmatization spread across the three stages S1, S2, and S3

S1	T1.Virus	Virus	Spread	Country	Travel	Year	Control	Chinese	Ban	Corona	Show
	T2.China/Chinese	Chinese	Virus	Deadly	China	Situation	Mask	Stop	Animal	Source	Eat
	T3.Infection	People	Case	Health	Infect	Confirm	Death	Sar	Number	Report	Market
	T4.Outbreak	China	Coronavirus	Wuhan	Outbreak	City	Hospital	News	Patient	Put	State
	T5.Travel	World	China	Government	Make	People	Time	Day	Bad	Flight	Start
S2	T1.Emergency	Virus	Spread	Day	Year	Corona	Show	Emergency	Food	Kit	Supply
	T2.Globe	China	World	Time	Country	Report	Death	Global	Health	Travel	Confirm
	T3.Infection	People	Case	Call	ncov	Infect	Kill	Pack	State	Flu	Number
	T4.China	China	Coronavirus	Wuhan	Outbreak	Quarantine	Stop	Find	Man	Dead	Thing
	T5.Chinese	Chinese	Make	Mask	Government	News	Good	Work	Citizen	Start	Respirator
S3	T1.Government	China	World	Spread	Country	Lie	Pay	Communist	Government	ccp	Make
	T2.?	Time	Make	India	Good	Give	Work	Day	Back	Fight	Buy
	T3.China	China	Coronavirus	Case	Death	Covid	Country	Economy	War	Number	Wuhan
	T4.Chinese	Chinese	Virus	People	Call	Stop	Racist	Start	Die	Blame	Corona
	T5.US	American	Trump	State	Medium	President	America	News	Great	Propaganda	Show

Table 5 Extracted topics and their corresponding keywords for the category of offensiveness spread across the three stages S1, S2, and S3

S1	T1.?	Country	ccp	Citizen	Virus	Arrest	Live	System	Security	Foreign	Understand
	T2.Government	People	Government	Democracy	Support	Life	Year	Regime	Uyghur	Camp	Give
	T3.?	China	World	Spread	Stop	Communist	Happen	Taiwan	Wuhan	Govt	Ban
	T4.Muslim	Chinese	Make	Muslim	Good	Kill	Police	Terrorist	Bad	Party	Lie
	T5.Human right	World	Freedom	Hong_Kong	Human	Human_right	Time	Free	Stand	hk	Fight
S2	T1.Freedom	World	Stop	Freedom	Truth	Spread	Good	Free	hk	Speech	Life
	T2.Cep	China	Chinese	ccp	Virus	Happen	Wuhan	Evil	Communist	Time	Uyghur
	T3.People	People	Make	Kill	Lie	ppl	Trust	Camp	Police	Thing	Man
	T4.China	China	Country	Regime	Pay	Money	Outbreak	Start	Work	Force	Control
	T5.Human right	Government	Citizen	Human	Fight	Support	Hong_Kong	Taiwan	Give	Democracy	Death
S3	T1.Death	World	People	Pay	Lie	Kill	Truth	Fight	Life	Die	Humanity
	T2.Government	Time	Call	Government	India	Communist	Pandemic	Give	Global	Send	Real
	T3.Virus	Chinese	Virus	Spread	Wuhan	Corona	Product	Buy	Control	Big	Day
	T4.China	China	Country	Make	ccp	Stop	Good	Coronavirus	Human	Trust	Support
	T5.World	China	World	War	Case	Start	Covid	Economy	Death	State	Italy

Table 6 Extracted topics and their corresponding keywords for the category of blame spread across the three stages S1, S2, and S3

S1	T1.Lie	Lie	Spread	Virus	Autocracy	Deceit	Imagine	True	Horrible	Infect	Country
	T2.Death	China	Dead	Die	Day	Order	Monstrosity	True	Thing	Kong	High
	T3.Safety	Coronavirus	Move	Lot	Cvirus	Epicentre	Safety	March	Careful	Knowingly	Health
	T4.Time	Wuhan	Lunar_new	Sick	Year	Time	Absolutely	Medium	Mutate	Emperor	Truth
	T5.Infection	People	Chinese	Make	Online	Pandemic	Catch	Number	Infect	Community	Official
S2	T1.Government	Lie	Chinese	Coronavirus	Government	Wuhan	Cover	Day	Body	Thing	Care
	T2.Spread	World	Country	Spread	Happen	Trust	Kill	Threat	Steal	Dead	Face
	T3.China	China	Truth	Bad	Free	Money	Communist	Case	Find	Start	Move
	T4.Virus	Virus	Stop	Make	Control	Good	China	Fight	Live	Report	Human
	T5.Death	People	Time	Number	Die	Real	life	Entire	Back	Citizen	Death
S3	T1.World	World	China	Country	Pay	Pandemic	Kill	Global	Economy	War	
	T2.?	People	Stop	Human	American	Eat	Put	President	Market	Happen	Live
	T3.Lie	China	Lie	Coronavirus	Wuhan	blame	Die	Case	Cover	Truth	Number
	T4.?	Make	Time	China	Good	Start	Buy	Trust	Back	Thing	Country
	T5.Government	Chinese	Virus	China	Government	Call	Communist	ccp	Covid	Spread	Hold

Table 7 Extracted topics and their corresponding keywords for the category of exclusion spread across the three stages S1, S2, and S3

S1	T1.Government	support	gov	join	people	evil	time	stand	sanction	Government	money
	T2.Human right	product	world	stop	human_right	freedom	tag	good	challenge	ppl	economic_infiltration
	T3.Boycott	china	Hong_Kong	fight	regime	boycott	show	international	control	trust	communist
	T4.Trade	make	buy	ccp	day	thing	friend	Taiwan	Japan	hope	today
	T5.Virus	Country	Chinese	people	spread	Year	human	animal	protect	virus	eat
S2	T1.Nation	people	Chinese	animal	happen	Government	initiative	nation	show	economy	law
	T2.Virus	virus	control	truth	support	live	kill	boycott	start	stand	cover
	T3.Threat	China	time	lie	threat	company	trust	big	entire	spy	Wuhan
	T4.Human right	world	Country	freedom	spread	human_right	economic	thing	evil	steal	raise
	T5.Trade	make	product	stop	buy	day	China	good	ccp	challenge	coronavirus
S3	T1.Virus	China	virus	world	pay	spread	ccp	covid	corona	market	call
	T2.Pandemic	world	China	company	communist	coronavirus	pandemic	global	nation	trust	war
	T3.Trade	Chinese	make	product	buy	boycott	stop	good	economy	Indian	
	T4.Human right	people	lie	Government	human	life	back	animal	eat	bring	
	T5.China	China	Country	time	start	business	give	thing	app	sell	money

more attention to the world situation. In addition, “China” and “Chinese” were separated and became two main topics of stigmatization. To specify, the stigmatization around “China” included sub-topics such as “wuhan”, “quarantine”, “dead”, which reflected the attention drawn to the status of Wuhan. The stigmatization around “Chinese” was more likely to focus on the government, with sub-topics such as “mask”, “government”, “citizen”, and “news”.

Main stigmatization topics at the third stage included “government”, “China”, “Chinese”, “US”, and one topic that could not be identified due to its irrelevant sub-topics. Notably, “government” has become an independent topic of stigmatization at this stage. Especially, under the theme of “government”, “communist” and “ccp” were co-existed with “lie”. In addition, “US” emerged as a new topic, and Trump who was the president of the USA during that time became one critical sub-topic under “US”. This might happen after Trump’s twitter that referred to COVID-19 as Chinese virus. Also, it is worth noting that the conflation between virus and a race/ethnicity contributed to a rapid growth of stigmatization-oriented racist opinions from stage 2 to stage 3 (from 5687 to 107174).

4.2 Offensiveness

At the first and second stages, offensiveness was the second-most-frequently-mentioned theme of racist tweets (S1: 1722; S2: 1808, see Table 3). To specify, according to Table 5, at the first stage, the main topics of offensiveness included “government”, “muslim”, and “human right”, in addition to two unidentified themes. Under “government”, we discovered that some sub-topics were not directly related to COVID-19. Instead, they (e.g. “uyghur” and “camp”) were more likely to target the sensitive internal affairs of China. Similarly, the topics of “muslim” and “human right” also tended to emphasize the China’s internal affairs that had been heatedly discussed before COVID-19. While the offensive sub-topics under Muslim included “kill”, “police”, “terrorist”, “bad”, “party”, and “lie”, the discussion on “human right” centred on “freedom” and “hongkong”.

Offensive language at the second stage still targeted China’s internal affairs and political system. The main topics included “freedom”, “ccp”, “people”, “China”, and “human right”. First, “freedom” emerged as a new topic at this stage. Second, the political attack became more specified, transferring from the “government” to “ccp”. And the “ccp” related discussion still focused on “uyghur”. But “wuhan” became a new sub-topic under “ccp”. “Human right” was still a major topic. However, under “human right”, besides “hongkong”, “taiwan” became a new sub-topic.

At the third stage, the main topics of offensiveness changed to “death”, “government”, “virus”, “China”, and

“world”. Under all topics, “uyghur”, “hongkong”, and “taiwan” were out of the picture. This indicated a shifting focus of offensiveness that started to shed more illumination on the virus rather than the political debates around China’s internal affairs. This theme shift was accompanied by the reduced attention to offensive expression. According to Table 3 (S3: 6973), offensiveness became the least important theme of racist tweets. In general, we find that when COVID-19 was reported to be discovered in China, offensiveness was largely deployed to raise hatred by relating this virus with China’s internal affairs. However, alongside the global outbreak of COVID-19, less attention was drawn to these internal affairs. In the meanwhile, fewer opinions were expressed in an offensive way.

4.3 Blame

According to Table 3, tweets for blaming grew rapidly across the stages of COVID-19 (S1:31; S2 777; S3: 38957). Especially, at the first two stages, blame only occupied the smallest number of racist tweets. However, at the third stage, blame became the second leading theme following stigmatization. To specify, according to Table 6, at the first stage, the main topics to “blame” included “lie”, “death”, “safety”, “time”, and “infection”. We found that blaming reactions at this stage tended to target the negative consequence of COVID-19. For instance, as noted, “stigmatization” tended to associate lie with the political system. However, “blame” was more likely to focus on the consequence of “lie” such as “spread”, “deceit”, “horrible”, and “infect”. This focus can also be easily detected from the rest four topics - “death”, “safety”, “time”, and “infection” that are variously related to the threats brought about by COVID-19 to people’s health and safety.

The second stage of “blame” involved “government”, “spread”, “china”, “virus”, and “death” as the main topics. It is important to note that, besides the words describing the negative COVID-19 consequence, “government” and “china” emerged as two new topics. These two new topics indicated that racist reactions tended to increasingly blame COVID-19 on China and its management of COVID-19. Especially, “government” was associated with “lie”, and “china” was associated “truth”. This suggested that there might be an increasing number of tweets blaming China and government for telling lies and regarding the lying behaviour as the major reason resulting in the outbreak of COVID-19.

Main topics at the third stage included “world”, “lie”, “government”, and two unidentified topics. Akin to the second stage, “lie” and “government” indicated that racist tweets still tended to lay the blame on the “lie” of “government”. Notably, at this stage, “world” emerged as a new main topic. In the category of “world”, we found sub-topics

such as “kill”, “global”, “economy”, “war”, “china”, and “pay”, which indicated that “blame” might have been leveraged to emphasize the negative effects brought about by COVID-19 to the world.

In general, across the three stages of COVID-19, the escalation of blaming behaviours was accompanied by an increasingly specified target to blame. In addition, this target was blamed to contribute to COVID-19 as well as its constantly expanded negative influence across the globe. In so doing, as revealed in prior research, the blaming reactions continued to reinforce the processes of “othering” (Bailey and Harindranath 2005) to draw the boundary between different racial groups.

4.4 Exclusion

According to Table 3, exclusion remained as the second least mentioned theme of racist expression on twitter across the three stages (S1: 872; S2: 1341; S3:8080). However, the topics of exclusion kept on changing. To specify, according to Table 7, at the first stage, the main topics include “government”, “human right”, “boycott”, “trade”, and “virus”. Like other themes, exclusion also targeted political system of China government and its management of COVID-19. However, exclusion also included topics such as “boycott” and “trade”. “Boycott” indicated the purpose of exclusion that was expected to lead to a rejection, while “trade” specified exclusion in an economic way.

At the second stage, exclusion topics became “nation”, “virus”, “threat”, “human right”, and “trade”. Amongst the five topics, “nation” and “threat” were new. It is interesting to note that an extended scope of exclusion that has been transferred from “government” to “nation.” Second, “threat”, as a new focus of exclusion, included sub-topics such as “lie”, “trust”, “spy”, which suggested the unexpected threat that may arise from distrust had become an important reason for exclusion.

At the third stage, topics changed to “virus”, “world”, “trade”, “human right”, and “China.” Two new topics included “world” and “China”. “World” suggested an increasingly globalized discussion on exclusion, including sub-topics such as “global”, “nation”, “trust”, and “war”. Opinion expression concerning “China” tended to place the emphasis on the discussion on the business of China, including sub-topics such as “country”, “business”, “app”, “sell”, and “money”.

Different from other categories, exclusion paid more attention to trade and economy. This suggested that the expression of exclusion tended to focus on the economic aspects between China and the world. Additionally, boycott emerged as a new theme especially in the early stage suggesting that early approach of the discussion might focus on

abandoning Chinese manufactured products and putting less reliance on China. Moreover, the reason behind exclusion seems to be the feeling of threat that tends to be originated from the distrust on China government.

5 Discussion and conclusions

Our study makes unique contribution that enriches the scholarly discussion on deviant racist behaviours on social media. First, bridging computational methods with social science theories, we transcend a binary classification of racist tweets and instead, propose a multidimensional model for racism detection, classification, and analysis. This method, echoing the complicated and dynamic nature of racism, maps the evolvement of racist behaviours alongside the development of COVID-19. Furthermore, the multidimensional categorization of racist behaviours also enables the capturing of the diversity of topics - how different focuses of racist tweets fell under different categories, and how the discussion focus of each category kept on changing across time.

This leads to the second contribution that lies in policy implementation. To specify, the nuanced and dynamic understanding of the racist reactions in the context of COVID-19 will enable the policy makers to have a better interpretation of the possible motivations driving the racist reactions. For instance, as our findings revealed, compared to offensiveness, blame, and exclusion, stigmatization was more likely to act as the leading factor triggering the racist behaviours. Another example is that offensive language was normally deployed to attack the internal affairs of China which might be irrelevant to COVID-19. Better knowledge regarding the reasons behind the public racist reactions could lead to the enactment of more effective policies to prevent the escalation of the race-related deviant behaviours and hate speech.

Additionally, the stage-wise analysis contributes to the enactment of intervention policies with a more specified target at different stages of pandemic. For instance, at the third stage, blame became the most rapidly growing theme, while less and less people were interested in using offensive language. Therefore, the intervention policy can change the focus accordingly across the stages for a better detection and monitoring of racist posts on social media platforms.

Lastly but not least, our study contributes to providing insights into the possible route for interdisciplinary research in the domain of social network analysis and mining. Especially, our study points a direction of deploying social science theories to develop the computational methods for big social data analytics. Future research can also consider embracing the social science perspectives to advance the detection and analysis of linguistic features concerning a particular topic.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abderrouaf C, Oussalah M (2019) On online hate speech detection. Effects of negated data construction. In: 2019 IEEE International Conference on Big Data (Big Data), IEEE pp 5595–5602
- Bailey OG, Harindranath R (2005) Racialised ‘othering’. *J Critl Issues* pp 274–286
- BBC (2021) Covid ‘hate crimes’ against asian americans on rise. <https://www.bbc.com/news/world-us-canada-56218684>, Accessed: 2021-05-21
- Blei DM, Griffiths TL, Jordan MI, Tenenbaum JB, et al. (2003) Hierarchical topic models and the nested chinese restaurant process. *NIPS*, 16
- Blei DM, McAuliffe JD (2010) Supervised topic models. arXiv preprint [arXiv:1003.0783](https://arxiv.org/abs/1003.0783)
- Blei DM, Ng AY, Jordan MI (2003) Latent dirichlet allocation. *J Mach Learn Res* 3(Jan):993–1022
- Cohen R, Ruths D (2013) Classifying political orientation on twitter: It’s not easy!. In: Proceedings of the International AAAI Conference on Web and Social Media 7:91–99
- Coombs T, Schmidt L (2000) An empirical analysis of image restoration: Texaco’s racism crisis. *J Public Relat Res* 12(2):163–178
- Deerwester S, Dumais S, Landauer T et al (1990) Indexing by latent semantic analysis. *J Am Soc Inf Sci* 41:391–407
- Devlin J, Chang MW, Lee K, Toutanova K (2018). Bert: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805)
- Fahim M, Gokhale SS (2021) Detecting offensive content on twitter during proud boys riots. In: 2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA), IEEE pp 1582–1587
- Fan L, Yu H, Yin Z (2020) Stigmatization in social media: documenting and analyzing hate speech for covid-19 on twitter. *Proc Assoc Inf Sci Technol* 57(1):e313
- Fang Z, Zhao X, Wei Q, Chen G, Zhang Y, Xing C, Li W, Chen H (2016) Exploring key hackers and cybersecurity threats in chinese hacker communities. In: 2016 IEEE conference on intelligence and security informatics (ISI), IEEE pp 13–18
- Garland J, Ghazi-Zahedi K, Young JG, Hébert-Dufresne L, Galesic M (2020) Countering hate on social media: large scale classification of hate and counter speech. arXiv preprint [arXiv:2006.01974](https://arxiv.org/abs/2006.01974)
- Gebre BG, Zampieri M, Wittenburg P, Heskes T (2013) Improving native language identification with tf-idf weighting. In: Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications, pp 216–223
- Guntuku SC, Buffone A, Jaidka K, Eichstaedt JC, Ungar LH (2019) Understanding and measuring psychological stress using social media. In: Proceedings of the international AAAI conference on web and social media 13:214–225
- He B, Ziems C, Soni S, Ramakrishnan N, Yang D, Kumar S (2021) Racism is a virus: anti-asian hate and counterspeech in social media during the covid-19 crisis. In: Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, pp 90–94
- Hearst MA, Dumais ST, Osuna E, Platt J, Scholkopf B (1998) Support vector machines. *IEEE Intell Syst their Appl* 13(4):18–28
- Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput* 9(8):1735–1780
- Hofmann T (1999) Probabilistic latent semantic indexing. In: Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval, pp 50–57
- Jaidka K, Zhou A, Lelkes Y (2019) Brevity is the soul of twitter: the constraint affordance and political discussion. *J Commun* 69(4):345–372
- Jelodar H, Wang Y, Yuan C, Feng X, Jiang X, Li Y, Zhao L (2019) Latent dirichlet allocation (lda) and topic modeling: models, applications, a survey. *Multimed Tools Appl* 78(11):15169–15211
- Jeshion R (2013) Expressivism and the offensiveness of slurs. *Philos Perspect* 27(1):231–259
- Kapilashrami A, Bhui K (2020) Mental health and COVID-19: Is the virus racist? *Br J Psychiatry* 217(2):405–407
- Keum BT, Miller MJ (2018) Racism on the internet: conceptualization and recommendations for research. *Psychol Violence* 8(6):782
- Li J, Huang G, Fan C, Sun Z, Zhu H (2019) Key word extraction for short text via word2vec, doc2vec, and textrank. *Turk J Electr Eng Comput Sci* 27(3):1794–1805
- Lim M (2017) Freedom to hate: social media, algorithmic enclaves, and the rise of tribal nationalism in indonesia. *Crit Asian Stud* 49(3):411–427
- Lin CX, Zhao B, Mei Q, Han J (2010) Pet: a statistical model for popular events tracking in social communities. In: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining, pp 929–938
- Liu L, Cao Z, Zhao P, Hu PJH, Zeng DD, Luo Y (2022) A deep learning approach for semantic analysis of COVID-19-related stigma on social media. *IEEE Trans Comput Soc Syst*. <https://doi.org/10.1109/TCSS.2022.3145404>
- Lu R, Sheng Y (2020) From fear to hate: How the covid-19 pandemic sparks racial animus in the united states. arXiv preprint [arXiv:2007.01448](https://arxiv.org/abs/2007.01448)
- Matamoros-Fernández A (2017) Platformed racism: the mediation and circulation of an australian race-based controversy on twitter, facebook and youtube. *Inf Commun Soc* 20(6):930–946
- McCallum AK (2002). A machine learning for language toolkit, Mallet
- Mikolov T, Chen K, Corrado G, Dean J (2013) Efficient estimation of word representations in vector space. arXiv preprint [arXiv:1301.3781](https://arxiv.org/abs/1301.3781)
- Miller CT, Kaiser CR (2001) A theoretical perspective on coping with stigma. *J Soc Issues* 57(1):73–92
- Oboler A (2016) Measuring the hate: the state of antisemitism in social media. Online Hate Prevention Inst, Melbourne
- Pei X, Mehta D (2020) # coronavirus or# chinesevirus?!: Understanding the negative sentiment reflected in tweets with racist hashtags across the development of covid-19. arXiv preprint [arXiv:2005.08224](https://arxiv.org/abs/2005.08224)
- Pennington J, Socher R, Manning CD (2014) Glove: Global vectors for word representation. In: Proceedings of the 2014 conference

- on empirical methods in natural language processing (EMNLP), pp. 1532–1543
- Richeson JA (2018) The psychology of racism: an introduction to the special issue. *Curr Dir Psychol Sci* 27:148–9
- Srikanth M, Liu A, Adams-Cohen N, Wang B, Alvarez RM, Anandkumar A Finding social media trolls: dynamic keyword selection methods for rapidly-evolving online debates 2019
- Tahmasbi F, Schild L, Ling C, Blackburn J, Stringhini G, Zhang Y, Zannettou S (2021) “go eat a bat, chang!”: on the emergence of sinophobic behavior on web communities in the face of covid-19. In: *Proceedings of the web conference 2021*:1122–1133
- Wang S, Chen X, Li Y, Luu C, Yan R, Madrisotti F (2021) ‘i’m more afraid of racism than of the virus!’: racism awareness and resistance among chinese migrants and their descendants in france during the covid-19 pandemic. *Euro Soc* 23(sup1):S721–S742
- Zhai Z, Liu B, Xu H, Jia P (2011) Constrained lda for grouping product features in opinion mining. In: *Pacific-Asia conference on knowledge discovery and data mining*, Springer, pp 448–459
- Zhang Y, Jin R, Zhou ZH (2010) Understanding bag-of-words model: a statistical framework. *Int J Mach Learn Cybern* 1(1):43–52
- Ziems C, He B, Soni S, Kumar S (2020) Racism is a virus: Anti-asian hate and counterhate in social media during the covid-19 crisis. arXiv preprint [arXiv:2005.12423](https://arxiv.org/abs/2005.12423)