“Go eat a bat, Chang!”: On the Emergence of Sinophobic Behavior on Web Communities in the Face of COVID-19

Fatemeh Tahmasbi
Binghamton University
ftahmas1@binghamton.edu

Leonard Schild
CISPA
leonard.schild@cispa.saarland

Chen Ling
Boston University
ccling@bu.edu

Jeremy Blackburn
Binghamton University
jblackbu@binghamton.edu

Gianluca Stringhini
Boston University
gian@bu.edu

Yang Zhang
CISPA
zhang@cispa.de

Savvas Zannettou
Max Planck Institute for Informatics
szannett@mpi-inf.mpg.de

ABSTRACT

The outbreak of the COVID-19 pandemic has changed our lives in unprecedented ways. In the face of the projected catastrophic consequences, most countries have enacted social distancing measures in an attempt to limit the spread of the virus. Under these conditions, the Web has become an indispensable medium for information acquisition, communication, and entertainment. At the same time, unfortunately, the Web is being exploited for the dissemination of potentially harmful and disturbing content, such as the spread of conspiracy theories and hateful speech towards specific ethnic groups, in particular towards Chinese people and people of Asian descent since COVID-19 is believed to have originated from China.

In this paper, we make a first attempt to study the emergence of Sinophobic behavior on the Web during the outbreak of the COVID-19 pandemic. We collect two large datasets from Twitter and 4chan’s Politically Incorrect board (/pol/) over a time period of approximately five months and analyze them to investigate whether there is a rise or important differences with regard to the dissemination of Sinophobic content. We find that COVID-19 indeed drives the rise of Sinophobia on the Web and that the dissemination of Sinophobic content is a cross-platform phenomenon: it exists on fringe Web communities like /pol/, and to a lesser extent on mainstream ones like Twitter. Using word embeddings over time, we characterize the evolution of Sinophobic slurs on both Twitter and /pol/. Finally, we find interesting differences in the context in which words related to Chinese people are used on the Web before and after the COVID-19 outbreak: on Twitter we observe a shift towards blaming China for the situation, while on /pol/ we find a shift towards using more (and new) Sinophobic slurs.

CCS CONCEPTS

- **General and reference** → **General conference proceedings;**
- **Networks** → **Social media networks; Online social networks.**

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW ’21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8312-7/21/04.

<https://doi.org/10.1145/3442381.3450024>

KEYWORDS

COVID-19, Sinophobia, Hate Speech, Twitter, 4chan

ACM Reference Format:

Fatemeh Tahmasbi, Leonard Schild, Chen Ling, Jeremy Blackburn, Gianluca Stringhini, Yang Zhang, and Savvas Zannettou. 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 (WWW ’21)*, April 19–23, 2021, Ljubljana, Slovenia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3442381.3450024>

1 INTRODUCTION

The coronavirus disease (COVID-19) caused by the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) is the largest pandemic event of the information age. SARS-CoV-2 is thought to have originated in China, with the presumed ground zero centered around a wet market in the city of Wuhan in the Hubei province [47]. In a few months, SARS-CoV-2 has spread, allegedly from a bat or pangolin, to essentially every country in the world, resulting in over 39M cases and 1.1M deaths as of October 15, 2020 [27]. Humanity has taken unprecedented steps to mitigating the spread of COVID-19, enacting social distancing measures that go against our very nature. While the repercussions of social distancing measures are yet to be fully understood, one thing is certain: the Web has not only proven essential to the approximately normal continuation of daily life, but also as a tool by which to ease the pain of isolation.

Unfortunately, just like the spread of COVID-19 was accelerated in part by international travel enabled by modern technology, the connected nature of the Web has enabled the spread of misinformation [42], conspiracy theories [28], and racist rhetoric [30]. Considering society’s recent struggles with online racism (often leading to violence), and the politically charged environment coinciding with COVID-19’s emergence, there is every reason to believe that a wave of Sinophobia is not just coming, but already upon us.

In this paper, we present an analysis of how online Sinophobia has emerged and evolved as the COVID-19 crisis has unfolded. To do this, we collect and analyze two large-scale datasets obtained from Twitter and 4chan’s Politically Incorrect board (/pol/). Using temporal analysis, word embeddings, and graph analysis, we shed light on the prevalence of Sinophobic behavior on these communities, how this prevalence changes over time as the COVID-19

pandemic unfolds, and more importantly, we investigate whether there are substantial differences in discussions related to Chinese people by comparing the behavior pre- and post- COVID-19 crisis.

Main findings. Among others, we make the following findings:

- (1) We find a rise in discussions related to China and Chinese people on Twitter and 4chan's /pol/ after the outbreak of the COVID-19 pandemic. At the same time, we observe a rise in the use of specific Sinophobic slurs, primarily on /pol/ and to a lesser extent on Twitter. Also, by comparing our findings to real-world events, we find that the increase in these discussions and Sinophobic slurs coincides with real-world events related to the COVID-19 pandemic.
- (2) Using word embeddings, we look into the context of words used in discussions referencing Chinese people finding that various racial slurs are used in these contexts on both Twitter and /pol/. This indicates that Sinophobic behavior is a cross-platform phenomenon existing in both fringe Web communities like /pol/ and mainstream ones like Twitter.
- (3) Analyzing word embeddings over time, we discover new emerging slurs and terms related to Sinophobic behavior, as well as the COVID-19 pandemic. For instance, on /pol/ we observe the emergence of the term "kungflu" after January, 2020, while on Twitter we observe the emergence of the term "asshoe," which aims to make fun of the accent of Chinese people speaking English.
- (4) By comparing our dataset pre- and post-COVID-19 outbreak, we observe shifts in the content posted by users on Twitter and /pol/. On Twitter, we observe a shift towards blaming China and Chinese people, while on /pol/ we observe a shift towards using more and new Sinophobic slurs.

Disclaimer. Note that content posted on the Web communities we study is likely to be considered as highly offensive or racist. Throughout the rest of this paper, we do not censor any language, thus we warn the readers that content presented is likely to be offensive and upsetting.

2 RELATED WORK

Due to its incredible impact to everybody's life in early 2020, the COVID-19 pandemic has already attracted the attention of researchers. In particular, a number of papers have studied how users on social media discussed this emergency. Chen et al. [9] release a dataset of 50M tweets related to the pandemic. Cinelli et al. [11], Singh et al. [40], and Kouzy et al. [31] studied misinformation narratives about COVID-19 on Twitter. Lopez et al. [32] analyzed a multi-language Twitter dataset to understand how people in different countries reacted to policies related to COVID-19.

A number of papers studied racist activity on social networks. Keum and Miller [29] argued that racism on the Internet is pervasive and that users are likely to encounter it. Zimmerman et al. [51] focused on the influence that the anonymity brought by the Internet has increased the likelihood for people to take part in online aggression. Relia et al. [38] found that racist online activity correlates with hate crimes. In other words, users located in areas with higher occurrence of hate crimes are more likely to engage in racism on social media. Yang and Counts [49] studied how users who experienced racism on Reddit self-narrate their experience.

They characterize the different types of racism experienced by users from different demographics, and show that commiseration is the most valued form of social support. Zannettou et al. [50] present a quantitative approach to understand online racism targeting Jews. As part of their analysis, they present a method to quantify the evolution of racist language based on word embeddings, similar to the technique presented in this paper. Hasanuzzaman et al. [21] investigated how demographic traits of Twitter users can act as a predictor of racist activity. By modeling demographic traits as vector embeddings, they found that male and younger users (under 35) are more likely to engage in racism on Twitter. Other work performed quantitative studies to characterize hateful users on social media, analyzing their language and their sentiment [8, 39]. In particular, it focused on discrimination and hate directed against women, for example as part of the Pizzagate conspiracy [7, 10].

3 DATASETS

To study the extent and evolution of Sinophobic behavior on the Web, we collect and analyze two large-scale datasets from Twitter and 4chan's Politically Incorrect board (/pol/).

Twitter. Twitter is a popular mainstream microblog used by millions of users for disseminating information. To obtain data from Twitter, we use the Streaming API [46], which provides a 1% random sample of all tweets made available on the platform. We collect tweets posted between November 1, 2019 and March 22, 2020, and then we filter only the ones posted in English, ultimately collecting 222,212,841 tweets.

4chan's /pol/. 4chan is an anonymous imageboard, which is divided into several sub-communities called *boards*: each board has its own topic of interest and moderation policy. In this work, we focus on the Politically Incorrect board (/pol/) because it is the main board for the discussion of world events. We use the approach of Hine et al. [22] to collect all posts made on /pol/ between November 1, 2019 and March 22, 2020. Overall, we collect 16,808,191 posts.

Remarks. We focus on these two specific Web communities as we believe that they are representative examples of both mainstream and fringe Web communities. That is, Twitter is a popular mainstream community that is used by hundreds of millions of users around the globe, while 4chan's /pol/ is a notorious fringe Web community that is known for the dissemination of hateful or weaponized information [22].

4 TEMPORAL ANALYSIS

We start our analysis by investigating the temporal dynamics of words related to "china" and "chinese" on 4chan's /pol/ and Twitter. Also, we study the prevalence of several racial slurs targeted towards Chinese and Asian people.

Socio-Political Terms. Figure 1 shows the number of occurrences of "china" and "chinese," and the proportion of posts containing these two words on 4chan's /pol/ on a daily base. We also annotate (with vertical lines) real-world events related to the COVID-19 pandemic (see Table 1 for more details). We first observe a sudden increase in the usage of both words around January 23, 2020, the day the Chinese government officially locked down the city of Wuhan

Figure 1: Mentions of the terms “china” and “chinese” on 4chan’s /pol/.

#	Date	Event
1	2019-12-12	President Donald Trump signs an initial trade deal with China [43].
2	2020-01-23	The Chinese government announces a lock-down in Wuhan. [4].
3	2020-01-30	The World Health Organization declares a public health emergency [35].
4	2020-02-23	11 municipalities in Lombardy, Italy are locked down [19].
5	2020-03-09	Italy extends restrictions in the northern region of the country [18].
6	2020-03-16	Donald Trump referred to COVID-19 as “Chinese Virus” on Twitter [34].

Table 1: Major events, annotated on Figure 1– 4.

marking the first large-scale effort in China to combat COVID-19 (annotation 2 in Figure 1). After the Wuhan lock-down, the popularity of “china” and “chinese” declines until late February, right around the time that COVID-19 cases started to appear en masse in Europe. On February 24, 2020 (annotation 4 in Figure 1), 11 municipalities in Lombardy, Italy were put on lock-down in an attempt to slow the explosion of community spread cases, and we start to see the use of “china” and “chinese” slightly rise again. This rate increases dramatically around March 9, 2020 (annotation 5), which is when the Italian government extended the lock-down to the entirety of Italy. The second peak comes around March 16, 2020, when Donald Trump referred COVID-19 as “Chinese Virus” in a tweet. Interestingly, after this event, we observe a peak in activity where around 10% of all posts made on /pol/ are related to China (see Figure 1).

On Twitter (see Figure 2), we see the same high level trend: discussion about “china” and “chinese” has a large up tick when Wuhan is locked down, and then declines until COVID-19 hits Europe. There is one important difference however. The amount of relative discussion on Twitter during the first peak is much lower than the level of discussion once Europe comes into play. This may be due to the fact that discussion on Twitter is more geographically distributed, or that 4chan’s /pol/ is more easily inflamed by conspiracies and racism-related posts. The concept of “social distance” (not to be confused with “social distancing”) refers to how perception of others is shaped, in part due to cultural, national,

Figure 2: Mentions of the terms “china” and “chinese” on Twitter.

ethical, educational, occupational, etc. similarities [1, 6, 13, 45]. A certain degree of social distance is likely an explanation for the gap between the two peaks. Socio-geographic intimacy, as well as a much more similar cultural background, leads to more attention on COVID-19 outbreaks in Europe than the lock-down in Wuhan.

Racial Slurs. Besides “china” and “chinese,” we also analyze the temporal dynamics of Sinophobic racial slurs on 4chan’s /pol/ and Twitter. We pick a set of 8 Sinophobic slurs, including “chink,” “bugland,” “chankoro,” “chinazi,” “gook,” “insectoid,” “bugmen,” and “chingchong” (for the slur “chingchong,” we search for both “chingchong” and “ching chong”). Some of them are well-known racial slurs towards Chinese and Asian people [48], such as “chink,” “chingchong,” and “gook.” Others (e.g., “bugland”) are based on preliminary results of word embeddings that we use to discover racial slurs (see Section 5 for more details).

The results are depicted in Figures 3 and 4 for /pol/ and Twitter, respectively. For /pol/, we observe a general trend that is similar with the trend observed before for the mentions of “china” and “chinese” (see Figure 1). That is, we see two main peaks of activity, around January 23 and March 16, 2020, for most of the slurs. In particular, it is worrisome that the use of most slurs keeps increasing after the event where Donald Trump referred to COVID-19 as “Chinese Virus.” By looking at some posts on /pol/ posted after March 16, 2020, we find several examples of hateful rhetoric that call for violence: e.g., “*I really hope the world gets together to exterminate every last Chink on this planet.*” and “*Seconded, every chink needs to exterminated, the yellow Jew needs to be erased from this world.*” When looking at the popularity of these terms, we find that “chink” is the most popular Sinophobic slur on /pol/ with an order of magnitude more posts compared to other slurs like “gook” and “chingchong.”

On Twitter (see Figure 4), we observe a rise of Sinophobic slurs during December 2019, especially for “chink.” By manually examining the peak, we find that this happens because a tweet containing

Figure 3: Mentions of Sinophobic racial slurs on 4chan's /pol/.

Figure 4: Mentions of Sinophobic racial slurs on Twitter.

the slur “chink” went viral at that time.¹ Apart from this outlier, we observe an increase in the use of Sinophobic slurs after January 23; again the increase in the use of these slurs is lower than the one observed on /pol/ (see Figures 3 and 4). By the end of our dataset, we observe a substantial increase in the use of the slur “chingchong.” Looking at some examples of tweets including this specific term, we find several hateful comments: e.g., “All because Ching Chong had to chow a fucking frog raw for lunch. #coronavirussafety” and “...Fuck you’d just do to me you little shit I’ll break your fucking neck fucking ching chong corona viru...”

Discussion. A common theme among racist rethoric is that of an invading virus. History is rife with examples of diseases being attributed to specific races and nationalities, and there is no reason to believe that COVID-19 would buck this trend; the first identified COVID-19 cases *did* originate in China. However, the world today is much more diverse and connected than it was in the 15th century when Italians dubbed syphilis the “French disease.”

Figures 1 and 2 make it quite clear that 4chan and Twitter are heavily discussing China in relation to COVID-19, and that this discussion accelerated rapidly once the Western world became affected. The upswing is potentially related to the scapegoating

phenomenon [44]. The first cases originated in China, and the Chinese government was the first to take active and serious measures to combat its spread prompting a reasonable degree of discussion. When these measures were ineffective in preventing the spread to the Western world, however, China’s existing association with COVID-19, in particular China’s “failure” to prevent its spread make it a *just* scapegoat [2] in the face of a looming pandemic.

That said, we do see meaningful differences in the use of *slurs* on /pol/ and Twitter. /pol/’s use of slurs tracks with the use of “china” and “chinese” to a worrying degree, but this is much less pronounced on Twitter. This is not entirely unsurprising considering that /pol/ is well known to be a locus of racist ideology, however it is worthwhile discussing some of the theory around *why* it tracks so well. The clearly racist reaction fits the notion of *defensive denial*, which is a common strategy for coping with stress [3, 16, 23, 26, 41]. Essentially, the early stages of COVID-19 were exclusively a *Chinese* problem; “superior” Western society had nothing to worry about, even though experts were warning of a pandemic break-out even before Wuhan was locked down. This conforms with the *scapegoating* theory of clinical psychology, in which members of a group project unwanted self aspects onto another person or group, then attack the scapegoat believing that “this is not me” [12, 17, 37]. Political scientists have argued that scapegoating is a major driver for racism in a number of settings [14, 36].

¹<https://twitter.com/GoAwayWithJae/status/1203860502444748800>, which, is the culmination of outrage against the h3h3 podcast’s opinion of the K-pop music genre.

4chan's /pol/						Twitter					
Word (China)	Sim.	Word (Chinese)	Sim.	Word (Virus)	Sim.	Word (China)	Sim.	Word (Chinese)	Sim.	Word (Virus)	Sim.
chinas	0.773	chines	0.830	coronavirus	0.846	chinese	0.674	taiwanese	0.756	papilloma	0.702
chinkland	0.761	chink	0.818	covid	0.839	destabl	0.669	japane	0.705	spr	0.700
chinese	0.757	chinese	0.791	coronavirus	0.808	ccp	0.662	mainla	0.677	viruse	0.692
ccp	0.748	china	0.757	corona	0.798	uyghur	0.661	turkistani	0.676	mutate	0.681
nk	0.743	taiwanese	0.752	virius	0.783	fipa	0.660	china	0.674	corona	0.678
chyna	0.742	wuhan	0.725	vrius	0.782	renminbi	0.654	breifli	0.671	transmissible	0.670
bioattack	0.736	chinks	0.719	cornovirus	0.781	boycottbeij	0.653	learnchinese	0.671	ebol	0.661
biowepon	0.719	ccp	0.713	wuflu	0.780	fentanylchina	0.653	stillnoinfo	0.668	mononucleosi	0.657
chicomms	0.718	chinese	0.701	cornavirus	0.779	eastturkistan	0.652	vietnamese	0.667	desease	0.655
bugland	0.712	chinamen	0.699	convid	0.778	xinj	0.651	xijingp	0.661	flue	0.652
wuhan	0.711	japanese	0.693	hivs	0.777	xijinp	0.647	manchuria	0.660	wuhanflu	0.651
chinkistan	0.711	korean	0.690	pathogen	0.774	falung	0.647	putonghua	0.660	coronar	0.649
choyna	0.696	chingchong	0.687	supervirus	0.773	governemnt	0.646	cambodian	0.660	nucleotid	0.648
chine	0.682	mainlander	0.685	disease	0.764	xijingp	0.645	hainan	0.658	pesti	0.646
chiniggers	0.682	chinaman	0.680	sars	0.760	chinazi	0.644	pribumi	0.655	chikungunya	0.646
tradewar	0.682	cpc	0.678	viruse	0.759	xinjiang	0.644	kazakh	0.653	conoravirus	0.645
nambawan	0.681	chicom	0.676	biowepon	0.757	ccpchina	0.639	prc	0.651	commens	0.644
koreas	0.680	mainlanders	0.674	asymptomic	0.754	jinpin	0.637	qingpu	0.650	protozoa	0.644
chynah	0.678	chinkland	0.673	sras	0.753	beltandroad	0.630	laotian	0.644	dengue	0.642
chines	0.675	shina	0.671	megavirus	0.750	eastturkestan	0.630	shandong	0.643	antibi	0.642

Table 2: Top 20 most similar words, along with their cosine similarities, to the words “china,” “chinese,” and “virus” obtained from the word2vec models trained for the whole period (November 2019 - March 2020).

5 CONTENT ANALYSIS

To analyze the content, more specifically the *context* of the use of specific words, we train multiple word2vec models [33] for each Web community. In a nutshell, these models map words into a high-dimensional vector space so that words that are used in a similar way are close to each other. To do this, we use the skip-gram model, which is a shallow neural network aiming to predict the context of a specific word. In this work, we train three groups of word2vec models for each of Twitter and /pol/:

- (1) One word2vec model ($\mathcal{W}_A$) trained on all posts made during the period between October 28, 2019 and March 22, 2020. We denote the period by $\mathcal{T}$. This model allows us to study the use of words for the entire duration of our study.
- (2) One distinct word2vec model for each week between October 28, 2019 and March 22, 2020, which we denote by $\mathcal{W}_{t=i}$, $i \in \mathcal{T}$ (i is the i th week in $\mathcal{T}$). These models allow us to study *changes* in the use of words over time.
- (3) One word2vec model trained on historical data for all posts shared between July 1, 2016 and November 1, 2019 ($\mathcal{W}_C$). This model acts as a baseline and allows us to investigate the emergence of new terms during the period of our study.

Exploring the Context of Terms. First, we look into the overall use of words on 4chan’s /pol/ using the word2vec model trained on the period between October 28, 2019 and March 22, 2020 ($\mathcal{W}_A$). The left side of Table 2 reports the top 20 most similar words for the terms “china,” “chinese,” and “virus.” We make several observations. First, we note that there are many derogatory terms for Asian people (and Chinese people in particular), including “chink,” “chinkland” (referring to the land of chinks, i.e., China), and “chiniggers” (an offensive term created by combining “china” and “nigger”). This indicates that /pol/ users use a wide variety of derogatory terms to possibly disseminate hateful ideology towards Chinese and Asian people. Second, by looking at the most similar words of the term “virus,” we find several terms related to the COVID-19 pandemic [47].

This is evident since the four most similar words to the term “virus” are related to COVID-19; specifically, “coronavirus,” “covid,” “coronavirus,” and “corona.” This indicates that the overall use of words in /pol/ is highly affected by the COVID-19 pandemic and this event is likely to cause changes in the use of language by users.

The corresponding results for Twitter are depicted on the right side of Table 2. We first observe multiple political-related terms that are similar to “china” and “chinese,” such as “government” and “ccp” (Chinese Communist Party). Furthermore, we discover some potentially offensive terms like “chinazi.” This indicates that the use of Sinophobic content is not limited to fringe Web communities like 4chan, and it also exists in mainstream Web communities like Twitter. Also, many terms that are similar to “virus” are related to COVID-19, such as “corona” and “coronavirus.”

Visualizing the Use of Language. To better visualize the use of language related to Chinese people, we create graphs depicting words that are similar to “chinese” following the methodology of Zannettou et al. [50]. In a nutshell, we create a graph where nodes are words and an edge between two words exists if their cosine similarity (obtained from the trained word2vec model) is above a pre-defined threshold.² We limit the graph to nodes that are two hops away from a specific word of interest (in this case “chinese”). Then, we perform various tasks for visualizing the graph. First, the graph is laid out in the graph space with an algorithm that takes into account the weights of the edges [25]. That is, words that have high cosine similarity are laid out closer in the graph space. Second, the size of each node is relative to the number of occurrences of the word in our dataset. Third, we run the Louvain community detection method [5] on the graph and represent nodes that belong to the same community with the same color. The resulting graphs are depicted in Figure 5 and 6 for /pol/ and Twitter, respectively.

²The threshold differs for each resulted graph in a way that it maximizes the readability of the graph.

Figure 5: Visualization of a 2-hop graph from the word “chinese” on 4chan’s /pol/.

By inspecting the obtained communities of words in Figure 5, we observe several interesting themes around the use of words related to “chinese.” First, we discover a community that is highly related to the COVID-19 pandemic (blue community on bottom right). Interestingly, within this community, we also observe terms like “biowepon” (sic) and “bioattack,” likely indicating that /pol/ users are sharing probably false information about the pandemic, for instance claiming that the whole pandemic is a “bioattack” from the Chinese on the Western world. For example, a /pol/ user posted: *“Anyone that doesn’t realize this is a Chinese biowepon by now is either a brainlet or a chicom noodle nigger.”* Second, we observe two tightly-knit communities (red and yellow communities on left-side of the graph) that appear to predominantly include derogatory terms towards Asian, and in particular Chinese people. Some of the words in these communities are “ricenigger,” “chinksect,” “chankoro,” “chinks,” “yellowniggers,” and “pindick.” Interestingly, the most distant word in these communities is the word “batsoup,” which is closer to the community related to COVID-19 [42]. The rest of the communities in this graph are seemingly related to China in general (purple community) and to other countries in Asia (green community). Overall, this graph highlights that /pol/ users use a wide variety of derogatory terms to characterize Chinese people.

When looking at the graph obtained for Twitter (see Figure 6), we observe an interesting community of terms (blue), which includes words related to the COVID-19 pandemic. We observe a large number of words that are seemingly anti-China like the terms “makechinapay,” “blamechina,” and “chinoisasshoe.” At the same time, there are a lot of terms referring to the virus itself like “chinawuhanvirus,” “chinaflu,” and “coronacontrol,” as well as a few terms that aim to support Chinese people through this crisis like “staystrongchina.” For example, a Twitter user posted: *“How do you*

say “*Chi-com asshoe*”? #ChinesePropaganda #ChinaLiedPeopleDied #ChinaVirus #WuhanCoronavirus.” The other communities on the graph include various terms related to happenings in China and other Asian countries/regions.

Discussion. By taking a deeper look at profanities that appeared among the terms, we can roughly divide them into two groups: one is insults addressing Asian people, such as racist variations of “china” and “chinese” (e.g., “chinkland,” “chingchong,” and “chinksect”) or culturally oriented racist terms, including attacking dietary habits (e.g., “ricenigger”), skin tone (e.g., “yellownigger”), or sexual stereotypes (e.g., “pindick”). The frequent appearance of swear words among the terms can indicate an abreaction to the rising fear and stress in front of the disease [15, 20]. At the same time, the racist and targeted focus of these slurs can be explained with the mechanism of *defensive aggression*, either focusing on cultural taboos, such as sexuality [15, 24], or perpetuating societal oppression [24].

6 CONTENT EVOLUTION

Discussions on Web communities like 4chan’s /pol/ and Twitter are highly dynamic and respond to real-world events. For COVID-19, we expect the discussions of users on these Web communities to be dynamic as well as the pandemic unfolds over time. In previous sections, we have explored both the usage of some key terms related to Sinophobia, as well as a static understanding of content. However, these previous analyses do not help us understand how Sinophobic language is evolving over time. More specifically, we lack an understanding on how the context in which words are used changes, and also how new words are created. The former is important because it provides significant insights into the scope and breadth of the problem. The latter is important because the

Figure 6: Visualization of a 2-hop graph from the word “chinese” on Twitter.

language of online extremism has been shown to include memes and slang that have completely contradictory meanings to “normal” usage, or do not even exist outside of the communities that use them, yet encode violent rhetoric that has played a role in acts or mass violence. To this end, we first study the Sinophobic language evolution on 4chan’s /pol/, and in Section 6.2, we focus on Twitter.

6.1 Evolution on 4chan's /pol/

To study the evolution of discussions and use of language, we make use of the weekly word2vec models ($\mathcal{W}_{t=i}$, $i \in \mathcal{T}$) introduced in Section 5. To illustrate how these models are helpful, we initially compare the results from the model trained on the first week of our dataset ($\mathcal{W}_{t=0}$) with the model trained on the last week of our dataset ($\mathcal{W}_{t=-1}$). Table 3 reports the top 20 similar words to “china,” “chinese,” and “virus” for the first and last weekly word2vec models.

We make several key observations. First, when looking at the most similar words to the term “china” from the first weekly model (left side of Table 3), we observe words referring to other counties, mostly in Asia (e.g., “japan,” “singapore,” etc.), but also that the derogatory term “chinks” is among the top 20. This result indicates that 4chan’s /pol/ users typically use racial slurs targeted to Chinese people, and this was also happening even before the outbreak of the COVID-19 pandemic. Similar findings can be observed by looking at the most similar words to the term “chinese.” When looking at

the most similar words to the term “virus,” before the COVID-19 pandemic, we observe general terms related to diseases or other outbreaks, e.g., “ebola.”

Second, by comparing the most similar words from the first and last models, we observe that derogatory terms like “chink” have a higher cosine similarity to “china” in the last model, likely indicating a rise in the use of this term in discussions related to China. Furthermore, we observe terms like “chernobyl,” which may indicate that /pol/ users are comparing this outbreak with the Chernobyl disaster. For example, a /pol/ user posted: *“I can see China collapsing after all this, just as the Chernobyl incident was the beginning of the end for the USSR...”* We also see the term “childkiller,” which upon manual investigation is due to a particularly active user repeatedly posting that China created COVID-19 as a bioweapon. Specifically, we find multiple occurrences of the following sentence in multiple /pol/ posts: *“CHINA CREATED THE CHINA BIOWEAPON MURDER DEATH CHILDKILLER VIRUS IN CHINA!”* Interestingly, we also find some terms that seem to be sarcastic towards the way that Chinese people talk English. For instance, the term “numba” refers to the word “number” and “asshoe” refers to the term “asshole.” Some examples from /pol/ posts are: *“Don’t trust China, China is asshoe”* and *“TAIWAN NUMBA 1 CHINA NUMBA NONE!”*

Third, by looking into the most similar words to the term “chinese” on the last weekly model, we observe the term “bioterrorism”

First Word2vec model (week ending on 2019/11/03)						Last Word2vec model (week ending on 2020/03/22)					
Word (china)	Sim.	Word (chinese)	Sim.	Word (virus)	Sim.	Word (china)	Sim.	Word (chinese)	Sim.	Word (virus)	Sim.
japan	0.779	han	0.740	infectious	0.700	ccp	0.738	ccp	0.744	bioengineered	0.726
singapore	0.737	china	0.733	viruses	0.682	chinese	0.730	chink	0.730	wuflu	0.722
chinese	0.733	tibetans	0.656	infects	0.680	chinas	0.711	china	0.730	decease	0.714
ccp	0.730	chinks	0.641	inject	0.665	chinks	0.694	chinks	0.662	covid	0.712
taiwan	0.711	japan	0.633	pathogen	0.652	chink	0.662	ebright	0.635	supervirus	0.698
ruusia	0.708	singapore	0.630	ebola	0.650	chinkland	0.651	transparently	0.631	viruses	0.695
india	0.703	taiwanese	0.627	lice	0.630	whistleblowers	0.641	taiwanese	0.631	specimens	0.689
venezuela	0.697	chink	0.622	vectors	0.623	chernobyl	0.635	amerikkans	0.627	mutated	0.684
surpass	0.691	filipinos	0.618	spreads	0.618	childkiller	0.634	chinkoid	0.622	corvid	0.681
korea	0.670	payback	0.615	malignant	0.614	embargo	0.626	mainlanders	0.620	virulence	0.675
opium	0.660	ccp	0.613	outbreak	0.613	chyna	0.626	zainichi	0.619	inoculated	0.674
mainland	0.654	mainland	0.611	disposed	0.611	nk	0.624	bioterrorism	0.618	corona	0.674
geostrategic	0.654	cantonese	0.611	deficiency	0.611	chankoro	0.623	labelling	0.613	disease	0.673
surpassed	0.653	koreans	0.611	ensues	0.607	sanctions	0.619	cccp	0.612	chimera	0.671
steamrolled	0.650	hui	0.608	drawings	0.604	numba	0.619	originates	0.611	flu	0.670
hk	0.649	manchuria	0.607	carnage	0.603	asshoe	0.616	ideaia	0.608	transmittable	0.669
indonesia	0.648	paramount	0.604	disposition	0.603	retaliate	0.613	spies	0.606	infection	0.668
pakistan	0.647	mandarin	0.603	bioengineering	0.602	foothold	0.612	chines	0.605	transmits	0.667
manchuria	0.645	japs	0.602	abduction	0.596	vellly	0.609	wumao	0.605	chickenpox	0.666
chinks	0.644	dravidians	0.598	contagion	0.596	engrish	0.606	westerner	0.604	reinfecting	0.665

Table 3: Top 20 most similar words, along with their cosine similarities, to the words “china,” “chinese,” and “virus” for the first and last trained word2vec models from 4chan’s /pol/.

likely indicating that 4chan’s /pol/ users are calling Chinese people bioterrorists. This is related to conspiracy theories that COVID-19 was bioengineered. For example, a /pol/ user posted: “*THIS IS BIOTERRORISM NUKE CHINA NOW.*” By looking at the most similar words to the term “virus,” we find that the most similar one is the term “bioengineered,” indicating that the conspiracy theory went viral on /pol/ during that specific week and was discussed extensively. Finally, by looking at similar words to “virus” in the last weekly model, we clearly observe terms related to the COVID-19 pandemic, e.g., “wuflu” (created by combining Wuhan and Flu), “covid,” and “corona.”

These differences are made more evident by looking at the graph visualizations of the first and last weekly models for the term “chinese” in Figure 7 (we adopt the same methodology used for Figure 5). On the graph for the first model (see Figure 7(a)), we observe mostly innocuous terms related to Chinese people and other Asian people. By looking at the graph obtained from the last model (see Figure 7(b)), however, we observe an entirely different and more hateful behavior. Specifically, the two main tightly-knit communities (red and blue communities), are filled with slurs used against Chinese people like “ricenigger,” “fuckface,” “zipperhead,” “bugpeople,” “subhumans,” etc.

Our findings indicate that we are experiencing an explosion in the use of Chinese derogatory terms in fringe Web communities after the outbreak of the COVID-19 pandemic. These findings are particularly worrisome, since as the pandemic evolves, it is likely to have further rise in the dissemination of racist and hateful ideology towards Chinese people that might also have real-world consequences, such as physical violence against Chinese people.

Discovering new terms. Next, we aim to study how new terms, related to “chinese,” emerge on 4chan’s /pol/ and how their popularity changes over the course of our dataset. To achieve this, we make use of the terms extracted from the vocabularies of the trained word2vec models on 4chan’s /pol/. Specifically, we initially extract the vocabulary from the model trained on historical data ($\mathcal{W}_C$) and

treat it as our base vocabulary. Then, for each weekly trained model ($\mathcal{W}_{t=i}$, $i \in \mathcal{T}$), we extract the vocabulary and compare the terms with our base vocabulary: for each term that is new, we add it to our base vocabulary and treat it as a *new term*. Since we want to find new terms that are related to Chinese, we filter out all new terms that have a cosine similarity below 0.5 in the weekly trained model for which we discovered the new term. Overall, using the above methodology, we manage to discover a total of 50 new terms. Then, we visualize the popularity of the 20 most popular new terms of the course of our dataset in Figure 8.

We observe the emergence of several interesting words during the end of January, 2020. First, we see the rise of terms like “batsoup,” likely indicating that /pol/ users are discussing the fact that the COVID-19 outbreak allegedly started by Chinese people consuming bats. Second, at around the same time, we observe the emergence of “biolab” and “biowarfare.” The use of these words indicates that /pol/ users discuss various conspiracy theories on how the COVID-19 virus was created in a lab or how it can be used as a bioweapon. Interestingly, these terms are persistent from their emergence till the end of our datasets, indicating that these theories are generally appealing to 4chan’s userbase. Other interesting new terms include the terms “kungflu,” which an offensive term towards Chinese people related to the COVID-19 virus, and “heinsberg,” which is the center of the outbreak in Germany and indicates that /pol/ users was discussing about it, especially from the end of February to the beginning of March, 2020.

6.2 Evolution on Twitter

Now, we focus on the Sinophobic language evolution on Twitter. We follow the same methodology used in Section 6.1. The corresponding results are depicted in Table 4.

From Table 4, we can observe that during the first week covered by our Twitter dataset many similar terms to “china” and “chinese” are related to politics, e.g., “tradewar.” This is again quite different from the result on /pol/ (see Table 3). Meanwhile, for “virus” the

Figure 7: Visualization of a 2-hop graph from the word “chinese” on 4chan’s /pol/ using the first and last weekly word2vec models.

most similar terms are also related to diseases. However, when checking the last weekly Twitter model, we observe that many Sinophobic terms appear to be semantically similar to “china” and “chinese,” such as “chinazi.” As in 4chan’s /pol/ (see Table 3), newly created Sinophobic terms, including “chinavirus” and “kungflu,” appear to be close in context as well. For example, a Twitter user posted: *“I agree. Too specific. It’s obviously called the kungflu. It’s kicking all of our asses regardless of denomination.”* Moreover, many terms with similar contexts to “china” and “chinese” in the last weekly Twitter model are still about politics. In contrast to the first

weekly Twitter model, these political-related terms are related to COVID-19, e.g., “cpvirus,” and some of these terms even convey the meaning of revenge and punishment towards China, such as “boycottchina.” For instance, a Twitter user posted: “#ChineseVirus is chinesevirus. One name. #BoycottChina #ChinaLiesPeopleDie.”

By comparing the graphs obtained from the first and last weekly trained word2vec models (we remove these figures from the manuscript due to space constraints), we find substantial differences. The graph from the first model includes mainly words related to China and other Asian regions, as well as words used for discussing

First word2vec model (week ending on 2019/11/03)						Last word2vec model (week ending on 2020/03/22)					
Word (china)	Sim.	Word (chinese)	Sim.	Word (virus)	Sim.	Word (china)	Sim.	Word (chinese)	Sim.	Word (virus)	Sim.
xinjiang	0.735	quidpr	0.656	infect	0.787	ccp	0.766	desensit	0.739	corona	0.761
turkie	0.702	taiwanese	0.652	antibodi	0.773	wuhan	0.759	chinarvirus	0.735	viru	0.757
tibet	0.693	mainland	0.615	malnutrit	0.756	wuhanvirus	0.758	scapegoa	0.729	vir	0.716
tradewar	0.666	cantonese	0.613	mutat	0.751	chinarvirus	0.739	spokespeople	0.721	viruse	0.702
export	0.658	xijinp	0.609	measl	0.743	wuhancoronavirus	0.732	communist	0.718	coronavirus	0.695
dcep	0.658	xinjiang	0.608	pathogen	0.738	prc	0.731	ccp	0.713	corana	0.698
xijinp	0.655	sabahan	0.602	outbreak	0.731	ccpvirus	0.726	disinf	0.712	covd	0.696
chongq	0.651	china	0.602	inflamm	0.729	chinesevirus	0.726	wuhanvirus	0.711	caronavirus	0.690
sprat	0.642	uyghur	0.593	diseas	0.727	culpubl	0.722	incompe	0.705	carona	0.685
eros	0.638	danish	0.591	infecti	0.726	ashhoe	0.714	vrisu	0.700	coronaviri	0.685
whereshunt	0.630	dcep	0.590	immune	0.720	silkroad	0.711	chicom	0.699	nipah	0.685
sichuan	0.630	tibetan	0.587	urinari	0.714	kne	0.711	wuhancoronavirus	0.698	desensit	0.687
eastturkistan	0.629	bytedance	0.584	hvp	0.714	wuhanflu	0.705	reflexive	0.698	biochem	0.681
pork	0.627	cpec	0.575	influenza	0.713	boycottchina	0.705	chinesevirus	0.697	syphili	0.678
muslim	0.623	counterpart	0.574	vaccine	0.712	madeinchina	0.704	wumao	0.695	flue	0.676
munit	0.621	refut	0.573	bacteria	0.708	communist	0.703	prc	0.692	wuflu	0.675
mainland	0.621	cultu	0.571	cardiovascular	0.708	chinali	0.701	wuhanflu	0.686	wuhancoronavirus	0.673
shandong	0.619	tibet	0.570	cannabinoid	0.707	chinazi	0.697	kungflu	0.683	distrac	0.672
cpec	0.614	wechat	0.560	vaccin	0.703	wumao	0.694	volunte	0.682	monoronavirus	0.671
communist	0.610	warship	0.560	ebola	0.701	chinaissashoe	0.693	china	0.682	crono	0.669

Table 4: Top 20 most similar words, along with their cosine similarities, to the words “china,” “chinese,” and “virus” for the first and last trained word2vec models on Twitter.

Figure 8: Visualization of the emergence of new words related to “chinese” over time on 4chan’s /pol/.

matters related to China, e.g., “tradewar.” On the other hand, for the graph obtained from the last model, we find several terms related to COVID-19 like “chinavirus,” “chinesevirus,” “chineseflu,” and “chinoisasshoe.” This indicates a shift towards the use of racist terms related to Chinese people after the COVID-19 outbreak on Twitter. We also find some terms that appear related to the behavior of Donald Trump. For instance, we find the term “racistinchief” that is likely related to the fact that Donald Trump calls the COVID-19 virus the “Chinese Virus,” and this was discussed on Twitter. For instance, a Twitter user posted: “*Trump’s a real asshole, just in case y’all forget #TrumpPandemic #TrumpVirus #RacistInChief.*”

Discovering new terms. To discover new terms from Twitter, we follow the same methodology with /pol/, as documented in Section 6.1. Overall, we discover a total of 713 new terms between October 28, 2019 and March 22, 2020. Figure 9 visualizes a sample of 40 of the new terms according to their popularity and cosine similarity with the term “chinese.” We observe a lot of new terms relating to the Hong Kong protests emerging during November 2019, such as “freehongk” and “hongkongpoliceb.” Also, after the outbreak of the COVID-19 pandemic we observe the emergence of a wide variety of terms around the end of January 2020. Some notable examples include terms like “chinavirus,” “chinesevirus,”

Figure 9: Visualization of the emergence of new words related to “chinese” over time on Twitter.

“wuhanpneumonia,” “wuhancorono,” etc. These findings highlight that during important real-world events, such as the COVID-19 pandemic, language evolves and new terms emerge on Web communities like Twitter. At the same time, it is particularly worrisome that we observe the appearance of new terms in mainstream social media that can be regarded as Sinophobic like “chinesevirus,” which can possibly lead to hate attacks in the real-world, and almost certainly harm international relations.

Figure 10: Cosine similarities between various terms over time.

6.3 Semantic Changes between Words

As the last part of our analysis, we set out to assess how the semantic distance between words change over the course of our datasets. To do this, we use the weekly trained word2vec models ($W_{t=i}, i \in \mathcal{T}$): for each word2vec model, we extract the cosine similarity between two terms and then we plot their similarities over time. This allow us to understand whether two terms appear more often in a similar context as COVID-19 unfolds. We show some examples in Figure 10; the terms are selected based on our previous analysis.

We indeed observe several interesting changes in the similarities between terms. For the terms “chinese” and “virus” (see Figure 10(a)), we see a substantial increase in cosine similarity between these two terms over time, especially after the week ending on January 19, 2020. The cosine similarity on both Twitter and /pol/ is below 0.5 in the early models, while after January 19, 2020, it is mostly over 0.5, with the last model having a similarity over 0.6. This indicates that the terms “chinese” and “virus” are used in more similar ways over time on both Twitter and /pol/.

Another example is the terms “chinese” and “chink” (see Figure 10(b)). We observe that for both Twitter and /pol/, the similarity

between these terms increases over the course of our datasets. Interestingly, the increase in cosine similarity between these terms is *larger* for Twitter, likely indicating that Twitter users are more affected by COVID-19 with regards to sharing Sinophobic content, while on /pol/ the difference is smaller which indicates that /pol/ users were affected less by COVID-19 when it comes to sharing Sinophobic content. I.e., /pol/ started and ended using “Chinese” in the same context as the racial slur “chink,” however, Twitter exhibited a much larger swing towards racist connotations. Finally, we illustrate also the cosine similarity differences between the terms “chinese” and “bat” in Figure 10(c). We observe that the cosine similarity is low from our first models and substantially increases after the week ending on January 26, 2020. This indicates that both Twitter and /pol/ users have started discussing the fact that the virus allegedly originates from “bats” around that specific time frame and they continue doing so until through the end of our datasets.

7 CONCLUSION

To combat the COVID-19 pandemic, many governments have implemented unprecedented measures like social distancing and even government enforced, large-scale quarantines. This has resulted in the Web becoming an even *more* essential source of information, communication, and socialization. Unfortunately, the Web is also exploited for disseminating disturbing and harmful information, including conspiracy theories and hate speech targeting Chinese people and people of Asian descent.

Part of this can be attributed to scapegoating, a basic psychosocial mechanism to deal with stress. Building upon the well known in-group favoritism/out-group hostility phenomenon, racist ideology has a long history of scapegoating. A common scapegoating theme has been to equate the targeted people with a disease, either figuratively or literally. When threatened by events outside our control, it is only “natural” to seek for external blame. In the case of COVID-19, the entire world is threatened, and there is a “natural” external actor to blame.

In this paper, we make a first attempt to understand Sinophobic language on the social Web related to COVID-19. To this end, we collect two large-scale datasets from 4chan’s /pol/ and Twitter over a period of five months. Our results show that COVID-19 has indeed come with a rise of Sinophobic content mainly on fringe Web communities like /pol/ and to a lesser extent on mainstream ones like Twitter. Relying on word embeddings, we also observe the semantic evolution of slurs. Moreover, our study also shows that many new Sinophobic slurs are created as the crisis progresses.

Our study has several implications for both society and the research community focusing on measuring, understanding, and mitigating emerging social phenomena on the Web. First, we show that the dissemination of hateful content, and in particular Sinophobic content, is a cross-platform phenomenon that incubates both on fringe Web communities as well as mainstream ones. This prompts the need to have a multi-platform point-of-view when studying such emerging social phenomena on the Web. Second, we show that Sinophobic behavior evolves quickly and substantially, especially after world changing events like the COVID-19 pandemic. This highlights the need to develop new techniques and tools to understand these changes in behavior and work towards designing and

deploying counter-measures with the goal to prevent or mitigate real-world violence stemming from these behaviors.

While the COVID-19 crisis does provide a unique opportunity to understand the evolution of hateful language, our study should be also be taken as a call for action. The Web has enabled much of society to keep going, or at least to maintain social connections with other humans, but it has also allowed the proliferation of hateful language, disinformation and conspiracy theories, and calls to violence at a time where we can afford it the least.

Acknowledgments. We thank the anonymous reviewers for their comments. This work was supported by the NSF under Grants 2046590 and 1942610.

REFERENCES

- [1] George A Akerlof. 1997. Social distance and social decisions. *Econometrica: Journal of the Econometric Society* (1997), 1005–1027.
- [2] Gordon Willard Allport, Kenneth Clark, and Thomas Pettigrew. 1954. The nature of prejudice. (1954).
- [3] Kathryn Freeman Anderson. 2013. Diagnosing discrimination: Stress from perceived racism and the mental and physical health effects. *Sociological Inquiry* 83, 1 (2013), 55–81.
- [4] BBC. 2020. China coronavirus: Lockdown measures rise across Hubei province. <https://www.bbc.com/news/world-asia-china-51217455>.
- [5] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* 2008, 10 (2008), P10008.
- [6] Emory Stephen Bogardus. 1933. A social distance scale. *Sociology & Social Research* (1933).
- [7] Despoina Chatzakou, Nicolas Kourtellis, Jeremy Blackburn, Emiliano De Cristofaro, Gianluca Stringhini, and Athena Vakali. 2017. Hate is not binary: Studying abusive behavior of #gamergate on Twitter. In *ACM HyperText*.
- [8] 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 *World Wide Web Companion*.
- [9] Emily Chen, Kristina Lerman, and Emilio Ferrara. 2020. COVID-19: The First Public Coronavirus Twitter Dataset. *arXiv preprint arXiv:2003.07372* (2020).
- [10] Shira Chess and Adrienne Shaw. 2015. A conspiracy of fishes, or, how we learned to stop worrying about #GamerGate and embrace hegemonic masculinity. *Journal of Broadcasting & Electronic Media* (2015).
- [11] Matteo Cinelli, Walter Quattrociocchi, Alessandro Galeazzi, Carlo Michele Valensise, Emanuele Brugnoti, Ana Lucia Schmidt, Paola Zola, Fabiana Zollo, and Antonio Scala. 2020. The covid-19 social media infodemic. *arXiv preprint arXiv:2003.05004* (2020).
- [12] Arthur J Clark. 2002. Scapegoating: Dynamics and interventions in group counseling. *Journal of Counseling & Development* 80, 3 (2002), 271–276.
- [13] Patrick W Corrigan, Annette Backs Edwards, Amy Green, Sarah Lickey Diwan, and David L Penn. 2001. Prejudice, social distance, and familiarity with mental illness. *Schizophrenia bulletin* 27, 2 (2001), 219–225.
- [14] Margaret Denike. 2015. Scapegoat racism and the sacrificial politics of “security”. *Journal of international political theory* 11, 1 (2015), 111–127.
- [15] Jean-Marc Dewaele. 2004. The emotional force of swearwords and taboo words in the speech of multilinguals. *Journal of multilingual and multicultural development* 25, 2-3 (2004), 204–222.
- [16] Bruce P Dohrenwend, Yuval Neria, J Blake Turner, Nicholas Turse, Randall Marshall, Roberto Lewis-Fernandez, and Karestan C Koenen. 2004. Positive tertiary appraisals and posttraumatic stress disorder in US male veterans of the war in Vietnam: the roles of positive affirmation, positive reformulation, and defensive denial. *Journal of consulting and clinical psychology* 72, 3 (2004), 417.
- [17] Gary Gemmill. 1989. The dynamics of scapegoating in small groups. *Small group behavior* 20, 4 (1989), 406–418.
- [18] The Guardian. 2020. Coronavirus: Italy extends emergency measures nationwide. <https://www.bbc.com/news/world-europe-51810673>.
- [19] The Guardian. 2020. Italy imposes draconian rules to stop spread of coronavirus. <https://www.theguardian.com/world/2020/feb/23/italy-draconian-measures-effort-halt-coronavirus-outbreak-spread>.
- [20] Emre Güvendir. 2015. Why are males inclined to use strong swear words more than females? An evolutionary explanation based on male intergroup aggressiveness. *Language Sciences* 50 (2015), 133–139.
- [21] Mohammed Hasanuzzaman, Gaël Dias, and Andy Way. 2017. Demographic word embeddings for racism detection on twitter. (2017).
- [22] Gabriel Emile Hine, Jeremiah Onaolapo, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Riginos Samaras, Gianluca Stringhini, and Jeremy Blackburn. 2017. Kek, Cucks, and God Emperor Trump: A Measurement Study of 4chan’s Politically Incorrect Forum and Its Effects on the Web. In *ICWSM*.
- [23] B Kent Houston. 1973. Viability of coping strategies, denial, and response to stress. *Journal of Personality* (1973).
- [24] Geoffrey Hughes. 1998. *Swearing: A social history of foul language, oaths and profanity in English*. Penguin UK.
- [25] 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 (2014).
- [26] Ronnie Janoff-Bulman and Christine Timko. 1987. Coping with traumatic life events. In *Coping with negative life events*. Springer, 135–159.
- [27] Johns Hopkins. 2020. Johns Hopkins Coronavirus Resource Center. <https://coronavirus.jhu.edu/map.html>.
- [28] Jordan Kelly-Linden and Paul Nuki. 2020. Coronavirus myths, scams and conspiracy theories that have gone viral. <https://www.telegraph.co.uk/global-health/climate-and-people/coronavirus-myths-scams-conspiracy-theories-true-false-lab-leak/>.
- [29] Brian TaeHyuk Keum and Matthew J Miller. 2018. Racism on the Internet: Conceptualization and recommendations for research. *Psychology of violence* 8, 6 (2018), 782.
- [30] Catherine Kim. 2020. “They just see that you’re Asian and you are horrible”: How the pandemic is triggering racist attacks. <https://www.vox.com/identities/2020/3/25/21190655/trump-coronavirus-racist-asian-americans>.
- [31] 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).
- [32] Christian E Lopez, Malolan Vasu, and Caleb Gallemore. 2020. Understanding the perception of COVID-19 policies by mining a multilanguage Twitter dataset. *arXiv preprint arXiv:2003.10359* (2020).
- [33] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [34] NBCnews. 2020. Trump tweets about coronavirus using term ‘Chinese Virus’. <https://nbcnews.to/34E2Ge6>.
- [35] World Health Organization. 2020. Statement on the second meeting of the International Health Regulations (2005) Emergency Committee regarding the outbreak of novel coronavirus (2019-nCoV). <https://bit.ly/3luxfJX>.
- [36] Tracey Owens Patton and Julie Snyder-Yuly. 2007. Any four black men will do: Rape, race, and the ultimate scapegoat. *Journal of Black Studies* 37, 6 (2007), 859–895.
- [37] Charles Prudhomme. 1970. Reflections on racism. *American Journal of Psychiatry* 127, 6 (1970), 815–817.
- [38] Kunal Relia, Zhengyi Li, Stephanie H Cook, and Rumi Chunara. 2019. Race, Ethnicity and National Origin-based Discrimination in Social Media and Hate Crimes Across 100 US Cities. In *ICWSM*.
- [39] Manoel Horta Ribeiro, Pedro H Calais, Yuri A Santos, Virgílio AF Almeida, and Wagner Meira Jr. 2018. Characterizing and detecting hateful users on Twitter. In *ICWSM*.
- [40] Lisa Singh, Shweta Bansal, Leticia Bode, Ceren Budak, Guangqing Chi, Kornaphop Kawintiranon, Colton Padden, Rebecca Vanardsall, Emily Vraga, and Yanchen Wang. 2020. A first look at COVID-19 information and misinformation sharing on Twitter. *arXiv preprint arXiv:2003.13907* (2020).
- [41] Selcuk R Sirin, Lauren Rogers-Sirin, Jessica Cressen, Taveeshi Gupta, Sammy F Ahmed, and Alfredo D Novoa. 2015. Discrimination-related stress effects on the development of internalizing symptoms among Latino adolescents. *Child Development* 86, 3 (2015), 709–725.
- [42] Josh Taylor. 2020. Bat soup, dodgy cures and ‘diseaseology’: the spread of coronavirus misinformation. <https://www.theguardian.com/world/2020/jan/31/bat-soup-dodgy-cures-and-diseaseology-the-spread-of-coronavirus-bunkum>.
- [43] TIME. 2019. Trump Signs Off on Trade Deal With China to Avert December Tariffs. <https://time.com/5749191/trump-us-china-trade-deal/>.
- [44] Eugene Toker. 1972. The scapegoat as an essential group phenomenon. *International Journal of Group Psychotherapy* 22, 3 (1972), 320–332.
- [45] Yaacov Trope and Nira Liberman. 2010. Construal-level theory of psychological distance. *Psychological review* 117, 2 (2010), 440.
- [46] Twitter. 2020. Twitter’s Streaming API. <https://developer.twitter.com/en/docs/labs/sampled-stream/overview>.
- [47] Wikipedia. 2019. 2019-20 coronavirus pandemic. https://en.wikipedia.org/wiki/Coronavirus_disease_2019.
- [48] Wikipedia. 2019. Sinophobia. <https://en.wikipedia.org/wiki/Sinophobia>.
- [49] Diyi Yang and Scott Counts. 2018. Understanding self-narration of personally experienced racism on reddit. In *ICWSM*.
- [50] Savvas Zannettou, Joel Finkelstein, Barry Bradlyn, and Jeremy Blackburn. 2020. A Quantitative Approach to Understanding Online Antisemitism. In *ICWSM*.
- [51] Adam G Zimmerman and Gabriel J Ybarra. 2016. Online aggression: The influences of anonymity and social modeling. *Psychology of Popular Media Culture* (2016).