

Studying the Dynamics of COVID-19 Misinformation Themes using Topic Streams

Thomas Marcoux, Esther Mead, Nitin Agarwal
College of Engineering and Information Technology
University of Arkansas at Little Rock
 Little Rock, AR, USA
 {txmarcoux, elmead, nxagarwal}@ualr.edu

Abstract—The COVID-19 pandemic has seen the emergence of unique misinformation narratives in various outlets, through social media, blogs, etc. This online misinformation has been proven to spread in a viral manner and has a direct impact on public safety. In an effort to improve public understanding, we curated a corpus of 543 misinformation pieces whittled down to 243 unique misinformation narratives along with independent international organizations debunking these stories. Building upon previous applications of topic modeling to COVID-19 related material, we developed a tool leveraging topic modeling to create a chronological visualization of these stories. From our corpus of misinformation stories, this tool has shown to accurately represent the ground truth. This highlights some of the misinformation narratives unique to the COVID-19 pandemic and provides a quick method to monitor and assess misinformation diffusion, enabling policy makers (such as the Arkansas Office of the Attorney General - Arkansas, USA) to identify themes to focus on for communication campaigns. To further explore the potential of topic streams in understanding online opinion, we experiment with multiple topic models and also apply our methodology to YouTube data. The principal difference between our effort and other similar efforts by Google and social media companies is that we are paying special attention to cases of misinformation and scammers that are affecting our region, while also including global cases.

Index Terms—misinformation, disinformation, topic models, topic streams, COVID-19, misinfodemic, narratives.

I. INTRODUCTION

Social media is characterized as a powerful online interaction and information exchange medium. However, it has given rise to new forms of deviant behaviors such as spreading fake news, misinformation, and disinformation. Due to afforded anonymity and perceived diminished personal risk of connecting and acting online, deviant groups are becoming increasingly common. Online deviant groups have grown in parallel with Online Social Networks (OSNs), whether it is black hat hackers using Twitter to recruit and arm attackers, announce operational details, coordinate cyber-attacks [10], and post instructional or recruitment videos on YouTube targeting certain demographics; or state/non-state actors and extremist groups (such as the Islamic State of Iraq and Syria) savvy use of social communication platforms to conduct phishing operations, such as viral retweeting of messages containing harmful URLs leading to malware [6].

More recently, there is a surge in misinformation and scam cases pertaining to COVID-19. The problem of misinformation

is actually worse than the pandemic itself. That is why it is called infodemic or more specifically, misinfodemic. Like the pandemic, misinformation cases are also rising exponentially. These cases are more difficult to track than the epidemic, as they can originate in the dark corners of the Internet. To make matters worse, we cannot enforce lockdown on the Internet to stop the spread of this infodemic. This is in part because, during crises, the Internet is usually the first mode of communication and source of information. Although there are some quarantine efforts, for instance from social media companies such as Facebook, YouTube, and retail companies like Amazon are doing their best to block such content, by suspending bad actors or scammers who are spreading misinformation to further their political agenda or to try to profit off of this adversity. But such cases are simply too many and growing too fast. What makes this problem worse is the fact that the information spreads like a wildfire on the Internet, especially the false or misinformation. Many studies have concluded that misinformation travels faster than its corrective information, and the more questionable the misinformation is the faster it travels. This is simply because on social media people usually have a lot more virtual friends than they do in their real life. So, if they share or retweet some misinformation, wittingly or unwittingly, they expose all their virtual friends to the misinformation.

There are similarities between misinformation about COVID-19 and other misinformation cases that we have studied for NATO, US, EU, Singapore, and Canada, etc. Like in other cases, the motivation for spreading COVID-19 misinformation is monetization or to provoke hysteria. Bad actors or scammers are spreading misinformation to further their political agenda or simply trying to profit off of this adversity. For instance, there exists many cases of scammers selling fake masks, fake cures, using fake websites to ask for private/sensitive information from people by posing as government websites. However, there is a significant difference between COVID-19 and other misinformation campaigns that we have studied before. Being a global and rapidly evolving crisis, the nature of misinformation is also extremely diverse and super-fast. Other misinformation campaigns were specific to an entity, event, region, elections, military exercises. However, misinformation about COVID-19 has both global as well as regional narratives. While fake masks, fake cures,

etc. affect a global audience, the regional narratives include promoting medicines for bovine coronavirus as cure for human coronavirus affecting rural/agriculturalist regions. Moreover, the misinformation about COVID-19 ranges from health to policy to religion to geopolitical affairs, i.e., highly topically diverse. Given the volume, velocity, and variety of COVID-19 related misinformation, research is warranted to study such campaigns and their organization. As resources are stretched too thin, government and other regulatory bodies cannot afford to investigate all the misinformation campaigns and scams. Such research could help prioritize investigation of misinformation campaigns and scams.

Therefore, we propose a study of the themes and chronological dynamics of the spreading of misinformation about COVID-19. Our scope focuses on misinformation geographically relevant to us (Arkansas, USA) as well as some global stories, with our main corpus is a collection of unique misinformation stories manually curated by our team. In collaboration with the Arkansas Attorney General, we have shared our findings with their office and made all reports and misinformation stories publicly available online[8]. In addition, we have collected a variety of YouTube video titles and comments. This allows us to compare a curated corpus to a data set more chaotic and true to life. To highlight and visualize these misinformation themes, we use topic modeling, and introduce a tool to visualize the evolution of these themes chronologically.

The rest of this study is structured as follows. First, we will discuss the work done by other researchers in comparable pieces, then describe our methodology, including data collection, processing, and topic modeling methodology. Then in the results section, we will discuss the subjective findings of our misinformation team as well as the scientific topic streams visualizations that support them. Finally, we will briefly discuss our free online resource where our data can be found before presenting our conclusions.

II. LITERATURE REVIEW

The information community has been tackling the issue of misinformation surrounding the COVID-19 pandemic since early in the outbreak. We base the claims found in this paper on the findings that misinformation spreads in a viral fashion and that consumers of misinformation tend to fail at recognizing it as such [13]. In addition to this, we believe this research is essential as rampant misinformation constitutes a danger to public safety [11]. We also believe this research is helpful in curbing misinformation since researchers have found that simply recognizing the existence of misinformation and improving our understanding of it can enhance the larger public's ability to recognize misinformation as such [13]. In order to better understand the misinformation surrounding the pandemic, we look at previous research that has leveraged topic models to understand online discussions surrounding this crisis. Research has shown the benefits of using this technique to understand fluctuating Twitter narratives [17] over time,

and also in understanding the significance of media outlets in health communications [12].

To implement topic modeling, we use the Latent Dirichlet Allocation model (LDA). Within the realm of Natural Language Processing (NLP), topic modeling is a statistical technique designed to categorize a set of documents within a number of abstract "topics"[3]. A "topic" is defined as a set of words outlining a general underlying theme. For each document, which in this case, is an individual item of misinformation in our data set, a probability is assigned that designates its "belongingness" to a certain topic. In this study, we use the popular LDA topic model due to its widespread use and proved performances [4]. One point of debate within the topic modeling community is the elimination of stop-words: i.e., analysts should filter common words from their corpus before training a model. Following recent research claiming that the use of custom stop-words adds few benefits [16], we followed the researchers' recommendation and removed common words **after** the model had been trained.

Our model choice has seen use in previous research using LDA for short texts, specifically for short social media texts such as tweets [1, 7, 19]. Some other social media research using homogeneous social media sources such as tweets or blog posts use associated hashtags to provide further context to topic models [2]. We expand this research on social media corpora by focusing one of the largest information propagator on the web: YouTube.

In this paper, we propose to leverage topic models to understand the main underlying themes of misinformation and their evolution over time using a manually curated corpus of known fake narratives.

As a secondary goal, we observe the performances of different topic models for understanding online discourse. To accomplish this, we repeated our methodology on a secondary data set using a Hierarchical Dirichlet Process (HDP) model [18]. For our purposes, the major difference between the two models is that LDA models require a number of topics prior to training and will actively attempt to fit that number to the corpus, potentially leading to biased results. On the other hand, the HDP model infers the number of topics present in the corpus during training.

III. METHODOLOGY

This study uses a two-step methodology to produce relevant topic streams. First, through a manual curating process, we aggregate different misinformation narratives for later processing. We consider misinformation narratives, any narrative pushed through a variety of outlets (social media, radio, physical mail, etc.) that has been or is later believably disproved by a third party. This corpus constitutes our input data. Secondly, we use this corpus to train an LDA topic model and to generate subsequent topic streams for analysis. We describe these two steps in more details in the next sections.

A. Collection of Misinformation Stories

Initially, the misinformation stories in our data set were obtained from a publicly available database created by EU-

vsDisinfo in March of 2020 [9]. EUvsDisinfo’s database, however, was primarily focused on “pro-Kremlin disinformation efforts on the novel coronavirus”. Most of these items represented false narratives that were communicating political, military, and healthcare conspiracy theories in an attempt to sow confusion, distrust, and public discord. Subsequently, misinformation stories were continually gleaned from publicly available aggregators, such as POLITIFACT, Truth or Fiction, FactCheck.org, POLYGRAPH.info, Snopes, Full Fact, AP Fact Check, Poynter, and Hoax-Slayer. The following data points were collected for each misinformation item: title, summary, debunking date, debunking source, misinformation source(s), theme, and dissemination platform(s). The time period of our data set is from January 22, 2020 to July 22, 2020. The data set is comprised of 543 total stories and 243 unique misinformation narratives. For many of the items, multiple platforms were used to spread the misinformation. For example, oftentimes a misinformation item will be posted on Facebook, Twitter, YouTube, and as an article on a website. For our data set, the top platforms used for spreading misinformation were websites, Facebook, Twitter, YouTube, and Instagram, respectively. All the stories found by our team are made public through our partnership with the Arkansas Attorney General Office and can be found on our website.

B. Collection of YouTube Data

In order to observe results in uncontrolled, relevant social media environments, we also gathered YouTube data. Since there exists many studies that concern themselves with Twitter data, we chose YouTube because it is another principal vector of information and communication between users. Using the official YouTube API, we performed separate searches for the following keywords on April 19th: “Coronavirus, Corona, Virus, COVID19, COVID, Outbreak”. The result is a set of the most popular videos at that time, as determined by YouTube’s algorithm. From this search, we collected a total of 7,727 videos ranging mostly from January 1st to April 19th. For this particular study, in order to focus on the most relevant videos possible, we selected only videos published between March 1st and March 31st (included). This totals 444 videos, which is comparable to the number of narratives studied. For the purposes of this study, we will only look at the video titles. After selecting this corpus, we used the same API to collect comments posted in these videos and gathered a total of 652,120 comments. In order to comply with YouTube’s terms of service, this data cannot be made public.

C. Topic Modeling

In order to derive lexical meaning from this corpus, we built a pipeline executing the following steps. First, we processed each document in our text corpus. All that is needed is a text field identified by a date. Because in most cases of word of mouth or social media it is impossible to pinpoint the exact date the idea first emerged, we use the date of publication of the corresponding third party “debunk piece”. We trained our LDA model using the Python tool Gensim, with the

methodology and pre-processing best practices as described by its author [14] as well as best stop words practices as described earlier [16]. In this study, we found that generating 20 different topics best matched the ground truth as reported by the researchers curating the misinformation stories.

Still using Gensim, we also trained an alternative topic model using hierarchical Dirichlet process (HDP) [18]. The process is the same except for the number of topics. HDP infers the number of topics in a corpus (with a default threshold of 150). Therefore, we only select the first 20 topics, ordered by α , the weight of each document to topic distribution.

Once the models have been trained, we ordered the documents by date and created a numpy matrix where each document is given a score for each topic produced by the model. This score describes the probability that the given document is categorized as being part of a topic, i.e., if a probability score is high enough (more details below), the document is considered to be part of the topic. Through manual observations we noticed that many documents retain “noise probability”, giving them a probability to be in every topic of around 1% to 5%. For this reason, we set the probability threshold to a comfortable 10% and noticed consistent results. This allowed us to leverage the Python Pandas library to plot a chronological graph for each individual topic. We averaged topic distribution per day and used a moving average window size of 20 unless otherwise specified. This helped in highlighting the overarching patterns of the different narratives. Note, however, that this process hides some early and late data in our set as there are less data points around that time.

IV. RESULTS

In this section, we discuss the thoughts of our data collection team and the ground truth as they were observed, and compare these with the results obtained through our topic modeling visualization tool.

A. Prominent Misinformation Themes Over Time

Although a variety of misinformation themes were identified, particular dominant themes stood out, changing over time. These themes were considered as dominant based on a simple sum of their frequency of occurrence in our data set. During the month of March, the prominent misinformation theme was the promotion of remedies and techniques to supposedly prevent, treat, or kill the novel coronavirus. During the month of April, the prominent themes still included the promotion of remedies and techniques, but additional prominent themes began to stand out. For example, several misinformation stories attempted to downplay the seriousness of the novel coronavirus. Others discussed the anti-malaria drug hydroxychloroquine. Others promoted the idea that the virus was a hoax meant to defeat President Donald Trump. Others consisted of various attempts to attribute false claims to high-profile people, such as politicians and representatives of health organizations. Also in April, although first signs of these were seen in March, the idea that 5G caused the novel coronavirus began to become more prevalent. During the

month of May, the prominent themes shifted to predominantly false claims made by high-profile people, followed by attempts to convince citizens that face masks are either more harmful than not wearing one, or are ineffective at preventing COVID-19, and how to avoid rules that required their use. The number and variety of identity theft phishing scams also increased during May. Misinformation items attempting to attribute false claims to high-profile people continued throughout May. Also becoming prominent in May were misinformation items attempting to spread fear about a potential COVID-19 vaccine, and items promoting the use of hydroxychloroquine. During the month of June, the prominent theme shifted significantly to attempts to convince citizens that face masks are either more harmful than not wearing one, and how to avoid rules that required their use. Phishing scams also remained prominent during June. During the month of July, the dominant themes of the misinformation items shifted back to attempts to downplay the deadliness of the novel coronavirus. Another prominent theme in July was the proliferation of attempts to convince the public that COVID-19 testing is inflating the results.

B. Topic Streams

After using the tool described in Section III-C, we generated the graphs and tables described and discussed in this section. Our data for this step contained 243 unique misinformation narratives spanning from January 2020 to June 2020, when we stopped data collection. The data was curated by our research team through the process described in the methodology. Each entry contains, among other fields, a “date” used as a chronological identifier, a “title” describing the general idea the misinformation is attempting to convey, and a “theme” field putting the story in a concisely described category. For example, a story given the title “*US Department of Defense has a secret biological laboratory in Georgia*” is categorized in the following theme: “*Western countries are likely to be purposeful creators of the new virus.*” Each topic was represented by an identification number up to 20 and a set of 10 words. We picked the three most relevant words that best represented the general idea of each topic. Notably, obvious words such as *covid* or *coronavirus* were removed from the topic descriptions since they are common for every topic.

In Tables I and II, we described some of the twenty topics found by each of our LDA models. These topics were chosen because they each described a precise narrative and have a low topic distribution (or proportion within the corpus). A low proportion is desirable because this indicates the detection of a unique narrative within the corpus; as opposed to an overarching topic including general words such as “world”, “outbreak”, or “pandemic”. Do note that topic inclusiveness is not exclusive and documents can be part of multiple topics. This becomes apparent in tables I: from our topic model, we found a dominant topic encompassing 68% of narratives. It includes words such as “Trump”, “outbreak”, “president”, etc. Some other narratives also included words such as “flu”, “news”, or “fake”. Because the evolution of these narratives are consistent across the corpus and show little temporal

fluctuation, we chose not to report on them further. For these reasons, the narratives we focused on below show a low percentage of distribution (Tables I & II).

TABLE I
MOST FREQUENT DOMINANT TOPICS FROM TITLES.

Topic ID	Word 1	Word 2	Word 3	Proportion
10	china	chinese	spread	2%
12	scam	hydroxy...	health	2%
17	state	donald	trump	2%
18	vaccine	gates	bill	5%

TABLE II
MOST FREQUENT DOMINANT TOPICS FROM THEMES.

Topic ID	Word 1	Word 2	Word 3	Proportion
3	fear	spread	western	2%
9	predicted	pandemic	vaccine	2%
16	phishing	hydroxy...	vaccine	2%

1) *Using narrative titles as a corpus*: The general narratives described by the topics were thus:

- Topic 10 described the narratives related to the Chinese government and its responsibility in the spread of the virus. These stories represented an estimated 2% of the 243 stories collected.
- Topic 12 described the narratives related to personal health and scams or misinformation such as the benefits of hydroxychloroquine. These stories represented an estimated 2% of the 243 stories collected.
- Topic 17 described the narratives related to the response of Donald Trump and his administration. These stories represented an estimated 2% of the 243 stories collected.
- Topic 18 described the narratives related to the involvement of Bill Gates in various conspiracies, mostly linked to vaccines. These stories represented an estimated 4% of the 243 stories collected.

Related studies have found that finger-pointing narratives usually lead to negative sentiment and toxicity in online communities [1, 5, 7].

Figure 1 shows the evolution of Topic 10, the topic describing China-related narratives. It shows that these narratives were already in full force from the beginning of our corpus and slowly came to a near halt during the month of April. We notice a short spike again towards the end of the corpus during the month of June. This is consistent with the ground truth of online narratives that focused on the provenance of the virus during the early stages.

Figure 2 shows the evolution of Topic 12, the topic describing narratives related to health, home remedies, and general hoaxes and scams stemming from the panic. We can see it was consistent with the rise of cases in the United States and panic increased as with the spread of the virus. It is interesting to




Fig. 1. Topic’s probability distribution of titles for topic 10 (keywords: china, chinese, spread)




Fig. 3. Topic’s probability distribution of titles for topic 17 (keywords: donald, trump, state)

note that this figure roughly coincides with the daily number of confirmed cases for this time period [15].

effort. As with Figure 1, these narratives were especially strong early on (albeit this narrative remained active for a slightly longer time), before coming to a near halt.




Fig. 2. Topic’s probability distribution of titles for topic 12 (keywords: hydroxychloroquine, health, scam)

Figure 3 shows the evolution of Topic 17. This topic described stories related to Donald Trump and his administration. These stories generally referred to claims that the virus was manufactured as a political strategy, or claims that various public figures were speaking out against the response of the Trump administration.

Figure 4 shows the evolution of Topic 18. This topic described stories such as Bill Gates and his perceived involvement with a hypothetical vaccine, and other theories describing the virus’ appearance and spread as an orchestrated




Fig. 4. Topic’s probability distribution of titles for topic 18 (keywords: bill, gates, vaccine)

We notice that, as theories about the origins of the virus slowed down, hoaxes and scams increased - as shown on Figure 2. This includes attempts at identity theft, especially toward senior citizens, and attempts to sell miracle cures and miracle personal protection items.

2) *Using narrative themes as a corpus:* For this section, we inputted narrative themes as the corpus. Note that the topic IDs are independent from the previous set of topics using

titles. Similarly to section IV-B1, we found a dominant topic encompassing 68% of narratives as well. This time including words such as “attempt”, “countries”, and “purposeful”. As for section IV-B1, we chose not to report on that topic as well as other smaller but general topics showing little fluctuation. Therefore, the narratives we focused on below show a low percentage of distribution. The general narratives described by the topics are thus:

- Topic 3 described the narratives related to the speculations on the spread of the virus, especially in an international relations context. These stories represented an estimated 2% of the 243 stories collected.
- Topic 9 described the narratives related to stories claiming the creation and propagation of the virus were either designed or predicted, along with voices claiming a vaccine already exists. These stories represented an estimated 3% of the 243 stories collected.
- Topic 16 described the narratives related to personal health and scams or misinformation such as the benefits of hydroxychloroquine. These stories represented an estimated 2% of the 243 stories collected.




Fig. 5. Topic’s probability distribution of themes for topic 3 (keywords: fear, spread, western)

Figure 5 shows the evolution of Topic 3. It is linked to early fear of the virus and presented narratives as opposing the western block with the East, notably China. It matched closely with Figure 1 and its China-related narratives. In both cases, we see an early dominance of the topic followed by a near halt as the virus touched the United States.

Figure 6 describes the evolution of narratives claiming the virus was predicted or even designed. This figure is consistent with the results shown by Figure 4 which shows claims regarding Bill Gates, early vaccines, etc. They both showed stories of early knowledge of the virus and peaked early, appearing more or less sporadically as time goes on and as cases increased.




Fig. 6. Topic’s probability distribution of themes for topic 9 (keywords: predicted, pandemic, vaccine)




Fig. 7. Topic’s probability distribution of themes for topic 16 (keywords: hydroxychloroquine, vaccine, phishing)

Figure 7 is parallel to Figure 2. Both showed hoax stories promoting scams and health-related misinformation. We noticed an early rise in Figure 7, most likely due to the inclusion of the keyword “vaccines” in the topic, which caused some overlap with Topic 9 as shown in Figure 6.

C. YouTube Data

In this section, we explore how different topic models affect our YouTube data set. We focus on a subset of data published during the month of March to limit the number of comments to process.

1) *YouTube Videos*: The first observation for this set is that our HDP model did not perform as well as the LDA model. Our HDP model identified one dominant topic present in

87% of videos, with seemingly unrelated identifying keywords (“cases”, “hindi”, “nyc”, “italy”). While the rest of the topics are present in around 1% of the videos. The second most dominant topic (1.8% of documents) also features contradicting words such as “plandemic” and “hospitals”. One would expect language connected to the plandemic narrative in this topic, such as mentions of “Bill Gates” like we saw in the previous sets, but it is missing. There are two possible explanations for this. One is that performance may be due to the size of the set (more in the next section) as there were only 444 video titles processed. The other is that the set features numerous multilingual titles, which may skew results.

TABLE III
RELEVANT TOPICS FROM VIDEO TITLES (LDA MODEL)

Topic ID	Word 1	Word 2	Word 3	Proportion
0	news	update	live	12.4%
17	outbreak	doctor	cases	7.6%
6	plandemic	dempanic	dem	2.7%

Our LDA model, however, behaved as expected and was able to identify major topics, mostly news videos (Topics 0 & 17), as well as what we suspect to be a vehicle of misinformation (Topic 6). As described in Table III and visualized in Figure 8. Figure 8 has been smoothed with a moving average equal to 15% of the total data set size (67) in order to improve legibility and reveal patterns. Due to most of the videos being published late in March, this has removed some granularity towards early March from the plot. However, we notice news topics staying fairly consistent while Topic 6 sees a decline, possibly as the number of covid cases makes maintaining the “fake pandemic” narrative more difficult and other misinformation narratives take over such as various scams and hoaxes as seen in section IV-B1.




Fig. 8. Topic’s probability distribution of topics 0, 17 & 6 (LDA model)

2) *YouTube Comments*: Contrary to the previous section, this is a much larger data set of 652,120 comments. This led to better performances, but still inferior to the LDA model. Our HDP model was able to identify non-English comments (11.4% German, 4.5% Spanish, 1.6% French). More importantly, the HDP model identified a topic that could be described as polarizing discourse, some of the most frequent terms including “Trump”, “China”, and “virus”. This topic accounts for 6.6% of the corpus. The evolution of this topic is shown by Figure 9 where we notice that topic is on an upward trend. A moving average equal to 3% of the set size is applied to better identify patterns.




Fig. 9. Topic’s probability distribution of Topic 4 (HDP model)

On this very large set, our HDP model somewhat outperformed LDA for our purposes as it was able to identify a probable topic for misinformation. When applied to our comments set, our LDA model mostly found general terms while also successfully isolating non-English comments. The model did identify a topic with some toxic language and some that could be used in a hostile way or communicate sinophobic sentiments (Topic 7 & 17). See Table IV. While discussion of China has so far been on a downward trend since the start of the pandemic, the mention of the term “virus” along with “china” suggests toxic behavior. See Figure 10.

TABLE IV
RELEVANT TOPICS FROM COMMENTS (LDA MODEL).

Topic ID	Word 1	Word 2	Word 3	Proportion
7	china	virus	made	3.5%
17	trump	dumb	bats	3.3%

D. Public website and citizen science

We have put together a website with known cases of misinformation about COVID-19. As of January 2021, we




Fig. 10. Topic's probability distribution of Topic 7 & 17 (LDA model)

have documented close to 600 cases that we identified from numerous sources (social media - Facebook, YouTube, Twitter, blogs, fake websites, robocalls, text/SMS, WhatsApp, Telegram, and an array of such apps) - see Figure 11 [8]. The principal difference between our effort and other similar efforts by Google and social media companies is that we are paying special attention to cases of misinformation and scammers that are affecting our region, while also including global cases. We update the database periodically with newly detected cases. Moreover, we have put together a list of over 50 tips on the website for people to learn how to spot misinformation. We have also provided a feature for people to report fake websites or scams that are not currently in our database.

Our website uses a three-pronged approach:

- We identify new cases of fake websites, misinformation content, and bad actors. We use social network analysis and cyber forensic methodologies to identify such cases.
- We believe in educating people to be self-reliant because we might not be able to detect all possible cases of misinformation. Therefore, we go through identified cases and prepare a list of common telltale signs to detect whether a piece of information is genuine or not.
- For the cases that are not in our database and people cannot distinguish, we provide a way for people to submit cases of misinformation that we have not captured in our database.

The database of known misinformation cases and scams is publicly available for the research community to use [8]. We envision a tremendous value of this research database to various disciplines. The website is available for regulatory bodies (Arkansas Office of the Attorney General) and any citizen, which serves as an invaluable resource to not only educate people of the misinformation and scams about COVID-19 but also assisting legal authorities in taking action

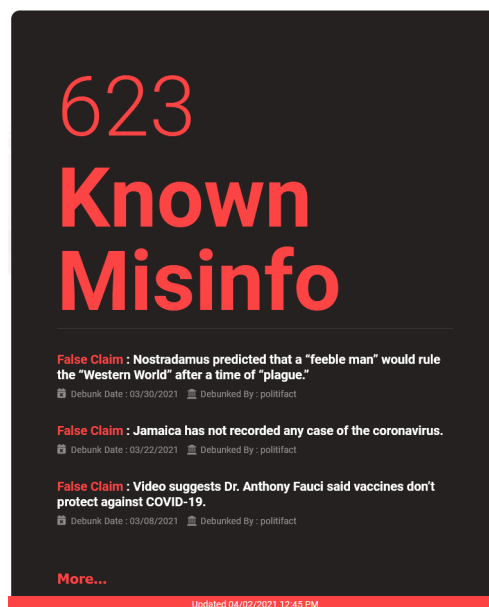


Fig. 11. COVID-19 Website Front page - Showing the latest misinformation stories

against malicious actors and groups. We are assisting the Arkansas' Attorney General's office by providing reports on cyber forensic evidence about scam/fake websites reported by people. The study presented in this paper will be developed into the system as a real-time campaign tracking feature. We will continue to work with Arkansas' Attorney General's office to assist in their effort to combat COVID-19 misinformation and scams to protect Arkansans.

V. CONCLUSION

This study has highlighted some of the narratives that surfaced during the COVID-19 pandemic. From January 2020 to July 2020, we collected 243 unique misinformation narratives and proposed a tool to observe their evolution. We have shown the potential of using topic modeling visualization to get a bird's eye view of the fluctuating narratives and an ability to quickly gain a better understanding of the evolution of individual stories. We have seen that the tool is efficient to chronologically represent actual narratives pushed to various outlets, as confirmed by the ground truth observed by our misinformation curating team and independent international organizations. Working with the Arkansas Office of the Attorney General, this study illustrates a relatively quick technique for allowing policy makers to monitor and assess the diffusion of misinformation on online social networks in real-time, which will enable them to take a proactive approach in crafting important theme-based communication campaigns to their respective citizen constituents. We have made most of our findings available online to support this effort.

We have also seen in this study that using carefully curated "themes" - which offer a lexical value close to the abstract topics provided by the LDA model - yields similar results to

using misinformation narratives “title”. For this reason, we scaled up our data and repeated our methodology on social media data: YouTube video titles, and their comments. Over concerns of our LDA topic model becoming difficult to scale, we experimented with a HDP (Hierarchical Dirichlet Process) model, which attempts to infer the number of topics.

We found promising but unsurprisingly less precise results. We notice that HDP was able to isolate a probable subset of polarizing comments. One possible way forward would be to use HDP to identify these subsets, filter out irrelevant comments, then apply the LDA model. This may reveal various narratives, some of them spreading misinformation, and further automate the process of identifying online misinformation in uncontrolled spaces.

ACKNOWLEDGEMENT

This research is funded in part by the U.S. National Science Foundation (OIA-1946391, OIA-1920920, IIS-1636933, ACI-1429160, and IIS-1110868), U.S. Office of Naval Research (N00014-10-1-0091, N00014-14-1-0489, N00014-15-P-1187, N00014-16-1-2016, N00014-16-1-2412, N00014-17-1-2675, N00014-17-1-2605, N68335-19-C-0359, N00014-19-1-2336, N68335-20-C-0540), U.S. Air Force Research Lab, U.S. Army Research Office (W911NF-20-1-0262, W911NF-16-1-0189), U.S. Defense Advanced Research Projects Agency (W31P4Q-17-C-0059), Arkansas Research Alliance, the Jerry L. Maulden/Entergy Endowment at the University of Arkansas at Little Rock, and the Australian Department of Defense Strategic Policy Grants Program (SPGP) (award number: 2020-106-094). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations. The researchers gratefully acknowledge the support.

REFERENCES

- [1] A. Abd-Alrazaq et al. “Top Concerns of Tweeters During the COVID-19 Pandemic: Inveillance Study”. In: vol. 22. 4. 2020, e19016. DOI: 10.2196/19016. URL: <http://www.ncbi.nlm.nih.gov/pubmed/32287039>.
- [2] M. H. Alam, W.-J. Ryu, and S. Lee. “Hashtag-based topic evolution in social media”. In: vol. 20. 6. Nov. 2017, pp. 1527–1549. DOI: 10.1007/s11280-017-0451-3. URL: <https://doi.org/10.1007/s11280-017-0451-3>.
- [3] D. M. Blei, J. D. Lafferty, and A. N. Srivastava. *Text Mining: Classification, Clustering, and Applications*. CRC Press, 2009, pp. 71–88.
- [4] D. M. Blei, A. Y. Ng, and M. I. Jordan. “Latent dirichlet allocation”. In: vol. 3. 2003, pp. 993–1022.
- [5] H. Budhwani and R. Sun. “Creating COVID-19 Stigma by Referencing the Novel Coronavirus as the “Chinese virus” on Twitter: Quantitative Analysis of Social Media Data”. In: vol. 22. 5. May 2020, e19301. DOI: 10.2196/19301. URL: <http://www.ncbi.nlm.nih.gov/pubmed/32343669>.
- [6] M. Calabresi. *Inside Russia’s Social Media War on America*. 2017. URL: <https://time.com/4783932/inside-russia-social-media-war-america/>. (accessed: 01.19.2021).
- [7] R. Chandrasekaran et al. “Topics, Trends, and Sentiments of Tweets About the COVID-19 Pandemic: Temporal Inveillance Study”. In: vol. 22. 10. Oct. 2020, e22624. DOI: 10.2196/22624. URL: <http://www.jmir.org/2020/10/e22624/>.
- [8] COSMOS. *COSMOS - COVID-19 Misinformation Tracker*. 2021. URL: <https://cosmos.ualr.edu/covid-19>. (accessed: 06.15.2021).
- [9] EUvsDisinfo. *EUvsDisinfo. March 16, 2020. The Kremlin and Disinformation About Coronavirus*. 2020. URL: <https://euvsdisinfo.eu/the-kremlin-and-disinformation-about-coronavirus/>. (accessed: 01.19.2021).
- [10] S. Al-khateeb et al. *Exploring Deviant Hacker Networks (DHM) on Social Media Platforms*. Vol. 11. Journal of Digital Forensics, Security and Law, 2016, pp. 7–20. DOI: 10.15394/jdfsl.2016.1375. URL: <https://commons.erau.edu/jdfsl/vol11/iss2/1>.
- [11] R. Kouzy et al. “Coronavirus Goes Viral: Quantifying the COVID-19 Misinformation Epidemic on Twitter”. eng. In: vol. 12. 3. Publisher: Cureus. Mar. 2020, e7255–e7255. DOI: 10.7759/cureus.7255. URL: <https://pubmed.ncbi.nlm.nih.gov/32292669>.
- [12] Q. Liu et al. “Health Communication Through News Media During the Early Stage of the COVID-19 Outbreak in China: Digital Topic Modeling Approach”. In: vol. 22. 4. Apr. 2020, e19118. DOI: 10.2196/19118. URL: <http://www.jmir.org/2020/4/e19118/>.
- [13] G. Pennycook et al. “Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for a Scalable Accuracy-Nudge Intervention”. In: vol. 31. 7. _eprint: <https://doi.org/10.1177/0956797620939054>. 2020, pp. 770–780. DOI: 10.1177/0956797620939054. URL: <https://doi.org/10.1177/0956797620939054>.
- [14] R. Řehůřek and P. Sojka. “Software Framework for Topic Modelling with Large Corpora”. In: May 2010, pp. 45–50. DOI: 10.13140/2.1.2393.1847.
- [15] H. Ritchie et al. *United States: Coronavirus Pandemic - Our World in Data*. 2020. URL: <https://ourworldindata.org/coronavirus/country/united-states?country=~USA>. (accessed: 07.29.2020).
- [16] A. Schofield, M. Magnusson, and D. Mimno. “Pulling Out the Stops: Rethinking Stopword Removal for Topic Models”. In: *15th Conference of the European Chapter of the Association for Computational Linguistics*. Vol. 2. Association for Computational Linguistics, 2017, pp. 432–436.
- [17] H. Sha et al. *Dynamic topic modeling of the COVID-19 Twitter narrative among U.S. governors and cabinet executives*. 2020. arXiv: 2004.11692 [cs.SI].
- [18] Y. W. Teh et al. “Hierarchical Dirichlet Processes”. In: vol. 101. 476. Publisher: Taylor & Francis _eprint: <https://doi.org/10.1198/016214506000000302>. 2006, pp. 1566–1581. DOI: 10.1198/016214506000000302. URL: <https://doi.org/10.1198/016214506000000302>.
- [19] Y. Zhang, W. Mao, and J. Lin. “Modeling Topic Evolution in Social Media Short Texts”. In: *2017 IEEE International Conference on Big Knowledge (ICBK)*. 2017, pp. 315–319.