Twitter User Sentiments Analysis: Health System Cyberattacks Case Study

Muhammad Abusaqer
Department of Math and Computer Science
Minot State University
Minot, ND, USA
muhammad.abusaqer@minotstateu.edu

M. Benaoumeur Senouci
The Faculty of Engineering
SDU Mechatronics (CIM)
University of Southern Denmark
Sønderborg, Denmark
senouci@sdu.dk

Kenneth Magel
Department of Computer Science
North Dakota State University
Fargo, ND, USA
kenneth.magel@ndsu.edu

2023 International Conference on Artificial Intelligence in Information and Communication (ICAIIC) | 978-1-6654-5645-6/23/$31.00 ©2023 IEEE | DOI: 10.1109/ICAIIC57133.2023.10067026

***Abstract*—Social media, such as Twitter, allow people to interact with ongoing events and share their sentiments. Therefore, people use social media to report and express their emotions about events they are experiencing. Furthermore, some officials take advantage of the popularity of social media to keep the public informed, especially during emergent events. Researchers have covered sentiment analysis on Twitter in many fields, such as movie reviews, stocks, politics, health, and sports. However, there is a research gap in studying the public's concerns on social media when a cybersecurity breach occurs and how people's sentiment changes over time. To fill the gap, The researchers selected the cyberattacks against Universal Health Services (UHS) during the late days of September 2020 and collected a large dataset of related tweets over five weeks. Live-streaming tweets and historical ones both were compiled. The focus while gathering tweets was in the context of cyberattacks on UHS using keywords and hashtags such as Universal Health System, UHS cyberattack, UHS Ransome, UHS security breach, and UHS locked. Then, the researchers determined tweets' sentiment classification on this developing event using deep learning of Long Short-Term Memory (LSTM) and Artificial Neural Networks (ANN) and their accuracies. Furthermore, the researchers performed exploratory data analysis for the dataset supplying information about how sentiment has changed over time to compare the sentiment per week since the start of these cyberattacks on UHS. This study is the first to provide an analysis of the public's sentiment toward a significant cybersecurity breach on a healthcare provider dealing with COVID-19 based on a large-scale dataset extracted from social media feeds.**

Keywords—Sentiment analysis, cyberattacks, ransomware, healthcare, security breach.

I. INTRODUCTION

Before the era of social media, individuals often shared their sentiments when they were physically close to people they know, such as friends and family members, using talks over face-to-face, on the phone, and writing letters. However, social media such as Twitter helps people broadcast new events, give further information about ongoing events, and shares their opinions and emotions on different issues. Consequently, Twitter currently hosts a gold mine of live streaming for raw data, which needs analysis. That being said, it provides motivation to collect and analyze tweets during important events of hacking the computer system of a healthcare provider in the USA. Then, the dataset was used to know how people reacted to the event, their emotions, and how their impressions changed over time, especially cyberattacks hitting a major health provider during a global pandemic.

At present, computer systems are involved almost in all operations of any organization providing services to the public. Therefore, computer systems failure will undoubtedly form a big concern for these institutions, especially if it belongs to a major healthcare provider that supplies a vital service during the outbreak of a global epidemic. This unique character of the healthcare industry makes it a desirable target for cybercriminals. The following quote indicates the severity and impact of such a cyberattack [1], "With an increased burden on the healthcare system due to COVID-19, cybercriminals know they have golden opportunities to make money from healthcare targets. During the first three months of 2020, the number of breached records in the healthcare sector exploded by 273% over the same period in 2019." Because of these characteristics of healthcare providers targeted by cyberattacks occurring in the middle of a global epidemic, it is important to study people's sentiments on social media such as Twitter and find out how it changes as time progresses. The following research questions are addressed in the study 1) How does the public's sentiment toward the cyberattacks against a primary healthcare provider, the United Health System – UHS, change over time? 2) Were the people's sentiments more positive or more negative? 3) Were people interested all along in this event? 4) What negative sentiments do the public present more during cyberattacks?

This research studies the Twitter feed associated with cyberattacks against the UHS hospital network. In the study, the researchers focused on tweets in the context of event-describing keywords such as UHS cyberattacks and UHS ransomware attacks to gather tweets. The researchers also kept related tweets to this cyberattack, such as warning messages by law-enforcement agencies of imminent similar cyberattacks, tweets reporting similar incidents, and explanations of what's happening. Then, the researchers studied the changes in emotions over one month, starting on the day of the cyberattack breach on UHS occurred.

The remainder of this paper is organized as follows. Related research about sentiment analysis over time is presented in section II. The Methodology of the data corpus that is used for this study is presented in section III. Sentiment analysis calculation is presented in section IV, followed by results and discussion V. Conclusion is outlined in Section VI

II. RELATED WORK

Mansoor et al. [2] presented how the global sentiment analysis of COVID-19-related tweets collected from different countries has changed over time. They also studied the consequences of COVID-19 on two life activities, working from home and online learning. The authors have found positive sentiment toward tweets related to COVID-19 was higher in Pakistan, South Africa, Mali, and Bangladesh. However, the negative sentiment was higher in the worst-hit

countries such as the USA, India, Australia, the United Kingdom, Brazil, and Turkey. Concerning work from home and online learning, positive sentiment was higher than negative sentiment in tweets work from home and online learning. Still, in online learning tweets, there are times when negative sentiment is higher. The paper used only a built-in Python library, VADER, to calculate the sentiment analysis scores of the collected tweets, which can be a starting point for grasping a basic understanding of the collected data. VADER stands for Valence Aware Dictionary for Sentiment Reasoning, a python library used to find text sentiment. VADER uses emotional words in a piece of text to determine its sentiment. However, analysis solely based on sentiment scores from VADER is not convincing enough. VADER does not consider the context of the sentence of the words of expression [3].

Sentiment analysis has been applied to various situations during the last few years [4] [5]. Trupthi, Pabboju, and Narasimha performed a sentiment analysis of live-streaming tweets based on Hadoop, which automatically analyzes large numbers of tweets [6]. Bergsma et al. [7] used tweets as a dataset to predict personal features such as ethnicity and gender. Bergsma and his colleagues applied clustering algorithms on declared user information posted on Twitter profiles, such as name, location, and friend list, to find hidden personal information. In [8], the authors used spark streaming to present a model for live-streaming tweets of Iraqi related to a recent controversy concerning the soccer player Bassam AlRawi.

The authors of [2] provided a global sentiment analysis of tweets related to the Coronavirus. They showed how people are changing their opinions over time. [9] presented a thorough study of how Twitter was used to disseminate information about the importance of physical activities during pregnancy. In [10], the authors demonstrated using the Twitter feed to detect power outages in certain regions.

Pano and Kashif introduced new text preprocessing methods to find connections between the sentiment scores and the prices of Bitcoin. They relied on VADER to find the sentiment scores [11]. According to [12], VADER performs better with slang and emojis, whereas TextBlob does better with formal English.

III. METHODOLOGY

A. Cyberattacks Against UHS Hospitals Chain

Early on Sunday, September 27, 2020, a group of 400 hospitals for UHS in the United States was cyber-attacked by ransomware. UHS is a major provider of hospital and healthcare services in the United States, with 89,000 employees. UHS operates in the USA and the United Kingdom and runs about 400 facilities[13].

UHS was targeted by a ransomware cyberattack which caused the regular services of the IT systems to become completely paralyzed. According to reporting by NBC News, the computerized systems of UHS began to collapse over the weekend to the level that some hospitals were forced to return to documenting patient information with pen and paper, which was reported as "potentially [the] largest in US history," according to NBC news [14]. UHS posted a statement the next day confirming that the IT network across its facilities was down [15].

The hospitals' cyberattacks carefully choose the time and target. For example, the cyberattack occurred on a weekend when primary IT personnel would not be present. In addition, all efforts at that time were dedicated to fighting the COVID-19 pandemic. In that period, the number of cyberattacks targeting hospitals increased significantly[16] [17] [18].

B. Data Collection

The researchers used two approaches to collect the dataset. In the first approach, Python snippets were coded using Twitter API to access the tweets streaming on Twitter. The live-streaming tweets were collected. The data was saved to a MySQL database system.

According to [19], Twitter's live-streaming API will not give all relevant tweets; it only brings a sample of about 1% of posted related tweets at a given moment. In addition to that, the 1% does not supply enough coverage of the topic because it is just a sample. Also, a collected sample might be biased, according to [20]. So, the researchers had to find a second complementary way to get more tweet coverage on this topic. Therefore, data scraping was used to gather historical tweets for better coverage [21].

Keywords were used standing for the event in both approaches to gathering relevant tweets. Examples of the keywords are #UniversalHealthServices, cyberattack, loss of control, and security breach. In addition, the researchers used the boolean operators and, or to ensure that they collected only concerned tweets, as shown in Fig. 1.

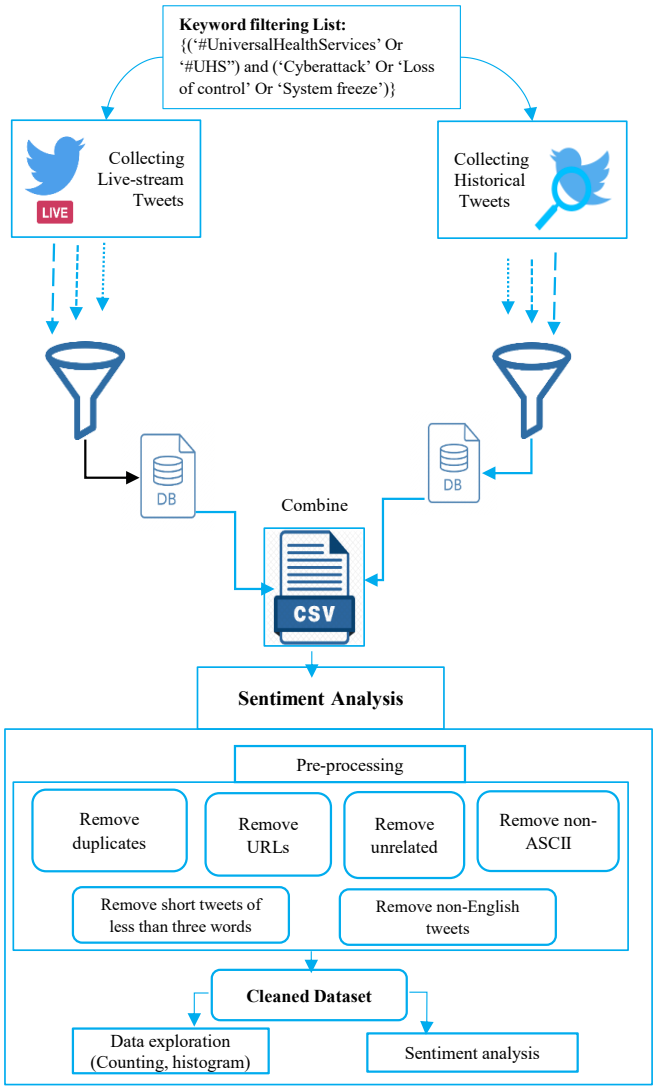

Fig. 1. Data collection and Preprocessing

Moreover, it is worth mentioning keywords and hashtags such as 'hospitals' and 'cyberattack' were so common on Twitter at that time, leading to retrieving many irrelevant tweets about COVID-19.

C. *Cleaning the Dataset*

Since the two datasets collected using two integrated approaches were merged, it resulted in duplicated tweets. Therefore, to clean the combined dataset, first the authors used a Python snippet to identify and delete duplicates taking advantage of the unique tweet id property [22]. However, duplicating retweeted tweets was kept because they had unique ids.

Further steps to clean the dataset were also performed using Python's regular expression - regex operations to clean each tweet by removing links and special characters. Regex matching is a text string used to describe a search query while extracting information from the text [23]. The links and special characters, such as white spaces, HTML tags, and punctuations, do not contain much information, which is why the researchers removed them [24]. Also, non-ASCII characters, such as emojis, were stripped from tweets. In addition, short tweets consist of three words or less and cannot give correct sentiment, so they were removed. Finally, all non-English tweets were removed to ensure they would not distract from the deep learning model during the training stage.

D. *Calculation Sentiment analysis using Deep Learning.*

The researchers obtained tweets' sentiment using deep learning of text transformer [25] [26] [27] [28] [29]. They used the deep learning model of RoBERTa, which is better suited to the model, as shown in the algorithm presented in Fig. 2.

The dataset's deep learning model was exclusively fine-tuned (transfer learning). RoBERTa stands for Robustly Optimized BERT Pre-training Approach [30]. RoBERTa is an AI model created by Facebook Research. Facebook Research team trained model RoBERTa on more than 124M tweets (from January 2018 to December 2021) for self-supervised natural language processing (NLP) [31] [32] [33] [34].

```
1. # Sentiment Analysis using Deep learning model of RoBERTa
2. initialization
3. import the necessaries libraries: transformers, pipeline
3. create an object from the class pipeline(model='cardiffnlp/twitter-
RoBERTa-base-sentiment')
4. read the dataset into a Panda dataframe object for manipulation
5. loop for each tweet T
6.    clean each tweet by removing URLs, white space, if not yet
7.    replace each mention with a general one of @user
8.    tokenize each word in the tweet
9.   do the training to transfer learning to the dataset and get the
sentiment
10.   separated the computed sentiment tuple into label and score
11. end the loop
12. save the sentiment as a new column/field in the dataframe.
13. save the dataframe into a dataset file of CSV format for later
processing.
```

Fig. 2. Pseudo algorithms calculate the sentiment using the deep learning model of RoBERTa.

This is one reason it was chosen, as it is tailored to Twitter sentiment. Researchers picked it after evaluating it with other available models. The model usually takes a significant amount of time in the training phase. In this case, it lasted for a day using a main-average computer of continuous run. Finally, this resulted in the sentiment label and score for each tweet. Researchers used python code to separate each tweet's sentiment label and score into two columns. The sentiment label and score were added as two new columns in the CSV version of the dataset, as described in Fig. 2. Sentiment score represents the intensity of the label. For example, a tweet with a score of 0.8 has stronger positive sentiment than another of 0.4. The same analogy applies to the sentiment score of negative and neutral labels.

IV. Results

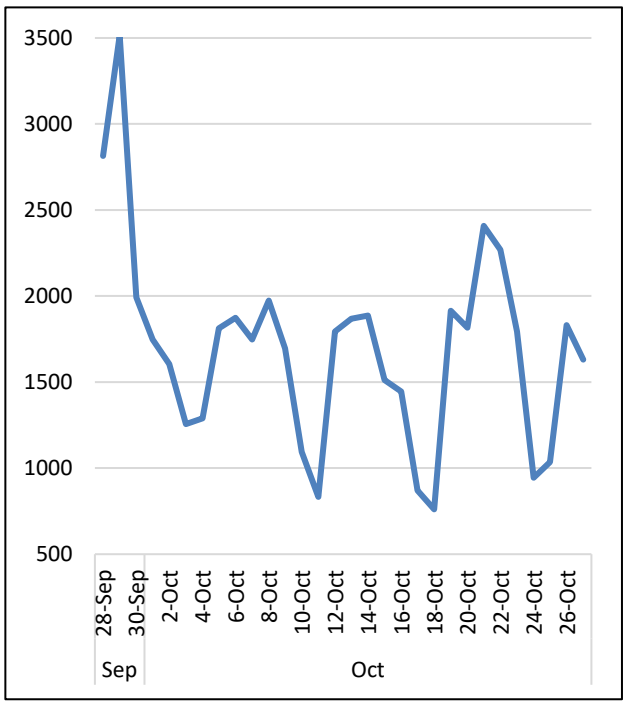

Fig. 3. The chart shows how the number of tweets per day. In addition, it shows how the count of tweets has changed over the days.

Fig. 3 shows the result of plotting the number of tweets per day.

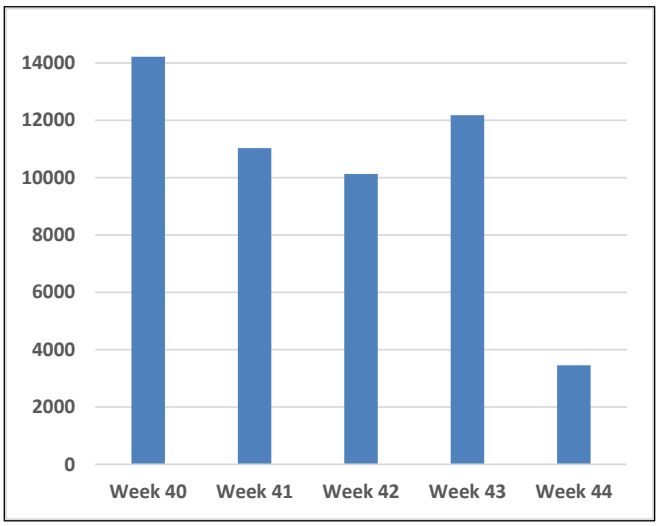

Fig. 4. The chart shows the count of tweets per week. For example, week 40 of the year 2020 spans from 9/272020 to 10/3/2020, and week 44 starts on 10/24/20 and ends on 10/31/2020.

Fig. 4 shows the number of tweets per week.

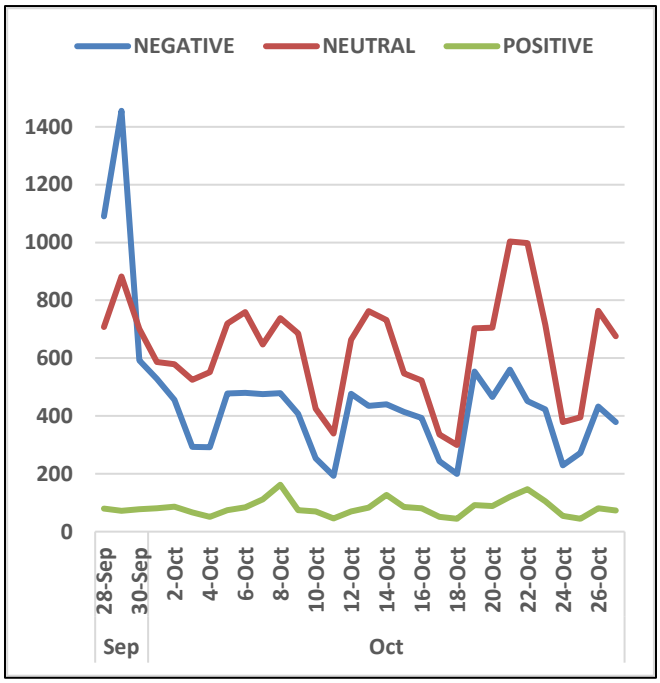

Fig. 5. It shows how positive, negative, and neutral sentiments vary over time.

Fig. 5 shows how sentiment has changed over time by counting the number of positive, negative, and neutral sentiments each week.

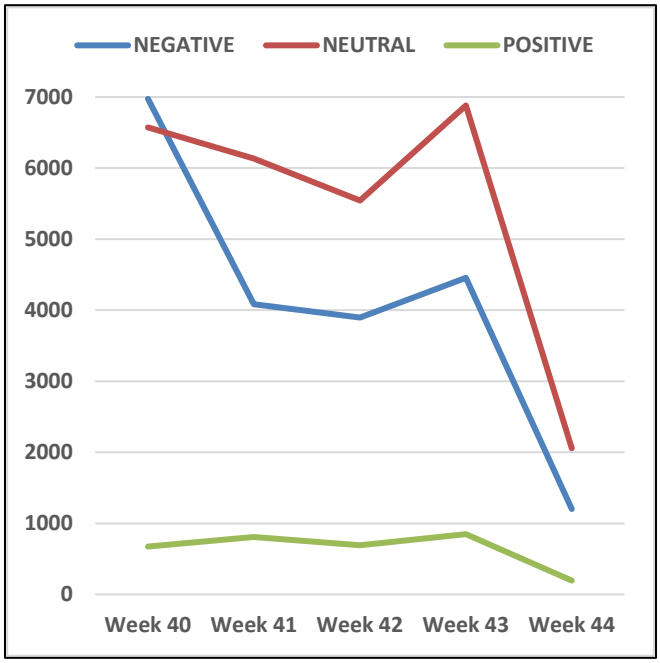

Fig. 6. The figure shows how sentiment changed per week.

Fig. 6 shows how the average sentiment changes per week.

The researchers plotted the public's weekly average of positive, negative, and neutral sentiment scores per day in Fig. 7.

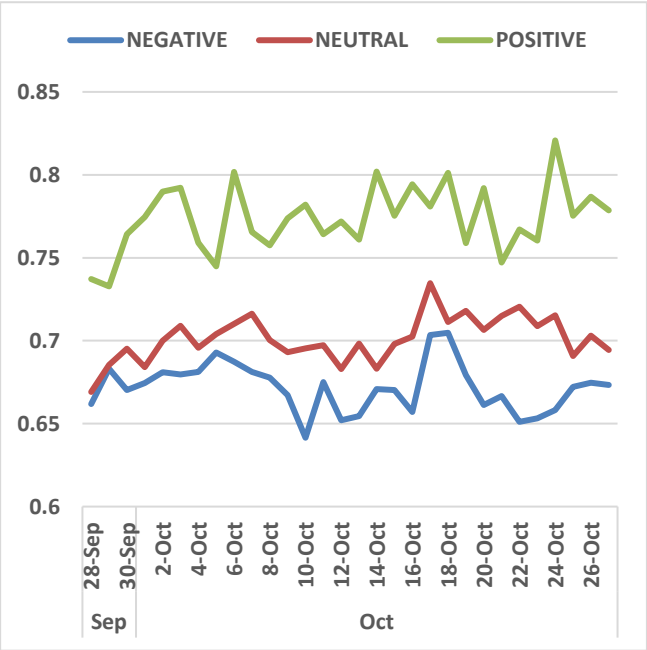

Fig. 7. It shows how the average score changes per day of the three classes of the sentiment.

Numbers in TABLE I are for the average of the sentiment in each week of the study and the grand average. The following Table shows how the average sentiment of tweets varied over weeks.

	Average Sentiment Score		
	Negative	Neutral	Positive
Week 40	0.67	0.69	0.76
Week 41	0.68	0.70	0.77
Week 42	0.67	0.70	0.78
Week 43	0.66	0.71	0.77
Week 44	0.67	0.70	0.78
Grand Total	0.67	0.70	0.77

V. Discussion

It can be seen from Fig. 3 there is a peak in the number of tweets during the first three days of the cyberattack incident. It occurred because of many tweets and retweets reporting the incident. Some related it to a potentially similar one during the upcoming presidential election in November 2022 in the USA. The number of tweets then significantly dropped. There is another local peak on October 21 due to warning messages issued by federal and local law enforcement agencies and other in charge departments, such as Homeland Security, of imminent similar cyberattacks. Many people retweeted those tweets since they came from authoritative sources. Others retweeted the UHS cyberattacks tweets to remind people what occurred and what can be expected.

The study aggregated the data in Fig. 4 to smooth it over weeks to avoid any sharp daily fluctuations and plotted the number of tweets per week in histogram columns. It is done to notice if specific forming patterns took place. Surprisingly, the researchers found a specific and noticeable pattern that formed showed the people's response on Twitter to these cyberattacks on UHS. First, the number of tweets during the first week of the cyberattacks was significantly larger than in the following five weeks. Then, another increase was observed in week 43 (of the year 2020), but the recorded increase is less than in the initial wave in the first week. This is because, in the first week of cyberattacks, the event was still somehow new and had never happened before on this scale. The researchers think the second raise was due to the following reasons. 1) People have become more vigilant of such security breaches targeting hospitals. Therefore, some individuals tweeted those info-tweets to draw the public's attention to this critical issue and raise their awareness. 2) The many warning messages issued by the FBI, Homeland Security Department, and others in charge agencies were found in the dataset, and many volunteers tweeted and retweeted those messages. 3) The fear and caution of the recurrence of new similar cyberattacks targeting other hospitals. 4) Some people retweeted old tweets to raise attention.5) Spreading some rumors of new cyberattacks occurred. 6) Some individuals take advantage of the cyberattacks and the later warnings to promote commercial ads for related products, such as training workshops to protect hospital IT systems. All these factors combined led to a spike in tweets a month after the attack.

After a week and no similar cyberattack occurred, the researchers noticed a drop in the concerned public interest in this cyberattack in week 44, as in Fig 4.

Fig. 5 shows how the number of positive, negative, and neutral sentiment labels changed over the study time. Even though initially there is a surge in the number of negative sentiments, the figure shows that the number of tweets classified as neural is more frequent than negative and positive ones in general. The lines took almost the bell shape in some parts of the chart, especially with the negative sentiment. The substantial number of neural tweets exceeding both the positive and the negative sentiment showed that many tweets were reporting the event without expressing any sentiment toward the cyberattacks.

Moreover, Fig. 5 shows the number of positive sentiment labels is always noticeably less frequent than the negative. This is because people usually conceive cyberattacks as harmful; the graph proves this fact. Surprisingly, more ransom cyberattacks targeted other institutions that provided vital services to society after the UHS cyberattack, such as Colonial Pipeline [35] [36]. More studies need to run to study the correlation between what seems to be people's optimism and what happened next in similar cyberattacks targeting more healthcare providers.

The researchers flat out the daily fluctuations in the number of positive, negative, and neutral sentiment labels by plotting them per week in Fig. 6. The figure enforces all the conclusions drawn from Fig. 5. In addition, Fig. 5 shows that as time passed, there was a decrease in the number of tweets. This is due to the public's lack of interest in this event and even almost forgotten.

The average numbers of positive, negative, and neutral sentiment score tweets of the UHS cyberattacks were computed and graphed in Fig. 7 to monitor how they changed over one month of the study as time went by. It might be because the owners of those positive tweets had confidence in the IT Personnel to overcome and get back to normal. This sentiment is rational since UHS resumed regular work shortly. However, it also could be interpreted as some people who tweet about the event might not realize the severity of the cyberattack because they are not professionals in computer

security and do not have an adequate background. Another explanation is that people were unaware of what happened after this cyberattack targeting UHS. Table I confirmed these results as it has the average weekly sentiment score of the three sentiment labels. It also reports the total average sentiment score.

A. Strength points

The researchers developed a model and workable process for collecting cybersecurity tweets and synthesizing them to remove irrelevant and/or duplicate tweets. They collected nearly one million tweets. They also use the recently developed Deep Learning RoBERTa model to obtain accurate sentiment scores of tweets. Finally, the authors plotted the measured sentiment to demonstrate how the public's sentiment changed over time toward a cyberattack incident.

B. Shortcoming or drawback

Although researchers did their best to refine the datasets, some of the gathered tweets may have been posted by bots. Furthermore, the analysis of sentiment in the study is based on the scores given by the deep learning model of RoBERTa, which may not be 100% perfect or even could be biased.

VI. Conclusion

This paper is the first to present the sentiment analysis of tweets related to the cybersecurity breach incident, aiming to study the public interests and views on social media when a cyberattack happened and how things changed afterward. Over three months, a large dataset of live-streaming and historical tweets related to cyberattacks was collected and analyzed. The cyberattacks on a group of UHS hospitals during the late days of September 2021 resulted in the misfunctioning of the computerized system. Then, the authors analyze the sentiment of the concerned public throughout three to measure how sentiment changes toward a vital and touching crisis of cyber attacking a group of hospitals during the outbreak of a global epidemic. Then, the authors did data analysis for the sentiment of the concerned public over one month to measure how sentiment changes toward a vital and touching crisis of cyber-attacking a group of hospitals during the outbreak of a global epidemic. The analysis of the tweets revealed people initially reacted with calm. However, as time went on, the interest of the concerned public faded toward security breaches of an important institution such as UHS. However, the study also revealed that people posted positive tweets about a critical security breach during stressful times. The study concluded that most people who cast their voices on Twitter trust people in charge of handling security breaches. The study is the first to survey Twitter's people and then measure their sentiment and how it changes during and after a cyberattack. Further analysis is still needed to know the lessons learned. The paper and the dataset collected by the authors can contribute to big data research and understanding of public perception.

References

[1] "UHS hospital chain hit with apparent ransomware attack," Healthcare IT News, Sep. 29, 2020. https://www.healthcareitnews.com/news/uhs-hospital-chain-hit-massive-ransomware-attack (accessed Feb. 18, 2021).

[2] M. Mansoor, K. Gurumurthy, A. R. U, and V. R. B. Prasad, "Global Sentiment Analysis Of COVID-19 Tweets Over Time," arXiv:2010.14234 [cs], Nov. 2020, Accessed: Feb. 23, 2021. [Online]. Available: http://arxiv.org/abs/2010.14234

[3] S. ES, "Sentiment Analysis in Python: TextBlob vs Vader Sentiment vs Flair vs Building It From Scratch," neptune.ai, Jul. 21, 2022. https://neptune.ai/blog/sentiment-analysis-python-textblob-vs-vader-vs-flair (accessed Dec. 25, 2022).

[4] F. Ferri, P. Grifoni, and T. Guzzo, Approaches, Tools and Applications for Sentiment Analysis Implementation.

[5] M. Farhadloo and E. Rolland, "Fundamentals of Sentiment Analysis and Its Applications," in Sentiment Analysis and Ontology Engineering: An Environment of Computational Intelligence, W. Pedrycz and S.-M. Chen, Eds. Cham: Springer International Publishing, 2016, pp. 1–24. doi: 10.1007/978-3-319-30319-2_1.

[6] M. Trupthi, S. Pabboju, and G. Narasimha, "Sentiment Analysis on Twitter Using Streaming API," in 2017 IEEE 7th International Advance Computing Conference (IACC), Hyderabad, India, Jan. 2017, pp. 915–919. doi: 10.1109/IACC.2017.0186.

[7] S. Bergsma, M. Dredze, B. Van Durme, T. Wilson, and D. Yarowsky, "Broadly Improving User Classification via Communication-Based Name and Location Clustering on Twitter," in Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Atlanta, Georgia, Jun. 2013, pp. 1010–1019. Accessed: Feb. 23, 2021. [Online]. Available: https://www.aclweb.org/anthology/N13-1121

[8] N. D. Zaki, N. Y. Hashim, Y. M. Mohialden, M. A. Mohammed, T. Sutikno, and A. H. Ali, "A real-time big data sentiment analysis for iraqi tweets using spark streaming," Bulletin EEI, vol. 9, no. 4, pp. 1411–1419, Aug. 2020, doi: 10.11591/eei.v9i4.1897.

[9] V. L. Meah, "Knowledge translation and social media: Twitter data analysis of the 2019 Canadian Guideline for Physical Activity throughout Pregnancy," Can J Public Health, p. 8.

[10] K. Bauman, A. Tuzhilin, and R. Zaczynski, "Using Social Sensors for Detecting Emergency Events: A Case of Power Outages in the Electrical Utility Industry," ACM Trans. Manage. Inf. Syst., vol. 8, no. 2–3, pp. 1–20, Aug. 2017, doi: 10.1145/3052931.

[11] T. Pano and R. Kashef, "A Complete VADER-Based Sentiment Analysis of Bitcoin (BTC) Tweets during the Era of COVID-19," BDCC, vol. 4, no. 4, p. 33, Nov. 2020, doi: 10.3390/bdcc4040033.

[12] B. White, "Sentiment Analysis: VADER or TextBlob?," Medium, May 27, 2020. https://towardsdatascience.com/sentiment-analysis-vader-or-textblob-ff25514ac540 (accessed Feb. 23, 2021).

[13] "Universal Health Services, Inc. | Healthcare Delivered with Passion," UHS. https://uhs.com/ (accessed Sep. 04, 2022).

[14] "Cyberattack hits major hospital system, possibly one of the largest in US history," NBC News. https://www.nbcnews.com/tech/security/cyberattack-hits-major-u-s-hospital-system-n1241254 (accessed Feb. 24, 2021).

[15] "Statement from Universal Health Services," UHS, Sep. 28, 2020. https://www.uhsinc.com/statement-from-universal-health-services/ (accessed Feb. 24, 2021).

[16] M. James, "FBI warns ransomware assault threatens US health care system: At least 5 hospitals have been hit this week," USA TODAY. https://www.usatoday.com/story/news/nation/2020/10/28/fbi-warns-ransomware-assault-threatens-us-healthcare-system/6065612002/ (accessed Mar. 09, 2021).

[17] D. J. Middaugh, "Cybersecurity Attacks during a Pandemic: It Is Not Just IT's Job!," p. 3.

[18] J. Klick, R. Koch, and T. Brandstetter, "Epidemic? The Attack Surface of German Hospitals during the COVID-19 Pandemic," arXiv:2101.07912 [cs], Jan. 2021, Accessed: Mar. 09, 2021. [Online]. Available: http://arxiv.org/abs/2101.07912

[19] F. Morstatter, J. Pfeffer, and H. Liu, "When is it biased? assessing the representativeness of twitter's streaming API," in Proceedings of the 23rd International Conference on World Wide Web, Seoul, Korea, Apr. 2014, pp. 555–556. doi: 10.1145/2567948.2576952.

[20] F. Morstatter, J. Pfeffer, H. Liu, and K. Carley, "Is the sample good enough? comparing data from twitter's streaming api with twitter's firehose," in Proceedings of the international AAAI conference on web and social media, 2013, vol. 7, no. 1, pp. 400–408.

[21] A. Khattri, A. Joshi, P. Bhattacharyya, and M. Carman, "Your Sentiment Precedes You: Using an author's historical tweets to predict sarcasm," in Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, Lisboa, Portugal, Sep. 2015, pp. 25–30. doi: 10.18653/v1/W15-2905.

[22] "Twitter IDs." https://developer.twitter.com/en/docs/twitter-ids (accessed Feb. 19, 2021).

[23] “re — Regular expression operations — Python 3.8.8 documentation.” https://docs.python.org/3.8/library/re.html (accessed Mar. 09, 2021).

[24] A. K. Dukare, “Data Cleaning for NLP of Social Media Text in 2 Simple Steps.,” Medium, May 18, 2020. https://towardsdatascience.com/data-cleaning-for-nlp-of-social-media-text-in-2-simple-steps-6ca48fa99c17 (accessed Mar. 09, 2021).

[25] H. Wachsmuth, “Text Analysis Pipelines,” in Text Analysis Pipelines: Towards Ad-hoc Large-Scale Text Mining, H. Wachsmuth, Ed. Cham: Springer International Publishing, 2015, pp. 19–53. doi: 10.1007/978-3-319-25741-9_2.

[26] “Pipelines.” https://huggingface.co/docs/transformers/main_classes/pipelines (accessed Oct. 11, 2022).

[27] T. Wolf et al., “Transformers: State-of-the-Art Natural Language Processing.” Association for Computational Linguistics, pp. 38–45, Oct. 2020. Accessed: Oct. 11, 2022. [Online]. Available: https://www.aclweb.org/anthology/2020.emnlp-demos.6

[28] A. Jaiswal, “All NLP tasks using Transformers Pipeline,” Analytics Vidhya, Dec. 27, 2021. https://www.analyticsvidhya.com/blog/2021/12/all-nlp-tasks-using-transformers-package/ (accessed Oct. 11, 2022).

[29] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” arXiv:1810.04805 [cs], May 2019, Accessed: Mar. 11, 2022. [Online]. Available: http://arxiv.org/abs/1810.04805

[30] “Overview of ROBERTa model,” GeeksforGeeks, Nov. 24, 2020. https://www.geeksforgeeks.org/overview-of-roberta-model/ (accessed Oct. 14, 2022).

[31] Y. Liu et al., “RoBERTa: A Robustly Optimized BERT Pretraining Approach.” arXiv, Jul. 26, 2019. doi: 10.48550/arXiv.1907.11692.

[32] “RoBERTa: An optimized method for pretraining self-supervised NLP systems.” https://ai.facebook.com/blog/roberta-an-optimized-method-for-pretraining-self-supervised-nlp-systems/ (accessed Oct. 12, 2022).

[33] F. Hamborg and K. Donnay, “NewsMTSC: A Dataset for (Multi-)Target-dependent Sentiment Classification in Political News Articles,” in Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, Online, Apr. 2021, pp. 1663–1675. doi: 10.18653/v1/2021.eacl-main.142.

[34] “cardiffnlp/twitter-roberta-base-sentiment · Hugging Face.” https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment (accessed Oct. 12, 2022).

[35] C. Bing, S. Kelly, C. Bing, and S. Kelly, “Cyber attack shuts down US fuel pipeline ‘jugular,’ Biden briefed,” Reuters, May 08, 2021. Accessed: Dec. 24, 2022. [Online]. Available: https://www.reuters.com/technology/colonial-pipeline-halts-all-pipeline-operations-after-cybersecurity-attack-2021-05-08/

[36] A. Hobbs, The Colonial Pipeline Hack: Exposing Vulnerabilities in US Cybersecurity. London, 2021. doi: 10.4135/9781529789768.