Brief communication

Mining social media data to inform public health policies: a sentiment analysis case study

Suzana N. Russell[1], Lila Rao-Graham[1], and Maurice McNaughton[1]

ABSTRACT

In the face of growing health challenges, nontraditional sources of data, such as open data, have the potential to transform how decisions are made and used to inform public health policies. Focusing on the COVID-19 pandemic, this article presents a case study employing sentiment analysis on unstructured social media data from Twitter (now X) to gauge public sentiment regarding pandemic-related restrictions. Our study aims to uncover and analyze Jamaican citizens' emotions and opinions surrounding COVID-19 restrictions following an outbreak at a call center in April 2020.

Machine learning sentiment analysis was used to analyze tweets from Twitter related to the lockdown. A total of 1 609 tweets were retrieved and analyzed, 76% of which expressed negative sentiments, suggesting that the majority of citizens were not in favor of the restrictions. The low compliance with the government-mandated policy may be related to the high percentage of negative sentiments expressed.

Insights from citizens' sentiments derived from open data sources such as Twitter can serve as valuable indicators for public health policymakers, providing critical input that will aid in tailoring interventions that align with public sentiments, thereby enhancing the effectiveness of and compliance with public health policies. This type of analysis can be useful to the health community and more generally to governments, as it allows for a more scientific assessment of public response to public health intervention techniques in real time. This study contributes to the emerging discourse on the integration of nontraditional data into public health policy-making, highlighting the growing potential for the use of these novel analytic techniques in addressing complex public health challenges.

Keywords COVID-19; quarantine; social media; sentiment analysis; data mining; health policy; Jamaica.

The advent of social media platforms, such as Facebook and Twitter (now X), has given rise to vast amounts of user-generated content that reflect real-time public opinions and experiences that can be used to inform public health policies. In the public health domain, public perception is critical because there is a strong association between risk perception and behaviors. Public health policies, particularly those implemented during a pandemic, will only be effective if they are adhered to by citizens. Therefore, it is important to gauge public sentiment so that appropriate policies and messaging can be implemented as means of increasing compliance.

Mining open data from social media platforms provides an opportunity to explore public sentiment using sentiment analysis. Open data are data that can be freely used, reused, and redistributed by anyone (1), and sentiment analysis is suited to unstructured data such as social media data. Sentiment analysis, also known as opinion mining, is a technique within the area of natural language processing, a field that is concerned with the interactions between computers and human languages and in particular with the programming of computers to process and analyze large amounts of natural language data (2). Sentiment analysis can be described as a method of understanding

[1] The University of the West Indies at Mona, Mona, Jamaica ✉ Suzana N. Russell, suzana.russell@uwimona.edu.jm

1

people's opinions, sentiments, attitudes, and emotions from written text (3). This type of analysis provides an instantaneous snapshot of the public's opinions and behavioral responses for a wide range of topics and social issues. Sentiment analysis from social media is not new and is already a widely researched subject, often used in business marketing to understand consumers' opinions toward a product.

Since the COVID-19 pandemic, research in the use of sentiment analysis in the health sector has increased significantly. For example, there has been an interest in understanding vaccination hesitancy by applying sentiment analysis to social media data (4). This study points out that policymakers could use the results to develop public health interventions that are responsive to the concerns of people who are hesitant to receive vaccines, as well as to develop public relations campaigns to encourage vaccination across the younger population. Sentiment analysis was also applied to Twitter feeds to understand public perceptions around health and care delivery in the United Kingdom as a result of COVID-19 (5). Even before the pandemic, researchers explored the potential of mining social network data to provide a tool for public health specialists and government decisionmakers to gauge the measure of concern expressed by Twitter users about public health issues (6).

In this study, we employ a case study methodology to characterize and analyze citizens' opinions on the government-mandated COVID-19 lockdown from 15 to 22 April 2020 in the parish of St. Catherine, Jamaica, arising from a rapid outbreak of cases traced to a call center operating in the city of Portmore. Because the data used in the study are open and in the public domain, research ethics approval was not required. The study uses sentiment analysis to analyze unstructured social media data obtained from Twitter to characterize public opinions on government interventions. Sentiment analysis is usually modeled as a classification problem, whereby a "classifier" is fed some text (e.g., a tweet) and returns a category (e.g., positive, negative, or neutral). As such, it usually requires building a classifier or using an existing pre-trained classifier related to the topic. For this study we built a classifier, using Python, because to our knowledge there were no existing classifiers related to this specific topic for Jamaica. Python was chosen as it is widely used in sentiment analysis due to its extensive libraries and tools that facilitate data collection, text processing, machine learning, and data analysis (7).

The workflow entailed three main steps. The first step involved using Python to collect tweets from Twitter. Twitter has several advantages over other social media platforms (like Facebook) that make it ideal for sentiment analysis, such as public availability of data (tweets are generally public by default); the extensive use of hashtags, which helps in categorizing and tracking topics, making it easier to filter and analyze sentiments on specific subjects or events; and an application programming interface (API) that is well-documented and widely used for data mining, unlike APIs such as Facebook that are more restricted, making it challenging for researchers to gather and analyze large datasets. We collected textual data, including posts, comments, and discussions using hashtags related to the St. Catherine lockdown and geofenced for only tweets originating within Jamaica, containing keywords or hashtags (#) including "Portmore lockdown," "St. Catherine lockdown," "Alorica," "call centre," and "outbreak."

After data retrieval, the next step involved preprocessing or cleaning of the data to remove irrelevant information, noise, and personally identifiable detail; e.g., stopwords, non-letters, punctuation, and usernames. The dataset was filtered to ensure that non-English content, such as emojis, was excluded. Emojis are often excluded from sentiment analysis because of the complexity and ambiguity in interpretation: a single emoji can be interpreted in multiple ways, making it challenging to assign a consistent sentiment value; encoding issues; and lack of standardization, which means there is no universally accepted standard for the sentiment classification of emojis (8).

The next step involved building the classifier, which used supervised machine learning and required three steps: training, validating, and testing (Figure 1). In the first step a set of training tweets was used to train the machine learning model. Here, the tweets were manually labeled as either positive or negative with the known sentiment (the grounded truth – GT).

The tweets were also classified into different content categories. For example, tweets classified as domestic movements refer to those that discussed the displacement that occurred within the interior of the country; government tweets refer to those that were addressed to or tagged a specific ministry or head of ministry and communication coming from a government entity; and health tweets discussed the toll of the lockdown on citizens' health, including mental health. Neutral and ambiguous tweets were not coded. These labels provided the input required to train the model.

The next step was validation, where the performance of the model was evaluated and fine-tuned. Here the classifier labeled a set of tweets and the results (predicted labels – P) compared to the grounded truths (the known sentiments) to determine the classifier's accuracy, measured as the proportion of correctly predicted labels out of the total labels. Various natural language processing techniques were used to determine which classification algorithm provided the highest accuracy. We chose the classifier that resulted in the highest accuracy (0.7874). An accuracy of 1 means the model perfectly classifies all positive and negative sentiments. After validating that the classifier had "learned" to accurately label tweets in terms of sentiment and content category, the final step was testing. At this stage the classifier was run on a set of independent, unlabeled tweets. This stage ensures that the classifier can generalize on unseen data (i.e., unlabeled data). It provides an unbiased estimation of how the model will perform in a real-world situation and must be done before the classifier is deployed in practice.

Building this classifier had an additional complexity because Jamaica has its own unique language, Jamaican Patois, that is used in informal contexts such as daily conversations, music, and social media posts. Most of the tweets in our dataset were written in Jamaican Patois and therefore, in order to establish the grounded truths to train the classifier, we had to manually translate these into standard English to facilitate the labeling of the tweets and to ensure the final classifier could account for Patois.

A total of 1 609 tweets were collected and used in the sentiment analysis. Some 76% of the tweets expressed negative sentiments, indicating that the majority of citizens were not in favor of the lockdown, even though it was seen by the government as a necessary measure to save lives. Table 1 shows the breakdown by content category and sentiment. Common themes from the analysis included concerns about personal

FIGURE 1. General overview of the classifier
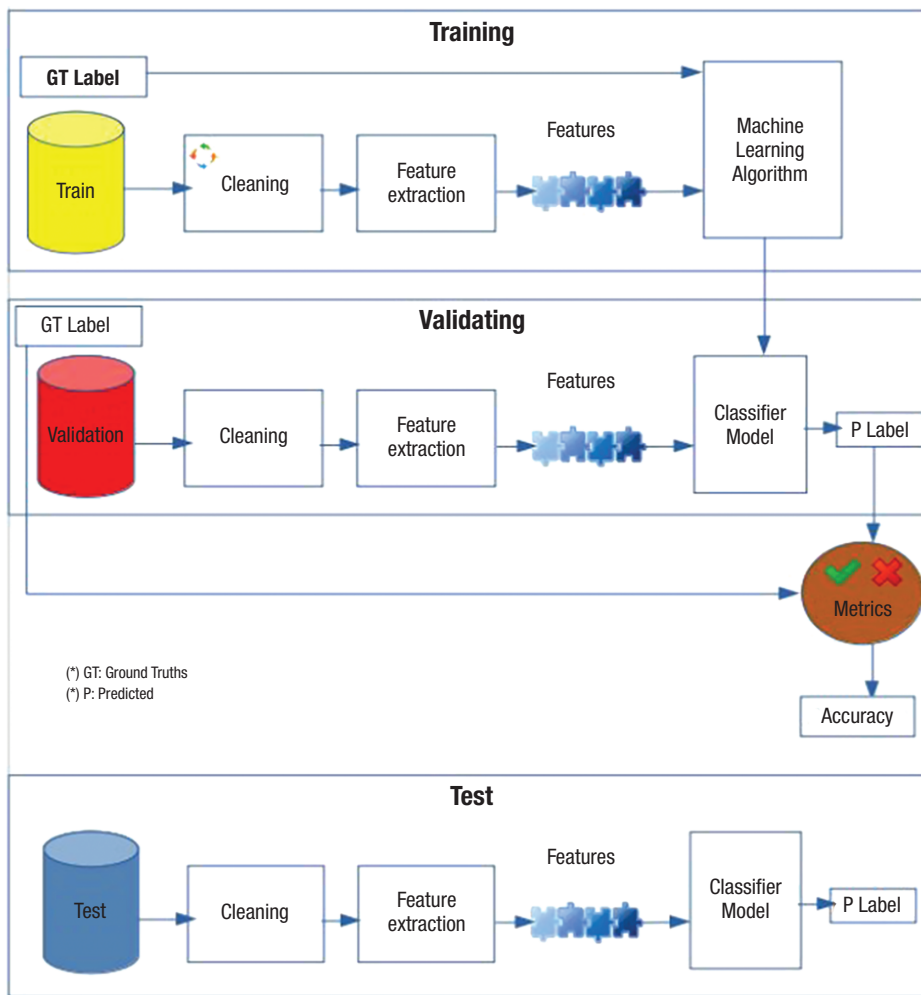
Source: Prepared by the authors.

freedoms, economic hardships, the impact on the provision of basic services, mental health, security, and education. The majority of the tweets (34%) focused on the impact of the lockdown on social and domestic movements leading to economic

TABLE 1. Breakdown of tweets by category and sentiment

Category	% of tweets	% classified as negative	% classified as positive
Social and domestic movements	34	29	5
Government	13	10	3
Commerce	13	11	2
Tourism/travel	12	5	7
Health	11	8	3
Transport	8	5	3
Basic services	3	3	0
Security	3	3	0
Economy	2	2	0
Education	1	1	0

Source: Prepared by the authors based on the study data.

hardships, while 1% focused on education. In general, the lockdown was discussed negatively across all categories, except for tourism/travel in which there were more positive than negative tweets on restricting incoming travel to the country from abroad. Citizens agreed that no one should be allowed to enter Jamaica while residents of St. Catherine were under lockdown.

During the one-week lockdown, the government struggled to get compliance from citizens as residents from St. Catherine crossed the northern border into neighboring parishes. The government resorted to calling in the police in an attempt to enforce compliance and stop residents from fleeing the lockdown. It is plausible to see the low compliance as being related to the overwhelming negative sentiments expressed by citizens.

This case study demonstrated that sentiment analysis of Twitter data can be used to understand public sentiments by capturing the diverse perspectives, emotions, and concerns of citizens and is a promising approach to inform evidence-based public health policies. The nuanced understanding of citizens' sentiments reflects their attitudes toward public health policies and may provide an indication of their potential compliance. Our findings suggest that the lockdown was viewed negatively

by the majority of the public, ultimately resulting in low compliance. Open data provide a more granular understanding of citizens' views in comparison to the traditional health data that are typically used in drafting public health policies. By monitoring Twitter conversations, public health policymakers will be able to readily identify emerging trends, concerns, and public reaction in real-time, allowing them to respond promptly and appropriately.

Qualitative analysis of sentiment data can help to identify key themes and concerns driving public perceptions. In our study, the 10 themes that emerged highlight areas of dissatisfaction, key concerns, and potential resistance to the lockdown and could have provided valuable insights for targeted intervention by the Jamaican government that could have enhanced public acceptance, thereby eliminating the need for police intervention to enforce compliance. Our findings underscore the importance of incorporating nontraditional data into real-time policy-making as a way of facilitating evidence-based decision-making that is responsive to public sentiments.

While sentiment analysis offers significant potential for developing countries like those in the Caribbean region, it is not without limitations. First, some tweets are difficult to categorize, such as sarcastic tweets. In this study the use of Jamaican Patois made some of the tweets even more ambiguous and difficult to label and demonstrates the challenges that can arise in sentiment analysis, as countries will need to navigate the day-to-day language colloquialisms. Large language models, like ChatGPT and MyAI on SnapChat app, are now able to understand different colloquial languages, including Jamaican Patois; however, the challenge of building a classifier to classify tweets written in non-standard English remains. Second, the small number of tweets extracted means the results may not be generalizable, and this limitation should be considered when interpreting the findings. In Jamaica only 4.5% of the population are Twitter users in comparison to 27.3% Facebook users (9). Notwithstanding, Twitter remains a simpler data source for this type of analysis due to the availability and transparency of the data. Another limitation of this approach is the inability to capture feedback from citizens without an online presence. This is particularly relevant for the Caribbean where access to the Internet remains a challenge for many. In 2023 Internet penetration in the Caribbean ranged from 27% to 97% across different countries and territories (10). Additionally, ethical issues including privacy protection and data anonymization need to be considered and handled in a robust manner in the undertaking of similar studies, to ensure compliance with the responsible use of Twitter data.

While age and gender analyses are possible in sentiment analysis, these were not considered in this study. It is largely acknowledged that analysis of such demographic data is challenging, as the process to ascertain gender and age of a user is quite complex. As a direction for future work, we recommend this type of analysis in order to develop more targeted public health messaging. Based on the low number of Twitter users, future studies should explore how to better use data from Facebook, which remains the most-used social media platform in the world.

Our case study underscores the potential of applying sentiment analysis to Twitter open data to provide timely, contextually rich insights into public sentiments. While the COVID-19 lockdown in Jamaica was used as the case study, using nontraditional data to complement traditional health data in public health policy-making is applicable to a range of applications in any Caribbean country; for example, vaccine mandates and healthcare delivery. Sentiment analysis can help to fill data gaps in a region that is often characterized as "data poor" and where the national statistical offices have difficulty in meeting the data needs of the region. This type of analysis can be extremely useful to the health community, and more generally to regional governments, as it allows for qualitative assessment of public response to public health intervention techniques in real time, thereby providing a critical input into planning to ensure public health policies are effective. These novel analytic techniques on open data can be used to complement the more traditional knowledge–attitudes–perceptions/practices (KAP) studies, which are robust and statistically representative, but costly.

Disclaimer. The opinions expressed in this manuscript are solely the authors' responsibility and do not necessarily reflect the views or policies of the *RPSP/PAJPH* or the Pan American Health Organization (PAHO).

REFERENCES

1. Open Knowledge Foundation. Open data handbook. London: OKF; [date unknown] [cited 4 June 2024]. Available from: https://opendatahandbook.org/.
2. Jurafsky D, Martin JH. Speech and Language Processing. 3rd edition, draft. [Stanford, CA]: 2024 [cited 2 June 2024]. Available from: https://web.stanford.edu/~jurafsky/slp3/.
3. Liu B. Introduction. In: Sentiment Analysis: Mining opinions, sentiments, and emotions. Studies in Natural Language Processing.

Cambridge: Cambridge University Press; 2020:1–17. https://doi.org/10.1017/9781108639286.002.
4. Griffith J, Marani H, Monkman H. COVID-19 vaccine hesitancy in Canada: content analysis of tweets using the theoretical domains framework. J Med Internet Res. 2021;23(4):e26874. https://doi.org/10.2196/26874.
5. Ainley E, Witwicki C, Tallett A, Graham C. Using Twitter comments to understand people's experiences of UK health care during

the COVID-19 pandemic: thematic and sentiment analysis. J Med Internet Res. 2021;23(10):e31101. https://doi.org/10.2196/31101.
6. Ji X, Chun SA, Wei Z, Geller J. Twitter sentiment classification for measuring public health concerns. Soc Netw Anal Min. 2015;5(1):13. https://doi.org/10.1007/s13278-015-0253-5.
7. Bird S, Klein E, Loper E. Natural language processing with Python. [Sebastopol, CA]: O'Reilly Media; 2009.
8. Bai Q, Dan Q, Mu Z, Yang M. A systematic review of emoji: current research and future perspectives. Front Psychol. 2019;10:2221. https://doi.org/10.3389%2Ffpsyg.2019.02221.
9. Statcounter GlobalStats. Social media stats Jamaica. [place unknown]: Statcounter; c2024 [cited 4 June 2024]. Available from: https://gs.statcounter.com/social-media-stats/all/jamaica.
10. Statista. Percentage of people online in Caribbean countries and territories as of January 2023. New York: Statista; 2023 [cited 4 June 2024]. Available from: https://www.statista.com/statistics/731275/internet-users-caribbean-countries/.

Minería de datos en las redes sociales para fundamentar las políticas de salud pública: estudio de caso de análisis de sentimientos

RESUMEN

Dados los crecientes desafíos en materia de salud, las fuentes de datos no tradicionales, como los datos de libre acceso (o datos abiertos), brindan la posibilidad de transformar la forma en la que se toman las decisiones y el modo como se usan para fundamentar las políticas de salud pública. En este artículo, centrado en la pandemia de COVID-19, se presenta un estudio de caso que emplea el análisis de sentimientos en los datos no estructurados de redes sociales procedentes de Twitter (ahora X) para evaluar los sentimientos del público ante las restricciones derivadas de la pandemia. Su objetivo es revelar y analizar las emociones y opiniones de los ciudadanos jamaicanos con respecto a las restricciones debidas a esta enfermedad tras un brote en un centro de llamadas en abril del 2020.

Se utilizó un aprendizaje automático para el análisis de sentimientos, a fin de analizar los tuits de Twitter relacionados con el confinamiento. Se recuperaron y analizaron un total de 1609 tuits, de los que el 76% manifestaba sentimientos negativos, lo que sugiere que la mayoría de los ciudadanos no era partidaria de las restricciones. Es posible que el bajo cumplimiento de la política impuesta por el gobierno esté relacionado con el porcentaje elevado de sentimientos negativos expresados.

Las perspectivas obtenidas respecto a los sentimientos de los ciudadanos a partir de fuentes de datos de libre acceso como Twitter pueden servir como indicadores útiles para los responsables de las políticas de salud pública, ya que proporcionan una información crucial para poder formular intervenciones adaptadas, acordes con los sentimientos de los ciudadanos, con la consiguiente mejora de la eficacia y el cumplimiento de las políticas de salud pública. Este tipo de análisis puede ser útil para la comunidad de atención de salud y, en términos más generales, para los gobiernos, ya que permite una evaluación más científica de la respuesta pública a las técnicas de intervención en materia de salud pública en tiempo real. Asimismo, este estudio constituye una aportación al discurso emergente sobre la integración de datos no tradicionales en la elaboración de políticas de salud pública al resaltar las posibilidades, cada vez mayores, que ofrece el uso de estas técnicas analíticas novedosas a la hora de abordar desafíos complejos en el ámbito de la salud pública.

Palabras clave

COVID-19; cuarentena; medios de comunicación sociales; análisis de sentimientos; minería de datos; política de salud; Jamaica.

Mineração de dados de redes sociais para subsidiar políticas de saúde pública: estudo de caso com análise de sentimentos

RESUMO

Diante dos desafios crescentes em saúde, fontes não tradicionais de dados (como dados abertos) têm o potencial de transformar a maneira como as decisões são tomadas e de serem aplicadas para subsidiar políticas de saúde pública. Trata-se de um estudo de caso que utiliza análise de sentimentos de dados não estruturados de redes sociais extraídos do Twitter (agora X), com enfoque na pandemia de COVID-19, com a finalidade de avaliar o sentimento do público em relação às restrições da pandemia. O objetivo foi expor e analisar as opiniões e as emoções da população jamaicana quanto às restrições impostas pela pandemia após um surto de COVID-19 em uma central de atendimento em abril de 2020.

Foi empregada uma análise de sentimentos usando aprendizado de máquinas para analisar textos postados no Twitter (tuítes) relacionados ao confinamento. Ao todo, foram identificados e analisados 1.609 tuítes. Desses, 76% expressavam sentimentos negativos, o que indica que a maioria não era favorável às restrições. A baixa adesão ao confinamento imposto pelas autoridades públicas pode estar relacionada ao elevado percentual de sentimentos negativos expressos pela população.

O entendimento adquirido a partir dos sentimentos da população com base em fontes de dados abertos, como o Twitter, pode ser uma contribuição útil para formuladores de políticas de saúde pública, fornecendo feedback essencial para adequar as intervenções aos sentimentos das pessoas, melhorando assim a efetividade e o cumprimento das políticas de saúde pública. Esse tipo de análise é vantajoso para a comunidade de saúde e, de modo geral, para os governos, porque permite analisar de maneira mais científica e em tempo real a resposta do público às estratégias de intervenção de saúde pública. Este estudo contribui para o discurso emergente de integrar dados não tradicionais à formulação de políticas de saúde pública, destacando o potencial cada vez maior de usar técnicas analíticas inovadoras para enfrentar os complexos desafios de saúde pública.

Palavras-chave

COVID-19; quarentena; mídias sociais; análise de sentimentos; mineração de dados; política de saúde; Jamaica.