

BRAIN. Broad Research in Artificial Intelligence and Neuroscience

e-ISSN: 2067-3957 | p-ISSN: 2068-0473

Covered in: Web of Science (ESCI); EBSCO; JERIH PLUS (hkdir.no); IndexCopernicus; Google Scholar; SHERPA/RoMEO; ArticleReach Direct; WorldCat; CrossRef; Peeref; Bridge of Knowledge (mostwiedzy.pl); abcdindex.com; Editage; Ingenta Connect Publication; OALib; scite.ai; Scholar9; Scientific and Technical Information Portal; FID Move; ADVANCED SCIENCES INDEX (European Science

Evaluation Center, neredataltics.org); ivySCI; exaly.com; Journal Selector Tool (letpub.com); Citefactor.org; fatcat!; ZDB catalogue; Catalogue SUDOC (abes.fr); OpenAlex; Wikidata; The ISSN Portal; Socolar; KVK-Volltitel (kit.edu)

2025, Volume 16, Issue 2, pages: 285-295

Submitted: April 5th, 2025 | Accepted for publication: May 24th, 2025

Unravelling the Power of Sentiment: How Emotions Shape Online Engagement in Public Health Discourse

Soufien Jaffali

Department of Management Information Systems and Production Management, College of Business and Economics, Qassim University. P.O. Box: 6640, Buraidah: 51452, Kingdom of Saudi Arabia.

s.jaffali@qu.edu.sa

<https://orcid.org/0000-0003-4612-0991>

Nesrine Khelifi

Department of Management Information Systems and Production Management, College of Business and Economics, Qassim University. P.O. Box: 6640, Buraidah: 51452, Kingdom of Saudi Arabia.

n.khelifi@qu.edu.sa

<https://orcid.org/0000-0002-7466-151X>

Abstract: *Social media platforms have evolved into powerful arenas for public discourse, significantly influencing opinions, behaviours, and policy decisions. This study explores the relationship between sentiment and user engagement in online conversations related to public health. Using advanced sentiment analysis techniques powered by the BERT model, we analyse how emotional tone affects interactions such as likes, replies, and shares. Our findings reveal that positive sentiments drive higher engagement in terms of likes and replies, while negative sentiments generate more shares, amplifying their reach. Neutral sentiments, in contrast, exhibit lower engagement across all metrics. Statistical analyses confirm significant differences in engagement based on sentiment, underscoring the critical role emotions play in shaping online conversations. These insights offer valuable guidance for policymakers, health communicators, and digital strategists aiming to craft more effective messaging strategies. By understanding how sentiment influences engagement, organisations can better connect with audiences, foster meaningful discussions, and enhance the impact of public communication campaigns.*

Keywords: *social media analysis; sentiment analysis; deep learning; natural language processing systems; vaccine acceptance; psychological factors.*

How to cite: Jaffali, S., & Khelifi, N. (2025). Unravelling the power of sentiment: How emotions shape online engagement in public health discourse. *BRAIN. Broad Research in Artificial Intelligence and Neuroscience*, 16(2), 285-295. <https://doi.org/10.70594/brain/16.2/21>



1. Introduction

The COVID-19 pandemic has presented an unprecedented global challenge that requires collective efforts to curb its spread and protect public health. Vaccination has emerged as a cornerstone of this response, playing a crucial role in mitigating the pandemic's impact. As immunisation campaigns continue worldwide, comprehending public sentiment and engagement toward vaccination is essential for crafting effective communication strategies, optimising outreach, and guiding public health interventions.

Social media platforms, particularly Twitter, have become influential arenas for public discourse, enabling users to share opinions, propagate information, and engage in discussions on critical topics such as COVID-19 and vaccination. These platforms provide a vast repository of real-time data, offering valuable insights into public perceptions, concerns, and attitudes toward vaccination efforts (Jaffali et al., 2019).

This study focuses on analysing vaccine-related tweets in Saudi Arabia, a country that took a proactive position in vaccine distribution. With its rapid and comprehensive immunisation program, Saudi Arabia became one of the first nations to ensure widespread vaccine access. Given this context, examining public sentiment and engagement in the country offers meaningful insights into the dynamics of vaccine acceptance and public discourse.

To achieve this, we leverage state-of-the-art natural language processing techniques, particularly the BERT transformer model, for sentiment analysis. BERT's ability to capture linguistic nuances and contextual depth makes it a powerful tool for analysing social media content. By categorising tweets into positive, negative, and neutral sentiments, we assess public perception and explore the affective determinants shaping vaccine-related discussions.

Beyond sentiment classification, we investigate user engagement metrics, including likes, retweets, and replies, to gauge the level of public interaction with vaccination discourse. By analysing the interaction between sentiment and engagement, we identify patterns that elucidate the role of emotional content in influencing online user behaviour.

The findings of this study hold significant implications for public health authorities, policymakers, and communication strategists, providing data-driven insights to optimise messaging strategies and intervention efforts. Ultimately, this research contributes to the broader goal of enhancing public health communication, fostering public trust, and promoting informed decision-making to improve vaccine acceptance and engagement.

2. Literature Review

Understanding user engagement on social media has become a focal point of research, with various studies examining the factors that drive interactions (Shahbaznezhad, Dolan, & Rashidirad, 2021). One key area of interest is the impact of emotional content. Weismueller et al. (2022) explore the effects of emotional appeal, perceived authority, and ideological alignment on the sharing of political content, revealing that posts evoking positive emotions are shared more frequently.

During the COVID-19 pandemic, digital interventions played a crucial role in mitigating the psychological distress among healthcare professionals. Blake et al. (2020) investigated the effectiveness of these interventions in providing targeted psychological support, emphasising the potential of technology-driven mental health solutions.

Artificial intelligence has also transformed healthcare, particularly in the personalisation of conversational agents. Kocaballi et al. (2019) examined how AI-driven interactions can enhance user engagement and improve patient outcomes, promoting a more patient-centric approach.

Social media engagement proved instrumental during public health crises, such as the COVID-19 pandemic. Zhang et al. (2022) analysed engagement patterns and found that interactions with government accounts varied across different stages of the crisis, underscoring the necessity of adaptive communication strategies.

Additionally, visual elements like emojis modulate user engagement. Ko, Kim, & Kim (2022) demonstrated how emotive visual cues enhance attention and foster greater interaction, offering valuable insights for digital marketing strategies.

Sentiment analysis has been a crucial tool in understanding public perceptions of COVID-19 vaccination. Ye et al. (2023) leveraged advanced analytical techniques to track sentiment variations across locations and periods, providing actionable insights for public health agencies.

Addressing vaccine hesitancy remains a persistent public health concern. Li, Wood, and Kostkova (2022) explored behaviour change theories and social media interventions to promote vaccine acceptance, offering strategic recommendations for policymakers.

Renuka & Radhakrishnan (2024) utilised a Kaggle-based dataset of manually annotated COVID-19 tweets to evaluate the effectiveness of the BERT model for sentiment classification, achieving an overall accuracy of 87% and outperforming traditional machine learning approaches like Multinomial Naive Bayes (MNB). Their study demonstrated BERT's strong performance in processing informal language, slang, and sarcasm prevalent on Twitter, particularly for COVID-19-related discussions. While their focus was primarily on improving classification accuracy and balancing performance metrics across sentiment classes, our research extends this line of inquiry by examining not only sentiment classification but also the relationship between emotional tones and user engagement metrics (likes, replies, retweets) in the context of Saudi Arabia's vaccination campaign. Furthermore, our work incorporates a regional lens by analysing sentiment geographically across major Saudi cities, offering insights into localised public health communication dynamics not addressed in previous general sentiment analyses.

In summary, research highlights the significant role of emotional content in shaping online engagement. Leveraging sophisticated sentiment analysis techniques can enhance our understanding of public discourse and inform the development of effective communication strategies.

3. Data Collection and Preprocessing

3.1 Data Collection Settings

We collected a dataset of tweets related to COVID-19 vaccination in Saudi Arabia using the TWINT API. The data collection focused on capturing public discourse from major Saudi cities during the country's vaccination campaign.

- **Geographic Scope:** The dataset includes tweets originating from major Saudi Arabian cities.
- **Time Frame:** Data was collected from December 17, 2020, to December 31, 2022, capturing discourse around the national vaccination rollout.
- **Keyword Selection:** The dataset was filtered based on specific vaccine-related terms, anti-vaccination phrases, and widely used hashtags.

The dataset focuses on tweets originating from major Saudi Arabian cities, specifically Riyadh, Jeddah, Mecca, Medina, Dammam, and Khobar. These cities were selected due to their high population densities, significant online activity, and strategic roles in the national vaccination campaign. The timeframe from December 2020 to December 2022 captures critical phases of the vaccination rollout, including initial public reactions, mid-campaign adjustments, and booster dose campaigns. Although the data spans two years, this range enables a comprehensive understanding of evolving public sentiment across different pandemic stages.

To comply with X (formerly Twitter) terms of service and developer agreements, which impose restrictions on data collection and sharing, we have implemented an anonymisation and data-sharing approach. Given that fully anonymising social media data is challenging due to its inherent searchability, we adhere to best practices by sharing only Tweet IDs and User IDs. This allows researchers to “re-hydrate” the dataset through X's API, ensuring ethical and compliant data

usage while maintaining reproducibility for academic research. The dataset is publicly available on Kaggle at the following link¹.

The dataset consists of 96 files, each corresponding to posts from a specific Saudi city. In total, it comprises 29,128 posts, offering a comprehensive snapshot of public discourse surrounding vaccination efforts during the COVID-19 pandemic.

Structured across 36 columns, the dataset captures a wide range of attributes, including:

conversation_id, created_at, date, time, timezone, user_id, username, name, place, tweet, language, mentions, URLs, photos, replies_count, reposts_count, likes_count, hashtags, cashtags, link, repost, quote_url, video, thumbnail, near, geo, source, user_rt_id, user_rt, repost_id, reply_to, repost_date, translate, trans_src, and trans_dest.

The dataset encompasses posts in 40 different languages (Figure 1). However, for methodological consistency, only tweets in Arabic and English—the most frequently used languages—were retained for analysis.




Figure 1. Language Distribution

3.2 Data Pre-processing

To enhance the quality and relevance of the dataset, a structured pre-processing workflow was implemented to filter and clean the tweets. Initially, the dataset comprised tweets in approximately 40 languages. However, for analytical consistency, only Arabic and English (the two most frequently used languages) were retained.

3.2.1 Language Filtering

The dataset originally contained tweets in over 40 languages, requiring a meticulous filtering process to ensure coherence and relevance. Given that Arabic and English accounted for approximately 70% of the dataset (Figure 1), they were prioritised for analysis. This selection aligns with their widespread use and significance in the study's context.

3.2.2 Tweet Cleaning

A standardised set of pre-processing steps was applied to both Arabic and English tweets to eliminate noise and elevate data quality:

- **Hashtag Removal:** All hashtags (words preceded by #) were removed, as they often contain nontextual characters or contextual markers that may introduce noise.

¹ <https://www.kaggle.com/datasets/jaffalisoufiene/covid-vaccine-tweets-from-saudi-arabia-20202022>

- **Hyperlink Elimination:** Any URLs (beginning with http or https) were excluded from the tweets, as they do not contribute to sentiment analysis and can be considered irrelevant.
- **Special Character Filtering:** Punctuation marks, emoticons, and other non-alphanumeric symbols that do not influence semantic meaning were removed.
- **Normalization of Repeated Letters:** Repeated characters within words were standardised to their base form. For example, 'vacciiiiine' was normalised to 'vaccine' to reduce text variability.
- **Lemmatization:** Words were reduced to their root or dictionary form, streamlining linguistic consistency and minimising redundancy.
- **Stop-word Removal:** Frequently occurring words such as articles, prepositions, and pronouns were eliminated, as they carry minimal semantic weight in sentiment analysis.

By implementing these pre-processing steps, the dataset was refined into a cleaner and more structured format, ensuring high-quality input for subsequent sentiment analysis. This process minimises irrelevant noise, enhances analytical accuracy, and facilitates more meaningful insights. After applying language filtering, text normalisation, and noise removal, the dataset was reduced from the initial 29,128 tweets to 19,980 tweets, ensuring that only relevant and high-quality data were retained for analysis.

4. BERT Transformer for Sentiment Analysis

4.1 Overview of NLP Techniques for Emotional Content Analysis

The analysis of emotional content in text has become a critical area of research within natural language processing (NLP), particularly for understanding public sentiment and affective reactions in social media data. Traditionally, sentiment analysis focuses on categorising text into three broad sentiment classes: positive, negative, and neutral. These early approaches relied heavily on lexicons or rule-based systems, which often lacked the flexibility to handle informal expressions, idioms, or context-dependent meanings common in user-generated content (Yang & Li, 2025).

However, recent advancements in NLP, particularly the development of transformer-based models, have revolutionised sentiment analysis. These models, especially the **Bidirectional Encoder Representations from Transformers (BERT)** model, can capture subtle emotional tones and complex contextual dependencies that older models typically miss (Saragih & Manurung, 2024).

BERT represents a significant leap forward due to its use of **bidirectional attention mechanisms**. Unlike traditional language models that read text in one direction (left-to-right or right-to-left), BERT processes each word in a sentence while attending to both its left and right context simultaneously. This means the model understands a word not just in isolation but in the full context of the sentence. For example, the word "shot" in the sentence "I am excited for my COVID-19 shot tomorrow" is interpreted differently than in "I am scared of getting the shot", thanks to BERT's contextual comprehension.

Moreover, BERT is pre-trained on a vast corpus of text using **masked language modelling (MLM)** and **next sentence prediction (NSP)**. In MLM, random words in a sentence are masked, and BERT is trained to predict them based on the surrounding words. In NSP, BERT learns to understand the relationship between sentence pairs, which enhances its ability to model coherent discourse.

These properties make BERT particularly well-suited for detecting emotional subtleties such as **sarcasm**, **negations**, and **ambiguous phrasing**, all of which are prevalent in social media data. As a result, BERT has become a state-of-the-art tool for many sentiment analysis tasks across industries, including public health, marketing, and political communication.

4.2 BERT's Application for Sentiment Analysis

BERT's advanced ability to interpret word meaning based on full sentence context makes it exceptionally well-suited for sentiment classification tasks. Unlike traditional machine learning classifiers, which often rely on keyword frequency or manually engineered features, BERT leverages a contextualised embedding of each word, capturing not only the word itself but also how its meaning changes depending on surrounding words.

For sentiment analysis, BERT is typically **fine-tuned on a labelled dataset**. This involves adding a **classification layer** on top of the pre-trained BERT architecture and training the model using sentiment-labeled text samples. During fine-tuning, the model learns to associate specific patterns in the text with sentiment classes (e.g., positive, negative, neutral) (Wu et al., 2024).

For example, when analysing a tweet like *"I can't believe how amazing this vaccine rollout is,"* BERT recognises that the sentiment is positive, despite the negation phrase *"can't believe,"* which may confuse simpler models. This makes BERT particularly powerful when applied to **real-world, noisy text** such as tweets, where abbreviations, emojis, and sentiment ambiguity are prevalent.

Additionally, **domain-specific fine-tuning** allows researchers to adapt BERT to the linguistic and emotional nuances of specialised fields, such as healthcare. In our study, BERT was fine-tuned on tweets related to COVID-19 vaccination in Saudi Arabia. This process allowed the model to better capture expressions, idioms, and tone unique to public health discourse within the sociocultural environment.

4.3 Applying BERT to Sentiment Analysis of Tweets

In this study, BERT was employed to analyse sentiment in tweets related to COVID-19 vaccination in Saudi Arabia. The dataset encompassed a diverse range of linguistic styles, including colloquial expressions, abbreviations, and emotional undertones, making it an ideal candidate for a context-aware NLP model.

To ensure optimal sentiment classification, we fine-tuned a pre-trained BERT model using a labelled sentiment dataset tailored to our research domain. This process allowed BERT to adapt to the unique characteristics of vaccination-related discussions on social media.

Once fine-tuned, the model was applied to classify tweets into three sentiment categories: positive, negative, and neutral. Table 1 presents sample tweets from the dataset along with their BERT-predicted sentiment labels, illustrating the model's ability to capture affective subtleties.

Table 1: Examples of tweets from each sentiment category

Sentiment Label	Example of tweets
Negative	"It's no secret where @... is... he's in an ICU on a ventilator after catching #COVID19! #VaccinesSaveLives #MasksWork"
Neutral	"First dose: Pfizer. Second dose: AstraZeneca. Third dose: Pfizer" "According to what I read, the third dose is only Pfizer or Moderna"
Positive	"So thankful my kids are vaccinated, masking, distancing, and washing their hands!" "The best way to end the year! I'm so thankful to have been able to receive my COVID-19 booster shot today!"

5. Results and Analysis

5.1. Sentiment Distribution

To better understand sentiment trends and user engagement in COVID-19 vaccination-related tweets, a descriptive analysis was conducted. The sentiment distribution in Figure 2 reveals that negative sentiment was the predominant class, representing 12,139 tweets. This was followed by positive sentiment with 4,501 tweets and neutral sentiment with 3,340 tweets.




Figure 2: Sentiment Distribution

In addition, engagement metrics, including likes, replies, and retweets, were examined to assess the interaction levels associated with different sentiment categories. Table 2 provides a comprehensive summary of these engagement statistics, including the count, mean, standard deviation, and distribution measures.

Table 2: Engagement Statistics

Measure	Likes	Replies	Retweets
Count	19,980	19,980	19,980
Non-zero Count	9,855	3,354	6,760
Mean	5.85	0.81	1.05
Std. Deviation	36.78	2.67	7.49
Maximum	1,019	77	341

The analysis revealed several key insights:

- The average number of likes per tweet was 5.85, with a standard deviation of 36.78, reflecting substantial variability in engagement.
- Replies averaged 0.81 per tweet, while retweets had a mean of 1.05, indicating relatively lower interaction compared to likes.
- Among all tweets, 9,855 tweets (49.3%) received at least one like, 3,354 tweets (16.8%) received at least one reply, and 6,760 tweets (33.8%) were retweeted at least once, demonstrating notable differences in engagement types.
- Maximum engagement levels varied significantly, with some tweets receiving up to 1,019 likes, 77 replies, and 341 retweets, highlighting the broad range of engagement across the dataset.

These insights offer valuable perspectives on how sentiment influences engagement, providing a foundation for further analysis of user behavior and public discourse on COVID-19 vaccination.

5.2. Relationship Between User Engagement and Sentiments

An analysis of the association between sentiment labels and user engagement metrics revealed notable trends (Table 3). The average number of retweets varied across sentiment categories, with negative sentiment tweets receiving the highest average retweets (1.39), followed by positive (0.54) and neutral tweets (0.52). However, interaction trends diverged when considering

likes and replies. Tweets with positive sentiment garnered the highest average likes (6.66) and replies (0.89), followed by negative sentiment tweets, which received an average of 6.54 likes and 0.80 replies. Neutral sentiment tweets exhibited the lowest engagement levels, averaging 2.24 likes and 0.75 replies per tweet.

Table 3: Average User Engagement Metrics by Sentiment Label

Sentiment	Avg. Likes	Avg. Replies	Avg. Retweets
Negative	6.54	0.80	1.39
Neutral	2.24	0.75	0.52
Positive	6.66	0.89	0.54

The observation that neutral sentiment tweets consistently receive lower engagement raises intriguing questions about the factors shaping user interactions on social media. While tweets conveying strong positive or negative emotions tend to attract more attention, the relatively muted response to neutral content suggests a potential gap in engagement strategies. This warrants further exploration into the cognitive-affective mechanisms that drive user interaction. Possible factors contributing to lower engagement with neutral tweets may include a lack of emotional resonance, reduced user interest, or the perceived informational value of the content. Understanding these dynamics could provide valuable insights for communication strategists seeking to optimise audience engagement in public health discourse.

To assess the statistical significance of differences in engagement across sentiment categories, an ANOVA test was performed on engagement metrics. As shown in Table 4, the p-value for likes (4.31e-09) suggests a highly significant difference in the average number of likes across sentiment categories, meaning that sentiment significantly predicts the number of likes a tweet receives. Similarly, the p-value for retweets (1.76e-14) is extremely low, indicating that sentiment significantly affects the likelihood of a tweet being shared, with negative tweets receiving the highest number of retweets. The p-value for replies (0.0446) is marginally significant, suggesting that sentiment may play a role in encouraging replies, but the effect is weaker than for likes or retweets. Overall, these results confirm that sentiment plays a crucial role in determining engagement levels on social media, particularly in terms of likes and retweets.

Table 4: ANOVA Test Results for Engagement Metrics Across Sentiment Categories

Engagement Metric	ANOVA p-value	Interpretation
Likes Count	4.31e-09	Highly significant ($p < 0.001$)
Replies Count	0.0446	Marginally significant ($p < 0.05$)
Retweets Count	1.76e-14	Highly significant ($p < 0.001$)

These findings highlight the predictive influence of sentiment in shaping user interactions on social media platforms. The results reveal that tweets with positive sentiment tend to attract more likes and replies, whereas negative sentiment tweets generate a higher number of retweets. These insights contribute to a deeper understanding of how sentiment drives user engagement and can inform strategic approaches to optimizing engagement levels, particularly in discussions related to COVID-19 vaccination on social media.

6. Interpretation

The findings of this analysis offer valuable insights into public health campaigns and immunisation strategies. Several key interpretations emerge from the results:

1. **Understanding sentiment and engagement:** The analysis highlights the relationship between sentiment and user engagement metrics, such as likes, replies, and retweets. While positive sentiment encourages direct interactions through likes and replies, negative sentiment spreads more widely via retweets, amplifying its

reach. This suggests that fostering a positive narrative around vaccination can enhance direct public engagement, while strategically addressing and reframing negative sentiment can expand message visibility. Public health campaigns should balance both approaches: encouraging positive engagement while mitigating misinformation and leveraging the viral nature of negative sentiment to drive awareness and constructive discourse.

2. **Tailored messaging strategies:** The findings underscore the need to customise communication strategies based on sentiment categories. Positive sentiment can be reinforced by sharing success stories and engaging with supportive voices, while neutral sentiment may require additional contextualisation to drive interest. For negative sentiment, public health campaigns should not only counter misinformation with factual information but also frame responses in ways that encourage discussion and trust-building, ensuring that widespread negative sentiment does not reinforce vaccine hesitancy but instead leads to productive conversations.
3. **Sentiment monitoring for proactive adjustments:** Regular sentiment analysis serves as an essential mechanism for monitoring public perceptions. By tracking sentiment trends, organisations can assess the impact of their messaging and refine their strategies accordingly. This proactive approach enables timely interventions to address emerging concerns and strengthen public trust.
4. **Collaboration and influencer engagement:** The correlation between sentiment and engagement underscores the need for collaboration among public health authorities, healthcare providers, and influential individuals or organisations (Centre for Tobacco Products, 2024). Engaging with trusted community influencers can help disseminate positive sentiment and boost public engagement. Social media platforms provide a powerful avenue for fostering effective communication and collaboration, allowing stakeholders to maximise outreach and impact (de Vere Hunt & Linos, 2022).
5. **Implications for public health strategies:** Incorporating sentiment analysis into public health initiatives can significantly enhance messaging effectiveness. By understanding public sentiment, campaigns can be tailored to resonate with the target audience, address concerns, and foster a more favourable attitude toward vaccination.

In summary, these findings reinforce the importance of integrating sentiment analysis into public health campaigns. Vaccination efforts can achieve greater reach and effectiveness by refining messaging strategies, addressing concerns, and fostering positive engagement.

6.1. Limitations

While this study provides valuable insights into the relationship between sentiment and user engagement in COVID-19 vaccination discussions on Twitter, several limitations should be acknowledged:

- **Data source and representativeness:** The dataset comprises tweets collected from Twitter, which may not fully reflect the sentiments and opinions of the broader population. Additionally, the focus on Saudi Arabia limits the generalizability of the findings to other regions or demographic groups.
- **Language and cultural variations:** The analysis primarily considers tweets in Arabic and English, overlooking sentiments expressed in other regional languages. Variations in language nuances, colloquial expressions, and cultural contexts may influence sentiment analysis results, potentially affecting interpretation.
- **Scope of user engagement metrics:** While this study examines likes, retweets, and replies, other crucial engagement indicators, such as shares or mentions, were not included due to data availability constraints. Future research should incorporate a

wider range of engagement metrics to offer a more comprehensive understanding of user behaviour on social media.

- **Inherent biases in social media data:** Social media data inherently contain biases influenced by user demographics, platform algorithms, and the nature of online discourse. Consequently, the sentiments expressed on Twitter may not always align with real-world attitudes and behaviours toward COVID-19 vaccination.
- **Interpretation of sentiment analysis results:** While sentiment analysis provides valuable insights, it may not capture the full complexity of human emotions and contextual nuances in text. Sentiment classification (positive, negative, neutral) remains inherently subjective and requires careful validation to ensure accuracy.

Recognising these limitations is crucial for accurately interpreting study findings and guiding future research to refine methodologies for analysing sentiment and user engagement in social media discourse.

7. Conclusion

In summary, this study leveraged advanced natural language processing techniques, specifically the BERT transformer model, to analyse sentiment in tweets related to COVID-19 vaccination in Saudi Arabia. The findings revealed a complex emotional landscape, with negative sentiments being the most prevalent. Notably, tweets expressing negative sentiment garnered higher user engagement, highlighting the intricate nature of public discourse. Additionally, an upward trend in engagement over time underscored the dynamic evolution of public sentiment and the critical need for timely, data-driven communication strategies.

Several key recommendations emerge from these insights. First, public health communication strategies must be tailored to address the widespread negative sentiment surrounding COVID-19 vaccination. This requires proactive dissemination of accurate information and targeted efforts to mitigate concerns. Second, collaboration between public health authorities, healthcare professionals, and influential figures is essential to fostering positive sentiment and driving engagement. Finally, continuous sentiment analysis using natural language processing can serve as a valuable tool for public health campaigns, enabling real-time adjustments and targeted interventions to address emerging public concerns effectively.

Looking ahead, future research should delve deeper into the factors driving engagement with negative sentiment to refine communication strategies further. Additionally, longitudinal studies tracking sentiment and engagement trends over extended periods would provide a more comprehensive understanding of shifting public perceptions. Comparative analyses across different regions, vaccination campaigns, or public health crises could offer further insights into variations in sentiment dynamics, informing more effective public health outreach efforts.

References

- Blake, H., Bermingham, F., Johnson, G., & Tabner, A. (2020). Mitigating the psychological impact of COVID-19 on healthcare workers: A digital learning package. *International Journal of Environmental Research and Public Health*, 17(9), 2997. <https://doi.org/10.3390/ijerph17092997>
- Center for Tobacco Products. (2024). *Influencers 101: Best practices for public health campaigns*. Retrieved from <https://digitalmedia.hhs.gov/tobacco/hosted/CTP-Influencers-101-JUN23.pdf>
- de Vere Hunt, I., & Linos, E. (2022). Social media for public health: Framework for social media-based public health campaigns. *Journal of Medical Internet Research*, 24(12), e42179. <https://doi.org/10.2196/42179>
- Jaffali, S., Jamoussi, S., Khelifi, N., & Ben Hamadou, A. (2019). Survey on social networks data analysis. In *International Conference on Innovations for Community Services* (pp. 100–119). Springer. https://doi.org/10.1007/978-3-030-37484-6_6

- Kocaballi, A., Berkovsky, S., Quiroz, J., Laranjo, L., Tong, H., Rezazadegan, D., Briatore, A., & Coiera, E. (2019). The personalization of conversational agents in health care: Systematic review. *Journal of Medical Internet Research*, 21(11), e15360. <https://doi.org/10.2196/15360>
- Ko, E., Kim, D., & Kim, G. (2022). Influence of emojis on user engagement in brand-related user-generated content. *Computers in Human Behavior*, 136, 107387. <https://doi.org/10.1016/j.chb.2022.107387>
- Li, L., Wood, C., & Kostkova, P. (2022). Vaccine hesitancy and behavior change theory-based social media interventions: A systematic review. *Translational Behavioral Medicine*, 12(2), 243–272. <https://doi.org/10.1093/tbm/ibab148>
- Renuka, O., & Radhakrishnan, N. (2024). BERT for Twitter sentiment analysis: Achieving high accuracy and balanced performance. *Journal of Trends in Computer Science and Smart Technology*, 6(1), 37–50. <https://doi.org/10.36548/jtcsst.2024.1.003>
- Saragih, H., & Manurung, J. (2024). Leveraging the BERT model for enhanced sentiment analysis in multicontextual social media content. *Jurnal Teknik Informatika C.I.T Medicom*, 16(2), 82–89. <https://doi.org/10.35335/cit.Vol16.2024.766.pp82-89>
- Shahbaznezhad, H., Dolan, R., & Rashidirad, M. (2021). The role of social media content format and platform in users' engagement behavior. *Journal of Interactive Marketing*, 53, 47–65. <https://doi.org/10.1016/j.intmar.2020.05.001>
- Weismueller, J., Harrigan, P., Coussement, K., & Tessitore, T. (2022). What makes people share political content on social media? The role of emotion, authority, and ideology. *Computers in Human Behavior*, 129, 107150. <https://doi.org/10.1016/j.chb.2021.107150>
- Wu, Y., Jin, Z., Shi, C., Liang, P., & Zhan, T. (2024). Research on the application of deep learning-based BERT model in sentiment analysis. *arXiv preprint*, arXiv:2403.08217. <https://doi.org/10.48550/arXiv.2403.08217>
- Yang, X., & Li, G. (2025). Psychological and behavioral insights from social media users: Natural language processing-based quantitative study on mental well-being. *JMIR Formative Research*, 9, e60286. <https://doi.org/10.2196/60286>
- Ye, J., Hai, J., Wang, Z., Wei, C., & Song, J. (2023). Leveraging natural language processing and geospatial time series model to analyze COVID-19 vaccination sentiment dynamics on Tweets. *JAMIA Open*, 6(1), ooad023. <https://doi.org/10.1093/jamiaopen/ooad023>
- Zhang, W., Yuan, H., Zhu, C., Chen, Q., & Evans, R. (2022). Does citizen engagement with government social media accounts differ during the different stages of public health crises? An empirical examination of the COVID-19 pandemic. *Frontiers in Public Health*, 10, 807459. <https://doi.org/10.3389/fpubh.2022.807459>