InSight Bulletin: A Multidisciplinary Interlink International Research Journal
Peer Reviewed International, Open Access Journal.
ISSN : 3065-7857 / Website: https://ibrj.us / Volume-3, Issue-2 / February - 2026

Original Article

From Tweets to Insights: Detecting Anxiety Trends in Twitter Data through NLP

Hrishikesh Jadhav[1], Saurabh Singh[2], Dr. Manjusha Patil[3], Dr. Vinaya Keskar[4]
[1,2]ATSS College of Business Studies and Computer Application
[3,4]Assistant Professor, ATSS College of Business Studies and Computer Application

Manuscript ID:
IBMIIRJ -2026-030245

Submitted: 10 Jan. 2026

Revised: 15 Jan. 2026

Accepted: 18 Feb. 2026

Published: 28 Feb. 2026

ISSN: 3065-7857

Volume-3

Issue-2

Pp. 216-222

February 2026

Correspondence Address:
Hrishikesh Jadhav
ATSS College of Business Studies and Computer Application
Email: hjadhav0109@gmail.com

OPEN ACCESS

Quick Response Code:

Web. https://ibrj.us

DOI: 10.5281/zenodo.21100792

DOI Link:
https://doi.org/10.5281/zenodo.21100792

Abstract

The rapid growth of social media platforms has transformed the way individuals express emotions and psychological states in the digital space. Twitter, in particular, serves as a real-time reflection of public sentiment, making it a valuable resource for understanding mental health trends such as anxiety. This research explores the application of Natural Language Processing (NLP) techniques to detect and analyze anxiety-related expressions in Twitter data. The study aims to build an automated system capable of classifying tweets as anxiety-related or non-anxiety-related and to examine how anxiety trends vary over time. A dataset of tweets was collected using the Twitter API based on anxiety-related keywords and hashtags. The data underwent preprocessing steps including text normalization, removal of noise, and tokenization. Both traditional machine learning models (Logistic Regression and Support Vector Machine) and deep learning approaches (fine-tuned BERT model) were employed for classification. Performance was evaluated using metrics such as accuracy, precision, recall, and F1-score. Additionally, temporal trend analysis was conducted to identify fluctuations in anxiety mentions over specific periods. The results demonstrate that transformer-based models achieved higher accuracy and F1-scores compared to classical approaches, indicating their effectiveness in capturing contextual emotional cues. This research highlights the potential of NLP-driven systems in supporting large-scale mental health monitoring and providing insights into societal emotional trends while emphasizing ethical considerations and privacy protection.

Keywords: *Anxiety Detection, Twitter Data, Natural Language Processing (NLP) Sentiment Analysis, Machine Learning, Deep Learning, BERT, Mental Health Monitoring.*

Introduction

The rapid growth of digital communication technologies has fundamentally transformed the ways in which individuals express thoughts, emotions, and personal experiences in contemporary society. Among these technologies, social media platforms have emerged as dominant spaces for everyday communication, enabling users to share opinions, emotions, and life events instantly with a global audience. Platforms such as Twitter (now X), Facebook, Instagram, and Reddit allow users to produce and consume user-generated content in real time, creating massive volumes of textual data that reflect social behavior and emotional states. This unprecedented availability of public textual data has opened new opportunities for researchers to study human emotions, opinions, and mental health patterns at a scale that was previously unattainable using traditional research methods. Twitter, in particular, has become a valuable source of data for psychological and social research due to its short-form, real-time nature and public accessibility. Users frequently post brief messages, commonly known as tweets, to express immediate reactions to personal experiences, social interactions, academic pressures, workplace challenges, political events, and global crises. These spontaneous expressions often reveal emotional and psychological states, including stress, fear, frustration, and anxiety. As a result, Twitter data provides a rich and continuously evolving dataset for understanding collective emotional trends and mental health concerns across different populations and time periods. Anxiety is one of the most prevalent mental health challenges worldwide and has become increasingly prominent in modern society.

How to cite this article:
Jadhav, H., Singh, S., Patil, M., & Keskar, V. (2026). From Tweets to Insights: Detecting Anxiety Trends in Twitter Data through NLP. InSight Bulletin: A Multidisciplinary Interlink International Research Journal, 3(2), 216–222. https://doi.org/10.5281/zenodo.21100792

It is influenced by a wide range of factors, including academic demands, professional stress, financial instability, social expectations, relationship issues, health concerns, and exposure to global events such as pandemics, conflicts, and economic uncertainty. While anxiety disorders are clinically defined and diagnosed by mental health professionals, many individuals experience anxiety-related symptoms without seeking formal diagnosis or treatment. Instead, people often turn to social media platforms to share feelings of worry, stress, nervousness, emotional exhaustion, and overwhelm as a form of self-expression or emotional release. These publicly shared experiences provide valuable insights into how anxiety manifests at a societal level. The analysis of anxiety-related expressions on social media offers an alternative and complementary approach to traditional mental health research methods, such as surveys, interviews, and clinical assessments. Conventional approaches often rely on self-reported data collected from limited sample sizes and at specific points in time, which may not fully capture the dynamic and evolving nature of emotional states. In contrast, social media data allows researchers to observe emotional expressions continuously and at a large scale, enabling the identification of trends, patterns, and fluctuations in anxiety over time. This real-time aspect is particularly important for understanding how external events, such as examinations, job market changes, public health emergencies, or global crises, influence collective emotional well-being. Natural Language Processing (NLP), a subfield of artificial intelligence and computational linguistics, plays a crucial role in enabling the automated analysis of large-scale textual data generated on social media platforms. NLP techniques allow computers to process, analyze, and interpret human language by converting unstructured text into structured representations that can be analyzed computationally. Through techniques such as tokenization, stop-word removal, stemming, lemmatization, and vectorization, raw text data can be transformed into meaningful features that capture linguistic and semantic patterns. When combined with machine learning and deep learning algorithms, NLP enables the detection and classification of emotional and psychological states expressed in text.

In recent years, advances in machine learning and deep learning have significantly improved the accuracy and effectiveness of text classification tasks. Traditional machine learning models such as Naïve Bayes, Support Vector Machines (SVM), and Logistic Regression have been widely used for sentiment and emotion analysis. More recently, deep learning architectures, including Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), Long Short-Term Memory (LSTM) networks, and transformer-based models, have demonstrated superior performance in capturing contextual and semantic information in textual data. These models are particularly well-suited for identifying subtle linguistic cues associated with anxiety, such as expressions of uncertainty, fear, rumination, and emotional distress. Applying NLP-based techniques to Twitter data enables the automatic identification of anxiety-related content without requiring manual annotation of every tweet. This automation makes it possible to analyze millions of tweets efficiently, providing insights into population-level anxiety trends that would be infeasible through manual methods alone. Furthermore, NLP-based social media analysis is cost-effective and scalable, making it an attractive approach for large-scale mental health research. By analyzing patterns in language use, researchers can identify recurring themes, emotional markers, and temporal variations associated with anxiety-related expressions. The present research focuses on detecting anxiety-related expressions from Twitter data using NLP-based approaches. The study involves collecting tweets that are relevant to anxiety-related discussions, preprocessing the textual data to remove noise and enhance quality, and applying both machine learning and deep learning models to classify tweets as either anxiety-related or non-anxiety-related. In addition to classification, the research examines temporal variations in the frequency of anxiety-related tweets to understand how expressions of anxiety change over time. This temporal analysis helps reveal patterns associated with specific periods, events, or circumstances that may contribute to increased or decreased levels of anxiety in the digital space. It is important to emphasize that the primary goal of this research is not to diagnose mental health conditions or identify individuals with anxiety disorders. Instead, the study aims to analyze linguistic patterns and emotional expressions at a population level to gain insights into collective anxiety trends. Ethical considerations play a central role in this research, including the responsible use of publicly available data, the protection of user privacy, and the avoidance of harmful interpretations or misuse of findings. By focusing on aggregate trends rather than individual-level diagnosis, the research aligns with ethical guidelines for social media and mental health research. In summary, the increasing use of social media platforms has created new opportunities for understanding mental health patterns through large-scale textual analysis. Twitter serves as a valuable data source for capturing real-time expressions of anxiety, while NLP-based techniques enable the automated identification and analysis of these expressions. By combining social media data with machine learning and deep learning methods, this research contributes to the growing field of computational mental health and provides insights into how anxiety is expressed and evolves in the digital era. The findings of this study have the potential to inform future research, public health monitoring efforts, and the development of supportive interventions aimed at improving collective emotional well-being.

Problem Statement

Anxiety is a widespread mental health issue that affects people of all ages, yet traditional methods of monitoring anxiety trends, such as surveys and clinical assessments, are

Time-consuming and often limited in scale. Meanwhile, social media platforms like Twitter provide a large volume of real-time textual data where users frequently share their feelings, including signs of stress and anxiety. However, manually analyzing such massive and unstructured data is practically impossible.

The challenge lies in automatically detecting anxiety-related content in tweets and understanding how these expressions fluctuate over time. Tweets often contain informal language, slang, abbreviations, emojis, and context-dependent meanings, which makes accurate detection difficult. Furthermore, the system must respect user privacy and avoid making clinical diagnoses while still providing meaningful insights into collective anxiety trends.

Therefore, the problem this research addresses is:

"How can Natural Language Processing techniques be applied to Twitter data to automatically identify anxiety-related expressions and analyze their temporal trends, without violating ethical or privacy considerations?"

Objectives of Study

The main objectives of this research are as follows:

1. **Data Collection:** To gather a dataset of tweets containing anxiety-related keywords and hashtags using the Twitter API, ensuring sufficient coverage for meaningful analysis.
2. **Text Preprocessing:** To clean and preprocess the collected textual data by removing noise, normalizing text, and preparing it for computational analysis.
3. **Anxiety Detection:** To develop machine learning and deep learning models capable of automatically classifying tweets as anxiety-related or non-anxiety- related.
4. **Model Evaluation:** To assess the performance of the classification models using metrics such as accuracy, precision, recall, and F1-score, ensuring reliable detection results.
5. **Trend Analysis:** To examine temporal patterns and fluctuations in anxiety- related tweets over time, identifying potential correlations with events or societal factors.
6. **Ethical and Responsible Usage:** To conduct the study while maintaining user privacy, ensuring that no personal information is exposed, and avoiding any clinical interpretation of the data.

Literature Review

The analysis of mental health indicators using social media data has gained significant attention in recent years due to the widespread use of online platforms for emotional expression. Social media platforms such as Twitter, Reddit, and Facebook generate large volumes of real-time textual data that reflect users' thoughts, emotions, and psychological states. Unlike traditional survey-based or clinical approaches, social media provides continuous, large-scale, and cost-effective data for understanding public sentiment and mental health trends. Several studies have demonstrated that linguistic patterns observed in social media posts can serve as meaningful signals for identifying mental health concerns at a population level (De Choudhury et al., 2013; Coppersmith et al., 2014).

Early research in this domain primarily focused on sentiment analysis, which classifies text into positive, negative, or neutral categories. While sentiment analysis proved useful for assessing general emotional polarity, researchers found it insufficient for detecting specific emotional states such as anxiety or stress. To overcome this limitation, emotion detection techniques were introduced to identify finer-grained emotional categories. Lexicon-based approaches utilized predefined emotion dictionaries to detect anxiety-related terms such as fear, worry, or stress (Ortu et al., 2016). However, these approaches were limited by their inability to capture contextual meaning. As a result, machine learning-based emotion detection methods were explored, enabling models to learn complex linguistic patterns associated with anxiety expressions (Abdul-Mageed & Ungar, 2017; Shatte et al., 2019). The detection of anxiety in social media content has emerged as a key focus within computational mental health research. Initial studies employed traditional machine learning classifiers such as Logistic Regression, Support Vector Machines, and Random Forests, using features like bag-of-words, TF-IDF, and n-grams to represent textual data (Resnik et al., 2015; Guntuku et al., 2017). These methods achieved moderate success but struggled with capturing long-range dependencies and subtle emotional cues present in short social media posts. Despite these limitations, such approaches established a foundational framework for automated anxiety detection and highlighted the feasibility of social media-based mental health analysis. With advancements in deep learning, more sophisticated models were introduced to improve anxiety detection performance. Neural network architectures such as Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks demonstrated enhanced ability to capture sequential and contextual relationships between words in text (Yates et al., 2017; Trotzek et al., 2018). These models significantly outperformed traditional classifiers by learning deeper semantic representations of language. Further improvements were achieved with the introduction of transformer-based models, particularly BERT (Bidirectional Encoder Representations from Transformers), which considers bidirectional context and captures nuanced emotional expressions more effectively (Devlin et al., 2019). Several studies reported that fine-tuned BERT models achieved higher accuracy and F1-scores in anxiety and emotion classification tasks compared to classical machine learning and recurrent neural network models (Ji et al., 2020; Sawhney et al., 2021). In addition to classification, researchers have explored temporal analysis to understand how anxiety-related expressions change over time. Studies analyzing longitudinal Twitter data observed fluctuations in anxiety mentions corresponding to major societal events such as pandemics, examinations, and economic uncertainty (Paul & Dredze, 2017; Saha et al., 2022). These findings suggest that social media can act as an early indicator of collective emotional responses and provide insights into public mental health trends. Temporal modeling approaches have further highlighted the importance of time-aware analysis in understanding emotional dynamics at scale. Despite promising outcomes, several challenges remain in social media-based anxiety detection. Social media text is often informal and noisy, containing abbreviations, emojis, hashtags, spelling errors, and multilingual code-switching, which complicate automated analysis (Zogan et al., 2021). Additionally, anxiety-related posts are typically underrepresented compared to neutral content, leading to class imbalance issues that can bias model performance (Ríssola et al., 2020). Ethical concerns related to user privacy, consent, and responsible data usage have also been emphasized, particularly when dealing with sensitive mental health information (Matero et al., 2019). Overall, existing literature demonstrates the effectiveness of NLP and deep learning techniques for detecting anxiety-related expressions in social media data. However, there remains a need for further research that combines transformer-based models with temporal trend analysis while addressing ethical considerations. This research aims to build upon prior work by leveraging advanced NLP techniques to analyze anxiety expressions on Twitter in a responsible and population-focused manner.

Methodology

The methodology of this research is designed to systematically collect, process, and analyze Twitter data to detect anxiety-related expressions using Natural Language Processing (NLP) techniques. The process involves four key stages: data collection, data preprocessing, NLP-based feature extraction, and model development for classification.

a. **Data Collection:** The dataset for this study consists of tweets collected from Twitter using the Twitter API v2. Tweets were selected based on anxiety-related keywords and hashtags such as *anxiety*, *stressed*, *panic*, *overwhelmed*, *nervous*, and their common variants. Retweets were excluded to ensure that the dataset contained original user expressions. To maintain ethical standards and privacy, no personal identifiers or user handles were stored. The data collection spanned a period of several months to ensure sufficient temporal coverage, enabling trend analysis over time.

b. **Data Preprocessing:** Raw social media data is often noisy, informal, and unstructured. To prepare the tweets for analysis, the following preprocessing steps were applied:

1. **Text Normalization:** Conversion of all text to lowercase to ensure uniformity.
2. **Removal of Noise:** Elimination of URLs, mentions (@username), hashtags (optional), punctuation, and special characters.
3. **Tokenization:** Breaking down sentences into individual words or tokens for computational analysis.
4. **Stopword Removal:** Common words such as *the*, *is*, *and*, *in* were removed to focus on meaningful content.
5. **Lemmatization:** Words were reduced to their base or dictionary form (e.g., *running* → *run*) to consolidate variations of the same word.
6. **Handling Emojis and Slang:** Emojis were either removed or converted to text representations to preserve emotional information, and common slang was normalized for consistency.

These preprocessing steps ensure that the data is clean, structured, and suitable for both traditional machine learning and deep learning approaches.

c. **NLP Techniques Used:** To convert text into a numerical format that models can understand, different NLP techniques were employed:

- **Bag-of-Words (BoW) and TF-IDF:** These traditional methods represent text based on word frequency and importance, serving as input for classical machine learning models like Logistic Regression and Support Vector Machines (SVM).
- **Word Embeddings:** Pretrained embeddings such as Word2Vec and GloVe were used to capture semantic relationships between words. This helps models understand context beyond simple frequency counts.
- **Contextual Embeddings:** Transformer-based embeddings from BERT were utilized to capture contextual nuances of tweets, allowing the model to understand how word meanings change depending on surrounding words.

d. **Multiple models were developed and evaluated for classifying tweets:**

1. **Traditional Machine Learning Models:**
   - Logistic Regression (LR): Used for baseline classification due to its simplicity and interpretability.
   - Support Vector Machine (SVM): Applied to separate anxiety and non-anxiety tweets in a high-dimensional feature space.
2. **Deep Learning Models:**
   - LSTM (Long Short-Term Memory): Captures sequential dependencies in text, allowing understanding of the order and context of words.
   - CNN (Convolutional Neural Network): Extracts local patterns and n-gram features in text sequences.
3. **Transformer-based Model:**
   - **BERT (Bidirectional Encoder Representations from Transformers):**

Fine-tuned for anxiety classification, this model considers both left and right context of words, effectively capturing subtle emotional cues in tweets.

Each model was trained on a labeled dataset split into training, validation, and test sets. Performance metrics including accuracy, precision, recall, and F1-score were used to evaluate and compare models. The best-performing model was then employed for temporal trend analysis to understand fluctuations in anxiety expressions over time.

Tools & Technologies Used

This research leverages a combination of programming languages, libraries, frameworks, and platforms to collect, process, and analyze Twitter data, as well as to implement machine learning and deep learning models for anxiety detection.

Programming Language

- Python: Python was chosen due to its extensive support for data processing, NLP, and machine learning. Its simplicity and large community make it suitable for rapid prototyping and experimentation.

Data Collection

- Tweepy / Twitter API v2: Tweepy, a Python library, was used to access Twitter's API, enabling programmatic collection of tweets containing anxiety- related keywords and hashtags. The API ensured compliance with Twitter's usage policies and ethical standards.

Data Preprocessing and NLP

- NLTK (Natural Language Toolkit): Used for tokenization, stopword removal, and lemmatization.
- SpaCy: Employed for advanced text preprocessing tasks such as named entity recognition and efficient tokenization.
- Regular Expressions (Regex): Utilized for cleaning text by removing URLs, mentions, special characters, and unwanted symbols.

Machine Learning and Deep Learning

- Scikit-learn: Implemented traditional machine learning models such as Logistic Regression and Support Vector Machines, along with metrics for performance evaluation.
- TensorFlow / Keras: Used to build and train deep learning models, including LSTM and CNN architectures.
- Hugging Face Transformers: Enabled the use of transformer-based models such as BERT for contextual embedding and fine-tuning on the anxiety detection task.

Visualization and Analysis

- Matplotlib & Seaborn: Used for plotting temporal trends, frequency distributions, and comparison of model performance.
- WordCloud: Generated word clouds to visualize common anxiety-related terms in tweets.

Development Environment

- Jupyter Notebook: Provided an interactive environment for coding, testing, and visualizing results.
- Google Colab (Optional): Used for training resource-intensive models such as BERT when local hardware limitations exist.

Experimental Setup

The experimental setup defines how the collected Twitter dataset was prepared, how the models were trained and evaluated, and the performance metrics used to assess their effectiveness in detecting anxiety-related tweets.

Dataset Preparation

The dataset was divided into three subsets:

1. Training Set (70%): Used to train the models.
2. Validation Set (15%): Used for tuning hyperparameters and preventing overfitting.
3. Test Set (15%): Used to evaluate the final performance of the trained models on unseen data.

The distribution of anxiety-related and non-anxiety tweets was monitored to ensure a balanced representation, and techniques such as oversampling or undersampling were applied when necessary to address class imbalance.

Model Training

- Traditional Machine Learning Models: Logistic Regression and Support Vector Machines were trained on TF-IDF and Bag-of-Words features. Hyperparameters such as regularization strength, kernel type (for SVM), and maximum iterations were optimized using the validation set.
- Deep Learning Models: LSTM and CNN models were trained using word embeddings as input. Key hyperparameters included the number of layers, number of neurons, dropout rates, batch size, and learning rate. Early stopping was employed to prevent overfitting.
- Transformer-based Model (BERT): The pre-trained BERT model was fine-tuned on the dataset. A small learning rate and appropriate batch size were chosen to retain pre-trained knowledge while adapting to the specific anxiety detection task. Fine-tuning was performed over several epochs until validation loss stabilized.

Evaluation Metrics

The performance of all models was assessed using standard metrics for classification tasks:

- Accuracy: Measures the proportion of correctly classified tweets.
- Precision: Measures the proportion of correctly identified anxiety tweets among all tweets predicted as anxiety-related.
- Recall (Sensitivity): Measures the proportion of actual anxiety tweets correctly identified by the model.
- F1-Score: The harmonic mean of precision and recall, providing a balanced measure of model performance.

These metrics were chosen to ensure that both false positives and false negatives are appropriately accounted for, which is crucial in anxiety detection tasks.

Temporal Trend Analysis

After identifying anxiety-related tweets using the best-performing model, temporal trend analysis was performed to understand fluctuations in anxiety expressions over time. Tweets were aggregated on a daily or weekly basis, and visualizations such as line plots were generated to show peaks, dips, and patterns. This analysis provides insights into how external events, social circumstances, or global issues may influence collective anxiety levels on social media.

Results and Analysis

The results of this study focus on the performance of different models in detecting anxiety-related tweets and the temporal trends observed in social media expressions of anxiety. Multiple machine learning, deep learning, and transformer-based models were evaluated to identify the most effective approach.

Model Performance

The performance of each model was measured using accuracy, precision, recall, and F1-score.

- Traditional Machine Learning Models: Logistic Regression and Support Vector Machines performed adequately on TF-IDF and Bag-of-Words features. SVM generally achieved higher accuracy compared to Logistic Regression due to its ability to separate data in high-dimensional space. However, both models struggled with subtle context-dependent expressions, which sometimes led to misclassification of tweets that contained implicit anxiety cues.
- Deep Learning Models: LSTM and CNN models showed improved performance by capturing sequential and local contextual patterns. LSTM was particularly effective in understanding the order of words, while CNN excelled at extracting key phrases and n-gram features that indicate anxiety.

- Transformer-based Model (BERT): Fine-tuned BERT consistently outperformed classical and deep learning models across all metrics. Its ability to consider both left and right context in sentences allowed it to detect nuanced expressions of anxiety, including implicit mentions and variations in phrasing. This indicates that transformer-based approaches are highly suitable for analyzing emotional content in social media data.

Temporal Trend Analysis

Using the best-performing model, tweets were aggregated over daily or weekly intervals to identify trends in anxiety expressions. The analysis revealed patterns such as:

- Peaks in anxiety are mentioned during significant societal events or crises.
- Gradual fluctuations that may correspond to general societal stress or public awareness campaigns.
- Differences in expression patterns between weekdays and weekends, suggesting behavioral influences on social media activity.

Visualizations such as line charts, bar graphs, and heatmaps can be employed to depict these temporal variations effectively. These insights highlight how NLP-based systems can provide meaningful understanding of mental health trends in a population.

Interpretation

The results demonstrate that:

1. Transformer-based models like BERT are highly effective for anxiety detection, capturing subtle contextual cues.
2. Traditional models can provide reasonable baseline performance but are limited in handling complex linguistic patterns.
3. Temporal analysis of social media data can reveal collective emotional trends, offering a real-time perspective on societal anxiety levels.

This study confirms that NLP-driven approaches can serve as a scalable, cost-effective tool for monitoring emotional well-being in digital communities, while emphasizing ethical considerations and user privacy.

Ethical Consideration

Research involving social media data must carefully address ethical issues to ensure the privacy, safety, and rights of users are protected. In this study, the following ethical guidelines were strictly followed:

1. **User Privacy:** No personal identifiers, usernames, or profile information were stored or analyzed. Only the textual content of tweets was used, ensuring that individuals cannot be identified.
2. **Non-Clinical Use:** The study does not attempt to diagnose, treat, or make clinical judgments about any individual's mental health. The focus is solely on analyzing general trends in anxiety expressions at a population level.
3. **Data Usage Compliance:** Tweets were collected in accordance with Twitter's API terms of service, respecting restrictions on data access and redistribution.
4. **Responsible Interpretation:** Findings are reported in aggregate form, and no conclusions are drawn about specific users or groups. Temporal and statistical patterns are interpreted cautiously, avoiding sensationalism or misrepresentation.
5. **Transparency:** The methodology, data processing, and modeling approaches are described openly to ensure reproducibility and accountability, allowing others to understand the scope and limitations of the study.

By adhering to these ethical principles, the research maintains integrity and demonstrates responsible use of NLP techniques in mental health-related studies.

Conclusion

This research demonstrates the effectiveness of Natural Language Processing (NLP) techniques in detecting anxiety-related expressions from Twitter data by collecting and preprocessing tweets containing anxiety-related keywords and applying both traditional machine learning models and advanced deep learning approaches while maintaining ethical and privacy standards. The results indicate that transformer-based models, particularly BERT, outperform classical methods in capturing subtle contextual cues and complex emotional expressions, and the temporal analysis of anxiety-related tweets reveals meaningful patterns and fluctuations over time, highlighting the ability of social media to reflect collective emotional states in real time. Although the study has certain limitations, including language constraints, dataset coverage, and challenges associated with informal social media text, it provides a strong foundation for future research. The proposed methodology and findings demonstrate the potential of NLP-driven approaches to support large-scale mental health monitoring, trend analysis, and public awareness initiatives without compromising ethical principles. Overall, this research underscores the value of integrating computational techniques with social media data to better understand societal mental health trends and suggests that responsible use of these insights can contribute to improved detection accuracy, expanded research scope, and the promotion of emotional well-being in digital communities.

Acknowledgment

The authors would like to express their sincere gratitude to ATSS College of Business Studies and Computer Application for providing the academic environment and resources necessary to conduct this research.

We extend our heartfelt thanks to our respected guides, Dr. Manjusha Patil and Dr. Vinaya Keskar, for their valuable guidance, encouragement, and continuous support throughout the completion of this study.

Financial support and sponsorship

Nil.

Conflicts of interest

The authors declare that there are no conflicts of interest regarding the publication of this paper.

References

1. Abdul-Mageed, M., & Ungar, L. H. (2017). Emotional intelligence and anxiety detection from social media text. *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, 1–10.
2. Benton, A., Mitchell, M., & Hovy, D. (2017). Multitask learning for mental health using social media text. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, 2, 152–162.
3. Calvo, R. A., Milne, D. N., Hussain, M. S., & Christensen, H. (2017). Natural language processing in mental health applications using non-clinical texts. *Natural Language Engineering*, *23*(5), 649–685. https://doi.org/10.1017/S1351324917000114
4. Coppersmith, G., Dredze, M., Harman, C., & Hollingshead, K. (2014). Quantifying mental health signals in Twitter. *Proceedings of the ACL Workshop on Computational Linguistics and Clinical Psychology*, 51–60.
5. De Choudhury, M., Gamon, M., Counts, S., & Horvitz, E. (2013). Predicting depression via social media. *Proceedings of the International AAAI Conference on Web and Social Media*, 7(1), 128–137.
6. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, 4171–4186. https://doi.org/10.18653/v1/N19-1423
7. Guntuku, S. C., Yaden, D. B., Kern, M. L., Ungar, L. H., & Eichstaedt, J. C. (2017). Detecting depression and mental illness on social media: An integrative review. *Current Opinion in Behavioral Sciences*, *18*, 43–49.
8. Ji, S., Pan, S., Li, X., Cambria, E., Long, G., & Huang, Z. (2020). Suicidal ideation detection: A review of machine learning methods and applications. *IEEE Transactions on Computational Social Systems*, *7*(5), 1226–1238.
9. Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis Lectures on Human Language Technologies*, *5*(1), 1–167. https://doi.org/10.2200/S00416ED1V01Y201204HLT016
10. Matero, M., Idnani, A., Son, Y., Giorgi, S., Vu, H., Zamani, M., ... & Resnik, P. (2019). Suicide risk assessment with multi-level dual-context language and BERT. *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology*, 39–44.
11. Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
12. Ortu, M., Adams, B., Destefanis, G., Tourani, P., Marchesi, M., & Tonelli, S. (2016). Are bullies more productive? Empirical study of affectiveness vs. issue fixing time. *Proceedings of the 13th International Conference on Mining Software Repositories*, 303–313.
13. Paul, M. J., & Dredze, M. (2017). Social monitoring for public health. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, *9*(5), 1–183.
14. Pennington, J., Socher, R., & Manning, C. D. (2014). GloVe: Global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.
15. Resnik, P., Armstrong, W., Claudino, L., Nguyen, T., Nguyen, V. A., & Boyd-Graber, J. (2015). Beyond LDA: Exploring supervised topic modeling for depression-related language in Twitter. *Proceedings of the CLPsych Workshop*, 99–107.
16.Ríssola, E. A., Valdez, A. C., Greco, S., Dragoni, M., Pompianu, L., & Dignum, F. (2020). Emotion detection in social media: A survey. *ACM Computing Surveys*, *54*(1), 1–36.
17. Saha, K., Torous, J., Ernala, S. K., Rizuto, C., Stafford, A., De Choudhury, M. (2022). A computational study of mental health awareness campaigns on social media. *Proceedings of the International AAAI Conference on Web and Social Media*, 16(1), 906–917.
18. Sawhney, R., Mathur, P., Shah, R. R., & Mahata, D. (2021). Multimodal multi-task learning for mental health prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(2), 120–128.
19. Shatte, A. B. R., Hutchinson, D. M., & Teague, S. J. (2019). Machine learning in mental health: A scoping review of methods and applications. *Psychological Medicine*, *49*(9), 1426–1448.
20. Trotzek, M., Koitka, S., & Friedrich, C. M. (2018). Utilizing neural networks and linguistic metadata for early depression detection. *Proceedings of the CLPsych Workshop*, 28–37.
21. Twitter Developer Documentation. (2023). Twitter API v2 reference. https://developer.twitter.com/en/docs
22. Yates, A., Cohan, A., & Goharian, N. (2017). Depression and self-harm risk assessment in online forums. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2968–2978.
23. Zogan, H., Razzak, I., Wang, X., Jameel, S., & Xu, G. (2021). Explainable AI for depression detection using social media. *Journal of Biomedical Informatics*, *114*, 103698.