

Reading Emotions Behind TikTok Text: Fine-Tuning IndoBERT for Nine-Class Emotion Classification in the Indonesian Language

Indri Anjar Kartika Sari¹, Catur Suratnoaji², Agus Widiyarta³

¹ Universitas Pembangunan Nasional “Veteran” Jawa Timur; 25067020011@student.upnjatim.ac.id

² Universitas Pembangunan Nasional “Veteran” Jawa Timur; catur_suratnoaji.ilkom@upnjatim.ac.id

³ Universitas Pembangunan Nasional “Veteran” Jawa Timur; agus_widiyarta.adneg@upnjatim.ac.id

ARTICLE INFO

Keywords:

IndoBERT, fine-tuning, emotion classification, transformer, NLP, TikTok, Indonesian, text computing.

Article history:

Received 2026-09-21

Revised 2026-09-22

Accepted 2026-12-01

ABSTRACT

This study focuses on fine-tuning IndoBERT for nine-class emotion classification of Indonesian TikTok text. A computational BAHAGIA pipeline is developed covering text preprocessing, label construction, class-imbalance handling, transformer fine-tuning, evaluation, and model export. The primary dataset contains 8,834 TikTok records, of which 2,775 contain emotion labels and form the labeled classification basis. Fine-tuned IndoBERT is compared with a TF-IDF–Logistic Regression baseline using accuracy, macro-F1, and weighted-F1. The documented repository results report 78.0% accuracy, 0.80 macro-F1, and 0.78 weighted-F1 for IndoBERT, compared with 73.7%, 0.73, and 0.74 for the baseline. The 0.07 macro-F1 improvement indicates better average performance across classes in the available evaluation. The main contribution is a contextual IndoBERT-based computational approach for fine-grained emotion classification in Indonesian social-media text. Heuristic labels and limited human validation constrain interpretation; the model is intended for text classification and social listening rather than psychological diagnosis.

This is an open access article under the [CC BY-NC-SA](https://creativecommons.org/licenses/by-nc-sa/4.0/) license.



Corresponding Author:

Indri Anjar Kartika Sari

Universitas Pembangunan Nasional “Veteran” Jawa Timur; 25067020011@student.upnjatim.ac.id

1. INTRODUCTION

TikTok has developed into one of the most influential digital communication spaces in contemporary society, functioning not only as an entertainment platform but also as a dynamic environment where users produce, consume, interpret, and circulate multimodal messages through short-video formats. Although audiovisual elements represent the dominant feature of TikTok content, textual components such as captions, hashtags, comments, and descriptions constitute an important layer of meaning construction. These textual elements provide valuable evidence for examining how individuals narrate personal experiences, express emotions, and negotiate social meanings within digital environments.

In digital communication research, social media texts should not be interpreted as transparent reflections of users’ internal psychological conditions; rather, they represent communicative artifacts through which individuals construct identities, frame experiences, and communicate emotional meanings to specific audiences. Therefore, analyzing social media text requires attention

to the social and linguistic contexts surrounding digital expression. This perspective highlights the potential of social media data as a rich source for understanding contemporary communication patterns, public discourse, and collective emotional expressions.

The rapid growth, accessibility, and diversity of user-generated content make platforms such as TikTok valuable resources for computational approaches aimed at identifying large-scale patterns of digital communication (Ceron & Negri, 2016; Stieglitz et al., 2018). Despite its potential, analyzing Indonesian-language social media text presents substantial computational and linguistic challenges. Unlike formal written language, TikTok captions and descriptions are characterized by highly dynamic forms of digital expression, including informal vocabulary, non-standard spelling, abbreviations, emoji usage, repeated characters, hashtags, slang expressions, and frequent code-mixing between Indonesian and other languages.

Moreover, emotional expressions in online communication are often implicit, context-dependent, and influenced by cultural and situational factors. These characteristics create difficulties for conventional text representation methods, particularly approaches based on frequency-based models such as Bag-of-Words and Term Frequency–Inverse Document Frequency (TF-IDF). Although these classical approaches remain useful for extracting statistical patterns from text, they often fail to capture semantic relationships and contextual meanings between words. Consequently, sentiment and emotion classification in Indonesian digital communication requires computational models capable of understanding linguistic context beyond surface-level word occurrence.

Transformer based architectures provide a significant advancement in this area because they generate contextualized representations that consider relationships between words within a sequence, enabling more sophisticated understanding of complex language patterns. This study focuses on IndoBERT as the primary computational model because it represents one of the most relevant pretrained language models specifically developed for Indonesian language understanding. Unlike traditional sparse representation methods, IndoBERT employs transformer-based contextual embeddings that allow words to have different meanings depending on their surrounding linguistic context.

This characteristic is particularly important for Indonesian social media analysis, where identical expressions may convey different emotional meanings depending on usage, tone, and interactional context. Previous developments in Indonesian Natural Language Understanding (NLU), including IndoNLU and IndoBERT, have demonstrated the effectiveness of transformer models for various downstream tasks such as sentiment analysis, text classification, and language understanding benchmarks (Koto et al., 2020; Wilie et al., 2020). Building upon these developments, the present study applies IndoBERT fine-tuning for multi-class emotion classification of TikTok textual content.

Rather than limiting analysis to positive and negative sentiment categories, this research investigates nine emotion classes to provide a more nuanced representation of how emotions are expressed in Indonesian digital communication. Such an approach contributes to a deeper understanding of emotional diversity within online discourse. The analysis of emotional expression on TikTok has gained increasing relevance within broader discussions of digital wellbeing, particularly as researchers examine the relationship between intensive social media engagement and users' psychological experiences.

Recent systematic reviews have reported associations between problematic TikTok usage patterns and several mental health indicators, including anxiety, depressive symptoms, and emotional distress (Jain et al., 2025; Li et al., 2025). However, these findings should be interpreted carefully because most available studies rely on cross-sectional designs and cannot establish direct causal relationships between platform usage and psychological outcomes. Similarly, computational

emotion analysis based on captions or textual expressions should not be understood as a diagnostic tool for identifying mental health conditions. Instead, emotion classification provides an exploratory analytical framework for mapping patterns of emotional communication, identifying dominant affective expressions, and understanding how users represent experiences within digital spaces.

Through this perspective, artificial intelligence-based text analysis can support interdisciplinary research connecting natural language processing, communication studies, and digital wellbeing research. The BAHAGIA project was developed as an effort to address these research opportunities by establishing an integrated pipeline for Indonesian TikTok emotion analysis. The project incorporates several methodological components, including TikTok text preprocessing, nine-class emotion annotation, IndoBERT fine-tuning, class weight optimization to address imbalance problems, comparison with classical machine learning baselines, and model deployment through ONNX export.

2. METHODS

This study employed a computational experimental research design based on Natural Language Processing (NLP) to develop an emotion classification model for Indonesian-language TikTok textual content using IndoBERT finetuning. The research aimed to classify TikTok captions into nine emotion categories: Happy, Sad, Anxious/Worried, Angry/Irritated, Bored, Tired/Burnout, Neutral, Envious/Insecure, and Satisfied/Grateful. The dataset consisted of publicly available Indonesian TikTok textual data collected from user-generated content, while personal identification information was excluded to maintain ethical standards. The collected texts underwent preprocessing, including normalization, removal of irrelevant elements, and tokenization using the IndoBERT tokenizer while preserving contextual features such as informal expressions, hashtags, and emojis. Each text sample was manually annotated according to predefined emotion guidelines, and class-weighting techniques were applied to address class imbalance among emotion categories. The proposed model was developed by fine-tuning IndoBERT with a supervised classification approach, using training, validation, and testing data partitions. Model optimization was conducted using the AdamW optimizer with cross-entropy loss, while hyperparameters were adjusted based on validation performance. The performance of IndoBERT was evaluated using accuracy, precision, recall, macro F1-score, and confusion matrix analysis, with a comparison against a TF-IDF-based classical machine learning baseline to examine the advantages of contextualized transformer representations. The final model was exported into ONNX format to support practical deployment and reproducibility. This methodological framework follows current practices in transformer-based NLP research by combining domain-specific language modeling, fine-grained emotion classification, and transparent computational evaluation (Devlin et al., 2019; Koto et al., 2020).

3. FINDINGS AND DISCUSSION

3.1. Data Profile And Emotion Distribution

The main dataset comprises 8,834 rows. Of these, 2,775 rows (31.41%) have emotion labels, whereas 6,059 rows have not been labeled. The distribution across the nine classes is imbalanced. Satisfied/Grateful is the largest class with 710 instances (25.59%), whereas Bored has only 45 instances (1.62%). This imbalance is important because a model can attain high accuracy by following the majority class too closely. Therefore, the use of class weighting and the reporting of macro-F1 are appropriate methodological choices for multi-class classification. This distribution also shows that digital wellbeing categories do not appear evenly.




Figure 1. Distribution of data across the nine emotion classes

The Anxious/Worried, Tired/Burnout, and Envious Insecure classes account for smaller proportions than Satisfied/Grateful, while Bored is the rarest class. This is consistent with the project's note that the Bored class is a particular challenge and requires additional data. From a communication perspective, this imbalance cannot be read directly as evidence that TikTok users are more often happy or grateful. The distribution primarily reflects the data that were successfully collected and labeled using the study's scheme.

3.2. Comparison Of Indobert And The Baseline Architecture And Role Of Indobert

In the BAHAGIA architecture, IndoBERT functions as a contextual encoder that maps a sequence of tokens into high-dimensional representations, after which a classification head produces probabilities for the nine classes. Fine-tuning allows the encoder parameters to adapt to the characteristics of TikTok language, while class weighting reduces the dominance of the majority class in the objective function. Conceptually, this pipeline can be written as: text IndoBERT tokenization contextual embedding classification head 9-class probabilities emotion label. This approach differs from TF-IDF, which represents a document as term weights without a contextual attention mechanism. The comparison of the two models is therefore used to test whether contextual representation yields a performance gain on the same task.

The documented evaluation shows that fine-tuned IndoBERT achieved an accuracy of 78.0%, a macro-F1 of 0.80, and a weighted-F1 of 0.78. The TF-IDF + Logistic Regression baseline achieved an accuracy of 73.7%, a macro-F1 of 0.73, and a weighted-F1 of 0.74. IndoBERT thus improved accuracy by 4.3 percentage points and macro F1 by 0.07 in the available evaluation. For computational studies, macro-F1 is an important indicator because the nine class dataset is imbalanced. The advantage of the contextual model can be explained conceptually by the informal character of TikTok text.

Model	Accuracy	Macro-F1	Weighted-F1
TF-IDF + Logistic Regression	73.7%	0.73	0.74
Fine-tuned IndoBERT	78.0%	0.80	0.78

Table 1. Model performance documented in the BAHAGIA repository

The same word can carry different emotional meanings depending on context, emoji, hashtags, and sentence structure. The TF-IDF approach primarily represents word frequency and weight, whereas transformers exploit contextual representations. This finding is consistent with the development of IndoBERT as an Indonesian language model and with the literature on emotion classification in Indonesian social media text (Koto et al., 2020; Glenn et al., 2023).




Figure 2. Comparison of accuracy and macro-F1

The repository notes that the largest improvement over the baseline occurred in the Bored class, which is the rarest class. This finding is methodologically interesting but must be read with caution because the available documentation does not include the full confusion matrix and complete per-class metrics for all nine categories. The README displays a classification report for only six classes. This study therefore does not conclude that IndoBERT outperforms the baseline on every class.

3.3. Interpretation For Digital Communication

The results indicate that emotion classification can be used to extend digital communication analysis beyond merely determining whether a text is positive or negative, toward a more granular mapping of affective categories. In the TikTok context, this approach enables researchers to identify the dominant types of expression in a given corpus and to observe how language, hashtags, emoji, and context are used to convey experience.

However, classification results must not be treated as a direct measurement of a creator's psychological state. A caption may be a quotation, song lyric, news item, promotion, humor, storytelling, or content created for engagement purposes. A single example that appears sad does not necessarily indicate that the content creator is experiencing sadness. Thus, the safest unit of interpretation is "emotional expression in text," not "the user's actual emotion." This boundary is consistent with literature emphasizing the need for caution when linking social media data to the conditions of populations or individuals (Ceron & Negri, 2016; Mellon & Prosser, 2017).

From a digital wellbeing standpoint, BAHAGIA can function as a social listening system that helps identify patterns of emotional communication at scale. Its potential lies primarily in aggregation and trend exploration, for example, observing changes in the proportions of

emotional categories over a given period or comparing content types. Such a system can serve as a research aid and as a tool for designing digital literacy interventions, but it cannot replace validated psychological instruments or professional assessment.

The findings also highlight the importance of label quality. Because 68.59% of the main data had no emotion label in the documented audit, and the populated labels derive from heuristics, model performance must be understood as performance relative to the available label scheme. This is a classic issue in supervised learning: a model may learn label consistency without any guarantee that the labels fully represent the theoretical construct to be measured. Human in the loop validation with independent annotators and Cohen's Kappa is an important step for improving construct validity.

DISCUSSION

The findings of this study demonstrate that emotion classification in Indonesian TikTok textual content is strongly influenced by the linguistic characteristics of digital communication, the distribution of emotion labels, and the representational capacity of the computational model. The dataset contains 8,834 textual records, but only 2,775 records (31.41%) have emotion labels, while 6,059 records (68.59%) remain unlabeled. Among the labeled data, Satisfied/Grateful is the largest category with 710 instances (25.59%), whereas Bored is the smallest with only 45 instances (1.62%). This imbalance is theoretically important because supervised classification models tend to learn patterns according to the empirical distribution of training data.

Consequently, the use of class weighting and macro-F1 is appropriate because it reduces the possibility that model performance is dominated by majority classes. However, the distribution of labels should not be interpreted as a direct representation of the emotional condition of TikTok users. Rather, it reflects the characteristics of the collected and annotated corpus. In digital communication theory, social media texts are communicative artifacts through which users construct meaning, identity, and interaction with audiences; therefore, the presence of a particular emotional label represents an expression embedded in discourse rather than an objective measurement of an individual's psychological state. The predominance of Satisfied/Grateful and the limited representation of Bored also illustrates the importance of distinguishing between data prevalence and emotional prevalence.

A larger number of Satisfied/Grateful texts does not necessarily indicate that gratitude or satisfaction is the dominant psychological experience among TikTok users. Similarly, the small number of Bored texts cannot be interpreted as evidence that boredom is uncommon among users. The observed proportions are affected by the data collection process, content availability, annotation procedures, and the operational definitions applied to each emotion category. This interpretation is consistent with the broader understanding of social media research that online behavioral data are shaped by self-selection, communication purposes, platform affordances, and audience considerations (Ceron & Negri, 2016; Mellon & Prosser, 2017).

Therefore, the findings are most appropriately interpreted as describing patterns of emotional expression within the analyzed corpus, rather than estimating the psychological distribution of the wider TikTok population. From the perspective of Natural Language Processing, the comparison between IndoBERT and the TF-IDF Logistic Regression baseline provides evidence for the importance of contextual language representation. IndoBERT achieved an accuracy of 78.0%, macro-F1 of 0.80, and weighted-F1 of 0.78, compared with 73.7% accuracy, 0.73 macro-F1, and 0.74 weighted-F1 for the TF-IDF baseline.

The improvement of 4.3 percentage points in accuracy and 0.07 in macro-F1 indicates that contextualized representation provides an advantage for the nine-class emotion classification task.

This result can be explained through the theoretical architecture of BERT, in which representations are generated according to relationships among tokens within their linguistic context rather than solely according to their frequency within a document (Devlin et al., 2019). IndoBERT is particularly relevant in this context because it was developed to capture Indonesian linguistic patterns and has demonstrated applicability across Indonesian language understanding tasks (Koto et al., 2020; Wilie et al., 2020).

The advantage of IndoBERT becomes particularly meaningful when considering the informal nature of TikTok communication. Digital texts frequently contain slang, abbreviations, repeated characters, hashtags, emojis, non-standard spelling, and expressions whose meanings depend heavily on context. TF-IDF can identify the statistical importance of words, but it does not inherently distinguish contextual changes in meaning. Transformer-based contextual representation, by contrast, enables the model to consider relationships among words in a sequence. From the perspective of linguistic theory, this distinction is important because emotional meaning is not always encoded in individual words but emerges from combinations of lexical, syntactic, pragmatic, and contextual elements.

Thus, the higher macro-F1 achieved by IndoBERT suggests that contextual information is particularly valuable when classifying nuanced emotional expressions in Indonesian digital discourse. The results also support the argument that emotion classification should move beyond conventional binary sentiment analysis. The nine categories used in this study distinguish between emotional states that may otherwise be grouped together under broad positive or negative categories. For example, Sad, Anxious/Worried, Tired/Burnout, and Envious/Insecure may all contain negatively valenced expressions but represent substantially different affective meanings. Similarly, Happy and Satisfied Grateful may share positive linguistic characteristics while representing different forms of emotional expression.

The use of multiple categories therefore provides a more fine-grained analytical framework for investigating affective communication on TikTok. This approach is theoretically consistent with the view that emotions expressed through digital communication are multidimensional and contextually constructed rather than adequately represented by a simple positive-negative polarity. The documented improvement in the Bored category is also noteworthy because Bored is the smallest class in the labeled dataset. This observation suggests that contextual modeling may provide benefits for minority categories that contain distinctive linguistic patterns. Nevertheless, this result should be interpreted cautiously because the available documentation does not provide a complete confusion matrix and complete per-class metrics for all nine categories.

The relatively small number of Bored samples also increases the possibility that performance for this category may be sensitive to the particular train-test partition. Therefore, the result should be regarded as an indication rather than definitive evidence that IndoBERT consistently performs better for every minority class. The findings can further be understood through the perspective of digital communication theory. TikTok captions and other textual elements are not necessarily spontaneous representations of users' internal experiences. They can function as instruments of self-presentation, humor, storytelling, entertainment, persuasion, social commentary, or audience engagement. A caption containing sadness-related vocabulary, for instance, may represent a quotation, song lyric, fictional narrative, or humorous expression rather than the creator's actual emotional state.

Consequently, the appropriate analytical unit is emotion expressed in text, rather than the assumed emotion experienced by the user. This distinction is essential for avoiding psychological overinterpretation of computational social media data. This theoretical boundary is particularly important when the findings are connected to digital wellbeing. The BAHAGIA framework has potential as a social listening and exploratory research system because it can identify patterns of

emotional communication across large quantities of textual data. Researchers may use such information to examine changes in emotional expression over time, compare different content categories, or identify recurring patterns in digital discourse. However, emotion classification should not be treated as a diagnostic instrument. A classification of Anxious/Worried or Tired Burnout indicates that the textual content contains linguistic patterns associated with those categories; it does not establish that the content creator has an anxiety disorder, burnout, or another psychological condition.

This distinction preserves the methodological boundary between computational text analysis and validated psychological assessment. Another important theoretical implication concerns construct validity and annotation quality. Since 68.59% of the main dataset remains unlabeled and the available labels are based on the project's existing labeling scheme and heuristics, the performance of IndoBERT should be interpreted as performance relative to the available annotation framework. In supervised learning, a model learns patterns embedded in its training labels. Therefore, even a high-performing model can reproduce inconsistencies or limitations contained within the annotation process.

The macro-F1 score of 0.80 demonstrates that IndoBERT performs well in reproducing the available labels, but it does not by itself establish that the model has perfectly captured the theoretical construct of human emotion. For this reason, the development of the BAHAGIA framework should place equal emphasis on model development and label validation. Human-in-the-loop annotation involving independent annotators would strengthen the reliability of the emotion categories, particularly for semantically overlapping classes such as Sad, Anxious/Worried, Tired/Burnout, and Envious Insecure.

The use of inter-rater agreement, such as Cohen's Kappa where appropriate, can provide evidence that different annotators consistently interpret the predefined emotion categories. Clear annotation guidelines, representative examples, and adjudication procedures would further improve construct validity and reduce the possibility that the model merely learns idiosyncratic labeling patterns. Taken together, these findings indicate that the contribution of the BAHAGIA framework is not limited to achieving higher classification performance. Its broader contribution lies in demonstrating how contextualized Indonesian language modeling can be integrated with fine-grained emotion classification to examine emotional expressions in digital communication. IndoBERT provides the computational capacity to model contextual linguistic information, while the nine class framework provides a more nuanced representation of affective expression than conventional sentiment polarity.

At the same time, the findings emphasize that computational predictions must be interpreted within the social, linguistic, and methodological contexts in which digital texts are produced. Therefore, future research should prioritize expansion of minority-class data, improvement of annotation quality, independent human validation, and external testing across different TikTok content and temporal contexts. Such developments would strengthen the reliability of emotion classification while maintaining an appropriate distinction between computational identification of emotional expressions and claims about users' actual psychological states.

4. CONCLUSION

Reading Emotions Behind TikTok Text: Fine-Tuning IndoBERT for Nine-Class Emotion Classification in the Indonesian Language In the evaluation documented in the repository, fine-tuned IndoBERT obtained an accuracy of 78.0%, a macro-F1 of 0.80, and a weighted-F1 of 0.78, whereas TF-IDF-Logistic Regression obtained 73.7%, 0.73, and 0.74, respectively. These results indicate a performance improvement for the contextual model on the study dataset, particularly in

macro-F1. The main contribution of this study lies in developing a computational framework for mapping emotional expression on Indonesian social media in greater detail and linking it to digital communication studies and digital wellbeing. However, the results cannot be interpreted as a diagnosis or a measure of an individual's psychological condition. The validity of the heuristic labels, data representation bias, and the limitations of the evaluation documentation constitute the main boundaries. Future research should prioritize human-in-the-loop validation, complete reporting of hyperparameters and the confusion matrix, temporal testing, and the integration of emotion classification with topic/discourse analysis. With these improvements, BAHAGIA can develop from a classification prototype into a more auditable research instrument for social listening and digital communication studies.

REFERENCES

- BAHAGIA. (2026). BAHAGIA: Behavioral Analysis of Happiness And Grief In social-media Addiction—Emotion classification engine for Indonesian-language TikTok content. GitHub. <https://github.com/indri007/BAHAGIA>
- Ceron, A., & Negri, F. (2016). The “Social Side” of Public Policy: Monitoring Online Public Opinion and Its Mobilization During the Policy Cycle. *Policy & Internet*, 8(2), 131–147. <https://doi.org/10.1002/poi3.117>
- Glenn, A., LaCasse, P., & Cox, B. (2023). Emotion classification of Indonesian Tweets using Bidirectional LSTM. *Neural Computing and Applications*, 35, 9567–9578. <https://doi.org/10.1007/s00521-022-08186-1>
- Jain, L., Velez, L., Karlapati, S., Forand, M., Kannali, R., Yousaf, R. A., Ahmed, R., Sarfraz, Z., Sutter, P. A., Tallo, C. A., & Ahmed, S. (2025). Exploring problematic TikTok use and mental health issues: A systematic review of empirical studies. *SAGE Open Medicine*, 16. <https://doi.org/10.1177/21501319251327303>
- Koto, F., Rahimi, A., Lau, J. H., & Baldwin, T. (2020). IndoLEM and IndoBERT: A benchmark dataset and pre-trained language model for Indonesian NLP. *Proceedings of the 28th International Conference on Computational Linguistics*, 757–770. <https://doi.org/10.18653/v1/2020.coling-main.66>
- Li, Y., et al. (2025). Association between problematic TikTok use and mental health: A systematic review and meta-analysis. *International Journal of Public Health*, 70. <https://doi.org/10.3934/publichealth.2025027>
- Mellon, J., & Prosser, C. (2017). Twitter and Facebook are not representative of the general population: Political attitudes and demographics of British social media users. *Research & Politics*, 4(3), 1–9. <https://doi.org/10.1177/2053168017720008>
- Stieglitz, S., Mirbabaie, M., Ross, B., & Neuberger, C. (2018). Social media analytics—Challenges in topic discovery, data collection, and data preparation. *International Journal of Information Management*, 39, 156–168. <https://doi.org/10.1016/j.jinfomgt.2017.12.002>
- Wilie, B., Vincentio, K., Winata, G. I., Cahyawijaya, S., Li, X., Lim, Z. Y., Soleman, S., Mahendra, R., Fung, P., Bahar, S., & Purwarianti, A. (2020). IndoNLU: Benchmark and resources for evaluating Indonesian natural language understanding. *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, 843–857. <https://doi.org/10.18653/v1/2020.aacl-main.85>

- Zimbra, D., Abbasi, A., Zeng, D., & Chen, H. (2018). The state-of-the-art in Twitter sentiment analysis: A review and benchmark evaluation. *ACM Transactions on Management Information Systems*, 9(2), Article 5. <https://doi.org/10.1145/3185045>
- Public Sentiment toward Indonesia's Free Nutritious Meal Program on X: A Systematic Review. (2026). Manuscript/document provided for the present research context.