Enhancing Fan Base Engagement through Explainable Self-Learning Sentiment Analysis

Mikael Lye
School of Computing
Informatics Institute of Technology
Colombo, Sri Lanka
m.mikaellye@gmail.com

Nethmi Wijesinghe
School of Computing
Informatics Institute of Technology
Colombo, Sri Lanka
wt.nethmi@gmail.com

2024 Moratuwa Engineering Research Conference (MERCon) | 979-8-3315-2904-8/24/$31.00 ©2024 IEEE | DOI: 10.1109/MERCon63886.2024.10688963

***Abstract*—Individuals and brands with large fan bases face difficulty in understanding fan sentiment and its potential impact on fan engagement and performance. This is particularly pertinent within fast-paced sports such as *Formula 1*, where fan opinions can significantly influence driver and team morale. To address the above-mentioned problem, this study proposes the use of deep learning-based sentiment analysis techniques to enhance fan base engagement. The system would act as a tool to automate the process of marketing and public relations teams by analysing textual data provided by the fan base and providing meaningful insight/reports of the fan base's emotion of the posted content.This is achieved by a novel multi-class sentiment analysis system that utilizes a fine-tuned Distil-BERT model for classification, self-supervised techniques for self-improvement, and an explainable AI (XAI) approach for interpretability. The proposed system demonstrated strong performance during testing and evaluation, achieving an overall accuracy of 82% and an F1-score of 80%. Overall, the systems components and structure focuses on utilizing the least amount of resources while also maintaining a high prediction accuracy and speed. This ultimately results in a budget friendly and robust tool that can be integrated into bigger analytics systems.**

***Key Words*—Multi-Class Sentiment Analysis, Distil-BERT, Fan Engagement, Self-Supervised Learning, Explainable AI (XAI)**

I. Introduction

The emergence of digital communications has fundamentally transformed how businesses interact with their end users. Traditionally, collecting feedback involved direct interactions or paper-based surveys, but digital platforms dominate today's landscape. Feedback is now predominantly gathered through online surveys, social media platforms, and direct observation of behaviour and interactions [1]. Among these methods, social media is particularly significant due to its universal accessibility and popularity [2]. These platforms play a crucial role in shaping relationships between celebrities and their fans, facilitating a digital space where fans can express their support and preferences. However, for this relationship to work, the fan bases need to be understood and heard by the individual or organisation that they follow [3]. This creates the next challenge, where it gets quite difficult to understand the requirements of a large number of people. To overcome the challenge, sentiment analysis can be utilized to gather large amounts of data and categorise the data under a general topic, making it easier to understand the data [4]. In a field where the majority of the fan base exists online like *Formula 1*, sentiment analysis can be performed using the data of fans through their online posts and reactions. However, the current sentiment analysis methods, which primarily focus on binary positive or negative categorizations, fail to capture the complex emotional expressions prevalent among fans. This oversimplification significantly hinders the understanding of nuanced fan sentiments that include a spectrum of emotions such as joy, anticipation, and disappointment. Recognizing these limitations, this research identifies four primary gaps in the existing sentiment analysis frameworks:

1) Multi-class Emotion Sentiment Analysis: Current models do not sufficiently recognize or categorize a broad range of complex fan emotions [5]. The author provides a multi- class emotional sentiment analysis system that performs sentiment analysis on a granular level, where the model predicts the emotional sentiments of 28 diverse emotions.

2) Explainable AI (XAI) in Sentiment Analysis: There is a lack of transparency in how sentiment analysis models arrive at their conclusions, particularly in models dealing with multiple sentiment categories [6] [7]. To address the gap, this research aims to provide insight into the models predictions to increase the trust and reliability of the system.

3) Adaptive Learning Models: Existing systems do not adequately evolve with changing data landscapes, requiring manual adjustments to stay relevant [8]. The author addresses this gap by utilizing the Large Language Model's (LLM's) fine-tuning capabilities to continuously train itself off the data it predicts.

In summary, this paper provides a literature review on the existing work while addressing their strengths and weaknesses, dives into the architecture of the system explaining each core module within the system and their functionality. Furthermore, the paper discusses the models results and evaluation, and finally providing the future work and direction of the authors next few steps to improve on the system.

II. Literature Review

A. Sentiment Analysis

In the growing field of artificial intelligence, sentiment analysis is a common yet challenging task, aiming to understand the overall emotions of opinion holders and textually expressive users. Sentiment analysis utilizes natural language processing (NLP) to map sentences to sentiments based on the data it's trained on [9]. It identifies the opinions, emotions, and attitudes of the writer of the message or the text and classifies the sentiment [5]. As of now, there are three main types of sentiment analysis techniques: binary, multi-class, and granular sentiment analysis. Binary sentiment analysis categorizes text into one of two outcomes, either positive or negative, while multi-class sentiment analysis classifies text into more than two categories. Finally, granular sentiment analysis identifies specific emotional states for more in-depth sentiment understanding. These techniques can be implemented using machine learning-based methods, rule-based methods, or hybrid methods which combine both rule-based and machine learning-based methods by feeding a machine learning model rule-based data such as lexicons [10].

B. Explainable Artificial Intelligence (XAI)

As artificial intelligence (AI) technologies are increasingly adopted across different domains like healthcare, finance, and criminal justice, so does the requirement to understand the decisions made by these systems. Explainable Artificial Intelligence (XAI) is a growing field that aims to enhance the transparency and interpretability of AI systems to human users [11] [12]. XAI can be implemented or performed with the following identified methods: Interpretable Models, Post-hoc Explanation, Example-based Explanation, and Feature-based Explanation. These mentioned methodologies are the most recognised XAI techniques at the time of this research. The growing interest in XAI highlights the field's relative novelty and its significant potential for advancement and expansion [13].

C. Self-Supervised Learning Techniques

Self-supervised learning is a powerful approach in machine learning (ML) that leverages unlabeled data to train models, reducing the need for manual annotations and human intervention [8]. It involves creating pretext tasks that enable the model to learn useful representations from the data without explicit supervision. For instance, in the field of medical imaging, self-supervised learning has shown promise in developing robust classification models by extracting insights from large datasets without labels [14].

D. Fan Engagement in fan bases

Social media platforms empower individuals and organizations with a fan base to interact directly with their supporters. This has two main benefits: it strengthens relationships within the fan base, fosters loyalty and a sense of belonging, and allows them to gather valuable feedback and data about their audience's preferences. This information can then be used to refine their craft and deliver experiences that better resonate with their fans. [15]. Social media platforms have seen exponential user grown post-COVID-19 due to the physical isolation people faced. According to a newsletter by Backlinko, as of October 2023, the number of people using social media is over 4.95 billion worldwide, which means that approximately more than 60% of the world's population uses social media with the average user accessing 6.7 social media platforms monthly [16]. With this rise of social media users, feedback collection techniques shifted online to social media platforms. As of now, social media platforms do not provide inbuilt sentiment analysis analytics to verified big-time users or organisations [17]. Although there are third-party applications that provide this information for a subscription or separate fee, at the cost of losing privacy to these third parties.

E. Existing Work

The research conducted by Kotei and Thirunavukarasu [8] provides a comprehensive review of deep learning transformer architecture, specifically BERT and GPT. The research also focuses on self-supervised learning and pretraining concepts in large language models, including the model's adaptation to downstream tasks. While this paper only provides a review on the possible techniques, it does not implement the techniques to resolve any real-world problem.

Wang's research paper [18] compares the several deep learning models with diverse structural types to analyse their performance in sentiment analysis. The deep learning models analyzed include RNN, CNN, LSTM, BiLSTM, BERT, and FNN, from which the BERT model achieved the highest accuracy under the same experimental configurations as possible. The BERT model scored 0.9224 accuracy on the test set with only a 0.314 loss.

Souza and Filho [19] conducts a comprehensive study focusing on various strategies for aggregating features from the BERT output layer, particularly focusing on their effectiveness for sentiment analysis. The proposed aggregation approach over BERT outputs brought considerable gains over the more conventional scheme of selecting the first layer, i.e., the output related to the [CLS] token. For all datasets, except UTLC-Movies, some BERT embedding configurations achieved the highest results among all evaluated models, concluding that fine-tuned BERT-based models greatly outperform pre-trained models. This paper provides a solution to achieve better results but lacks the ability to adapt to vocabulary change with time, requiring manual maintenance.

Meanwhile Tabinda Kokab, Asghar, and Naz [20] uses transformer-based deep learning to propose a generalized sentiment analysis model that is capable of handling noisy data and out of Vocabulary words (OOV) often seen in social media data. The research proposes an effective model combining BERT and a Convolution Bi-directional Recurrent Neural Network (CBRNN) for syntactic, semantic, sentiment, and contextual analysis. It utilizes zero-shot classification for polarity scores, pre-trained BERT for semantics, and contextual embedding. The model employs dilated convolution for local and global contextual features and Bi-LSTM for sequencing sentences. This CBRNN model was evaluated based on accuracy, precision, recall, f1-score and AUC values, where it was able to achieve 0.93, 0.93, 0.92, 0.93, and 0.969 respectively on the IMDB dataset.

The research papers by Diwali et al. [7] and Jain et al. [6] explore the methods of incorporating explainability to sentiment analysis tasks. While Jain et al. [6] only provides locally interpretable model-agnostic explanations (LIME) of NLP-based sentiment analysis, Diwali et al. [7] reviews the majority of the explainable sentiment analysis methods along with an in-depth analysis of the techniques. Overall, these papers provide insight into both ante-hoc and post-hoc on a global and local scope along with references to the sources the data was acquired from for a comprehensive review.

III. Methodology

A. Dataset Preparation and Preprocessing

1) Utilized Dataset: The effectiveness of any deep learning model significantly depends on the quality and preparation of the used dataset. For this study, authors utilized the **GoEmotions** dataset, which consists of text data labelled across 27 emotion classes, designed specifically for fine-tuning the sentiment analysis model. This dataset is structured into four sentiment categories: positive, negative, ambiguous, and neutral, with 12 classes belonging to the positive sentiment, 11 classes belonging to the negative sentiment, and 4 classes categorised as ambiguous with the last class for neutral [21]. The dataset was modified with an additional neutral class to enhance the dataset's capability in distinguishing non-emotional texts.

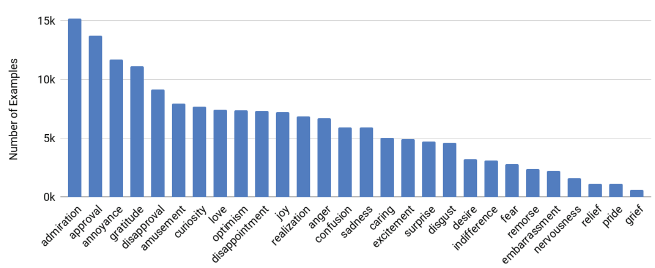

Fig. 1. The distribution of the emotional sentiment classes by the number of records [21].

2) Preprocessing the Dataset: The dataset was initially balanced to have the have an equal number of records for all 28 classes. This was archived by down-sampling and up-sampling techniques, where records under and above 5,500 records were balanced to fit the adequate sample count. To optimize the dataset for machine processing and improve the performance of the sentiment analysis model, several preprocessing steps were implemented:

- Text Cleaning: This step involved the removal of emojis, emoticons, URLs, user mentions, hashtags, and special characters that could potentially skew the model's learning process.
- Lemmatization: All texts were subjected to lemmatization to consolidate varying inflections of words into their base or dictionary form. For example, the phrase "Engineers" was lemmatized to "Engineer," focusing the model's attention on relevant nouns and thereby enhancing its learning accuracy.

B. System Architecture

The proposed architecture of the sentiment analysis system consists of three principal components: the Core Module, the Self-Learning Module, and the XAI Module. The core module is where the fine-tuned model predicts the sentiments of the input texts. In contrast, the self-learning module allows the system to fine-tune itself using the predictions it makes. Finally, the XAI module is where the predictions made by the system are interpreted. The Fig. 3 illustrates a schematic overview of the system, outlining the interactions between each component. Each module is discussed separately in detail below.

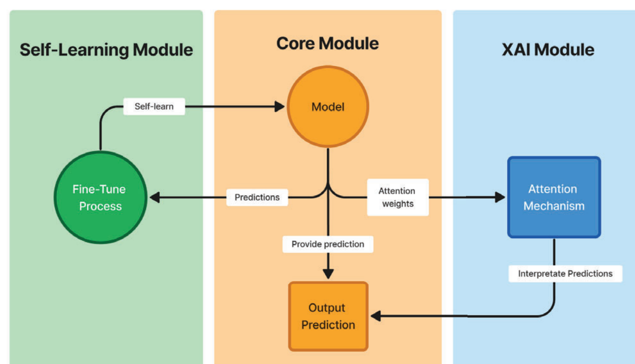

Fig. 2. A diagram of the proposed system's basic architecture.

1) Core Module: The model selected for multi-class sentiment prediction is DistilBERT, a pre-trained hugging face transformer model, chosen based on its demonstrated efficacy in similar tasks as outlined in recent studies [22]. This model is fine-tuned using the GoEmotions dataset to perform multi-class sentiment analysis on 28 diverse emotion classes. The fine-tuning employs Adam optimizer combined with grid search for hyperparameter optimization to enhance performance efficiently [23].

2) Self-Learning Module: This module incorporates a self-supervised learning strategy, enabling the model to independently adapt and improve by learning from the data it processes without human intervention [14]. This is achieved by setting up a loop for the model to get an input, predict its sentiment, and provide an output which is next fed back into the model to fine-tune itself on top of the already existing knowledge. To ensure the model isn't fine-tuning on incorrectly labeled data, the author has set rules where the model only trains on data that it predicts with 75% or higher confidence. This allows the model to keep its vocabulary updated overtime and reduces the manual maintenance once the system is deployed.

3) Explainable Artificial Intelligence (XAI) Module: The author opted for an attention mechanism approach from among the existing techniques in the XAI module. This approach leverages DistilBERT's ability to assign different weights to input tokens (tokenized words), allowing it to focus more on informative words while processing sequences and capturing contextual relationships effectively [24] [23]. The system utilises an attention mechanism to extract the most influential words from the input sentence and provide a visual depiction of them. This method not only aids in understanding the model's decision-making process but also in evaluating the contextual relevance of specific words within the input data. The use of attention mechanism provides users with the most influential word and also allows the user to connect associated words using the association rule to perform further analysis and get a better understanding on the fan bases common interests.

IV. Results and Discussion

A. Model Testing Process

The evaluation of the proposed system was conducted using the most relevant evaluation metrics for multi-class classification tasks as identified at the onset of this study. The model's performance was assessed comprehensively on both per-class and aggregate levels to derive detailed insights into its efficacy. The tests concluded resulting in the model achieving 82% overall accuracy, a weighted average of 81% for precision, 82% for recall, and 80% for F1-score of both positive and negative classification.

TABLE I
Class-Wise Test Results

Class	Precision	Recall	F1-Score
Amusement	0.77	0.81	0.79
Confusion	0.84	0.92	0.87
Disapproval	0.72	0.42	0.53
Excitement	0.78	0.9	0.84
Fear	0.86	0.96	0.91
...	...	...	...
Joy	0.79	0.82	0.8
Realization	0.85	1	0.92
Remorse	0.88	1	0.94
Sadness	0.81	0.99	0.89
Accuracy	0.82	0.82	0.82
Macro Avg.	0.81	0.82	0.81
Weighted Avg.	0.81	0.82	0.8

The class-specific testing process of the evaluation matrices is provided in Table 1. The table does not represent all the classes but provides a fraction of the emotional sentiment classes along with the evaluation metric values achieved. For additional context, the macro average is the calculation of a metric across all classes individually and then averaging the results, while the weighted average is the calculation of a metric for each class, with each class's contribution weighted by its support or number of true instances. These test processes were conducted to identify the model's performance on the test dataset that was created using the same dataset used for training the model.

B. Self-Learning Testing Process

The self-learning module of the proposed model capitalizes on the ability to refine its predictive accuracy over time through an iterative self-improvement process using unlabeled data. This adaptive mechanism allows the model to enhance and adjust its internal representations based on continuous feedback from its own predictions. To assess the effectiveness of the self-learning process, authors monitor the confidence levels of the model's predictions, both prior to and following the self-learning phase. These confidence levels are quantified using the probability scores assigned by the model to each prediction. For this study, authors establish a confidence threshold of 70%. Predictions with probability scores above this threshold are classified as having positive confidence, whereas those below are deemed to have negative confidence.

TABLE II
Self-Learning Module Test Results

Evaluations	Results Before	Results After
Average Confidence Score	73.24%	75%
Max Confidence Score	98.42%	99.67%
Min Confidence Score	11.56%	12.94%
Positive Confidence Count	10230	10690
Negative Confidence Count	3770	3310

This evaluation was completed by calculating the model's most and least confidently predicted scores, the average confidence score, and the counts of both positive and negatively predicted records. As provided in Table 2, the final results of the testing process conducted on the test dataset with 14,000 records to evaluate the self-learning module show increasingly improved sentiment predictions. This concludes that the self-learning module of the system maintains and slightly improves the model's predictions within the implemented domain.

C. XAI module Testing Process

The model's interpretations were evaluated by human observation to confirm the results generated. Additionally a qualitative analysis was conducted to get the qualitative measurements to validate the performance of this module. The visual interpretation of the system is provided in Fig. 4, where the words of the input test sentence are plotted along with their respective attention weight.

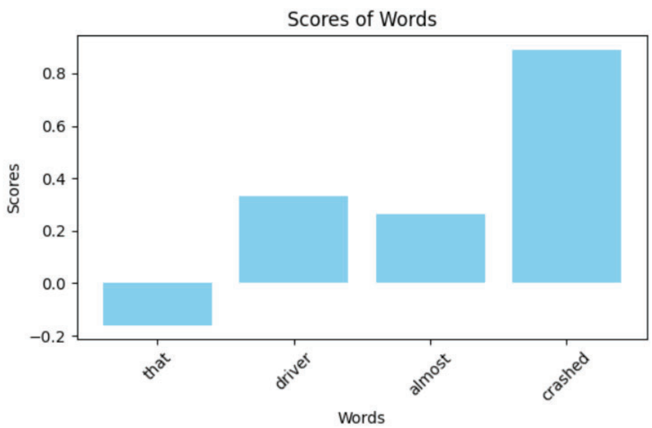

Fig. 3. The weight distribution of each word influence on the prediction.

For example, the model was provided with the input test sentence “That driver almost crashed!” and it was able to correctly identify the most influential words assigning the word “crashed” with 0.89 while “driver” got 0.33, the word “That” got -0.15, and the word “almost” was given 0.26. The words that have no direct impact on the sentence’s sentiment outcome are provided with a negative weight while the sentence with the highest impact is assigned the highest weight. To mitigate potential biases in input data, the model was presented with numerous test sentences generated by multiple evaluators. Despite this precaution, the model consistently yielded accurate and favourable results.

FUTURE WORK

The author intends to improve the pre-processing techniques used in the system to better fit LLM’s as majority of thees NLP based models are capable of understanding human language without the need for stop work removal, special character removal, etc. Another reason for this is due to the cleaning process removing emojis and emoticons from the input data where users tend to express their emotions through them. To address this challenge, a specialized module could be integrated to extract sentiment from emojis and emoticons, assigning a score based on their corresponding emotional sentiment class. This score would then be combined with the text-derived sentiment score, providing a more comprehensive understanding of the text’s overall sentiment. Furthermore, the system performance could be improved by including regularization techniques, cross-validation, and exploring hybrid models or implementing advanced fine tuning techniques to the existing model and its self-learning module. Finally as mentioned above, the system needs to incorporate a feature that allows the user to connect common words or expressions towards a common interest/despise to get a better understanding on what their fan base generally like/dislikes which can greatly help get deeper insight with the help of association rule mining.

CONCLUSION

In this paper, the author proposed an approach to classify emotional sentiment classes of social media texts on a granular level. This system aims to provide individuals and brands with a social media following the option to automate the process of analysing and understanding their followers’ overall emotional sentiment. The system is implemented by initially extracting text data from social media using their respective API and storing the collected data in the system database for cleaning and pre-processing. This stage removes URLs, user mentions, hashtags, special characters, and stop words to improve the readability of the data. The next stage consists of data tokenization and prediction of the emotional sentiments. In this stage, the DistilBERT model will perform tokenization and sentiment classification on the text data while also providing the system with the attention weights the model assigned each word for interpretation. Finally, the predicted class and the input text with the respective influence score per text are provided for sentiment analysis. Using the predicted data, the system self-improves by fine-tuning itself on the labelled class and input text allowing the model to adapt and enhance its representations over time, improving its performance and robustness without the need for human intervention.

REFERENCES

[1] S. Hermsen, J. Frost, R. J. Renes, and P. Kerkhof, “Using feedback through digital technology to disrupt and change habitual behavior: A critical review of current literature,” Computers in Human Behavior, vol. 57, pp. 61–74, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0747563215302946

[2] M. Reveilhac, S. Steinmetz, and D. Morselli, “A systematic literature review of how and whether social media data can complement traditional survey data to study public opinion,” Multimedia Tools and Applications, vol. 81, no. 7, pp. 10 107–10 142, Mar. 2022. [Online]. Available: https://link.springer.com/10.1007/s11042-022-12101-0

[3] G. Fuschillo, “Fans, fandoms, or fanaticism?” Journal of Consumer Culture, vol. 20, no. 3, pp. 347–365, Aug. 2020. [Online]. Available: http://journals.sagepub.com/doi/10.1177/1469540518773822

[4] M. Hajiali, “Big data and sentiment analysis: A comprehensive and systematic literature review,” Concurrency and Computation: Practice and Experience, vol. 32, no. 14, p. e5671, Jul. 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/cpe.5671

[5] N. Mirzh and Z. H. Ali, “Sentiment analysis techniques –survey,” Wasit Journal for Pure sciences, vol. 2, no. 2, pp. 282–290, Jun. 2023, number: 2. [Online]. Available: https://wjps.uowasit.edu.iq/index.php/wjps/article/view/152

[6] R. Jain, A. Kumar, A. Nayyar, K. Dewan, R. Garg, S. Raman, and S. Ganguly, “Explaining sentiment analysis results on social media texts through visualization,” Multimedia Tools and Applications, vol. 82, no. 15, pp. 22 613–22 629, Jun. 2023. [Online]. Available: https://link.springer.com/10.1007/s11042-023-14432-y

[7] A. Diwali, K. Saeedi, K. Dashtipour, M. Gogate, E. Cambria, and A. Hussain, “Sentiment Analysis Meets Explainable Artificial Intelligence: A Survey on Explainable Sentiment Analysis,” IEEE Transactions on Affective Computing, pp. 1–12, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10185138/

[8] E. Kotei and R. Thirunavukarasu, “A Systematic Review of Transformer-Based Pre-Trained Language Models through Self-Supervised Learning,” Information, vol. 14, no. 3, p. 187, Mar. 2023. [Online]. Available: https://www.mdpi.com/2078-2489/14/3/187

[9] A. Sharma, “Natural language processing and sentiment analysis,” International Research Journal of Computer Science, vol. 8, no. 10, pp. 237–242, Oct. 2021. [Online]. Available: https://www.irjcs.com/volumes/Vol8/iss-10/01.OCCS10080.pdf

[10] M. Wankhade, A. C. S. Rao, and C. Kulkarni, “A survey on sentiment analysis methods, applications, and challenges,” Artificial Intelligence Review, vol. 55, no. 7, pp. 5731–5780, Oct. 2022. [Online]. Available: https://link.springer.com/10.1007/s10462-022-10144-1

[11] A. Barredo Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. Garcia, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, and F. Herrera, “Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI,” Information Fusion, vol. 58, pp. 82–115, Jun. 2020. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1566253519308103

[12] R. Dwivedi, D. Dave, H. Naik, S. Singhal, O. Rana, P. Patel, B. Qian, Z. Wen, T. Shah, G. Morgan, and R. Ranjan, “Explainable ai (xai): Core ideas, techniques and solutions,” ACM Computing Surveys, vol. 55, 09 2022.

[13] W. Saeed and C. Omlin, “Explainable ai (xai): A systematic meta-survey of current challenges and future opportunities,” Knowledge-Based Systems, vol. 263, p. 110273, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0950705123000230

[14] S.-C. Huang, A. Pareek, M. Jensen, M. P. Lungren, S. Yeung, and A. S. Chaudhari, “Self-supervised learning for medical image classification: a systematic review and implementation guidelines,” npj Digital Medicine, vol. 6, no. 1, p. 74, Apr. 2023. [Online]. Available: https://www.nature.com/articles/s41746-023-00811-0

[15] S. E. Corp, “The impact of social media on music: How artists are building fanbases online,” Dec. 2023. [Online]. Available: https://medium.com/@systementcorp/the-impact-of-social-media-on-music-how-artists-are-building-fanbases-online-291ae71cb184

[16] B. Dean, “Social network usage & growth statistics: How many people use social media in 2024?” Dec. 2023. [Online]. Available: https://backlinko.com/social-media-users

[17] L. Yue, W. Chen, X. Li, W. Zuo, and M. Yin, “A survey of sentiment analysis in social media,” Knowledge and Information Systems, vol. 60, no. 2, pp. 617–663, Aug. 2019. [Online]. Available: http://link.springer.com/10.1007/s10115-018-1236-4

[18] G. Wang, “Analysis of sentiment analysis model based on deep learning,” Applied and Computational Engineering, vol. 5, no. 1, pp. 750–756, Jun. 2023. [Online]. Available: https://ace.ewapublishing.org/article.html?pk=0891596f4a3248afb5d3af5ea9a17c5e

[19] F. Souza and J. Filho, “BERT for Sentiment Analysis: Pre-trained and Fine-Tuned Alternatives,” Jan. 2022, arXiv:2201.03382 [cs]. [Online]. Available: http://arxiv.org/abs/2201.03382

[20] S. Tabinda Kokab, S. Asghar, and S. Naz, “Transformer-based deep learning models for the sentiment analysis of social media data,” Array, vol. 14, p. 100157, Jul. 2022. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S2590005622000224

[21] D. Demszky, D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, and S. Ravi, “GoEmotions: A Dataset of Fine-Grained Emotions,” Jun. 2020, arXiv:2005.00547 [cs]. [Online]. Available: http://arxiv.org/abs/2005.00547

[22] B. Chen, Q. Huang, Y. Chen, L. Cheng, and R. Chen, “Deep Neural Networks for Multi-class Sentiment Classification,” in 2018 IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS). Exeter, United Kingdom: IEEE, Jun. 2018, pp. 854–859. [Online]. Available: https://ieeexplore.ieee.org/document/8622880/

[23] H. Batra, N. S. Punn, S. K. Sonbhadra, and S. Agarwal, “Bert-based sentiment analysis: A software engineering perspective,” 2021, vol. 12923, pp. 138–148, arXiv:2106.02581 [cs]. [Online]. Available: http://arxiv.org/abs/2106.02581

[24] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter,” Feb. 2020, arXiv:1910.01108 [cs]. [Online]. Available: http://arxiv.org/abs/1910.01108