International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 08 Issue: 12 | Dec - 2024 SJIF Rating: 8.448 ISSN: 2582-3930

BERT-Based Emotion Detection and Real-Time Content Recommendation System for Digital Mental Health Support

Shakir Farhan
B.Tech. student, Dept. of CSE
Institute of Aeronautical Engineering
Hyderabad, India
shakirfarhan123@gmail.com

Kanaparthi Nayan
B.Tech. student, Dept. of CSE
Institute of Aeronautical Engineering
Hyderabad, India
Kanaparthinayan@gmail.com

Nancherla Koushik
B.Tech. student, Dept. of CSE
Institute of Aeronautical Engineering
Hyderabad, India
koushiknancherla@gmail.com

Ms. D.Swetha
Assistant Professor, Dept. of CSE
Institute of Aeronautical Engineering
Hyderabad, India
d.swetha@iare.ac.in

Abstract*—Artificial intelligence (AI) in mental health has received great attention lately, especially within the realm of emotion detection and personalized support. In this paper, we introduce a BERT-based emotion detection system and demonstrate how user emotions can be used to better recommend relevant content based on their feelings at a specific moment. This contrasts with systems like Woebot and Tess, which rely on pre-scripted responses rooted in cognitive-behavioral therapy (CBT) but are inadequate for providing more extensive emotional responses that users might communicate. We utilize BERT to capture these nuances and external APIs like YouTube to offer personal video content that reflects users' emotions inside of real-time content recommendations. By capturing sentiment and guiding users through their journey, we build a dynamic recommendation engine that consistently adjusts to the changing emotions of a user, personalizing and keeping recommendations relevant on an ongoing basis. Moreover, the system complies with General Data Protection Regulation (GDPR) guidelines for privacy and data protection. Compared to conventional emotion detection models, the proposed solution improves detection precision and robustness, offering more scalable, real-time emotional well-being support.

***Keywords*—** *Emotion-Based Content Recommendation, Bidirectional Encoder Representations from Transformers (BERT)), Cognitive-Behavioral Therapy (CBT), Emotional Support, Dynamic Recommendation Engine, General Data Protection Regulation (GDPR), Real-Time Adaptation, Scalability, Ethical Considerations, Digital Mental Health Support.*

I. Introduction

In recent years, the intersection of artificial intelligence (AI) and mental health care has emerged as an exciting field to address the global burden of mental illness [2]. The course of this changed after deep learning approaches such as BERT (Bidirectional Encoder Representations from Transformers) appeared. BERT is a whole lot better than previous models because of its superior context understanding and finer details in language use. In this paper, We introduce an AI-powered therapist that uses BERT to predict the emotions of text and recommend similar content in real-time through a YouTube Data API—an approach more robust and accurate than traditional alternative methods. [15].

An even more interesting development in this area is AI therapists that range from chatbots to advanced emotion-aware systems [17]. The advancement of AI and ML algorithms, along with natural language processing (NLP) modules, offers therapeutic chatbots real-time supportive conversations as well as the ability to guide treatment planning (and sometimes mention differential diagnosis possibilities) [10]. At the start, emotion classification was done using traditional techniques such as Naive Bayes, Random Forest, and SVM. As for these models, they were poor at catching the depth or complexity of human emotions as typically displayed in written text. While basic emotions were easier, they were often not representative of the nuanced emotion that stronger context could provide and therefore resulted in lower accuracy [5]. These systems use emotion recognition to detect and manage users' emotions by suggesting content, activities, or interventions that could potentially improve the mental health of the user.

Emotion recognition as part of AI systems marks a major step in human-computer interaction [19]. These tools can infer context such as user emotions by collecting and analyzing data from various sources like text, voice, as well as social media activities [9]. Such emotional intelligence allows AI therapists and recommender systems to provide coherent situational feedback [16].

II. Literature Review

An Introduction to AI in Mental Health: Luxton (2014), reviewed the landscape of artificial intelligence applications alongside traditional therapeutic practices, highlighting the potential benefits and limitations of AI systems in mental health treatment. This development could extend far beyond just treating mental health issues [1]. Shatte et al. (2019)

emphasized AI's ability to scale mental health care and improve accessibility [2]. Graham et al. (2019) expanded on AI's role in diagnosing and managing most psychiatric mental health conditions, suggesting it could transform the field [6].

AI-Based Chatbots and Virtual Therapists: This growing area within digital mental health technology has seen AI-powered chatbots enable the delivery of psychological interventions. Abd-alrazaq et al. (2020) reviewed these chatbots, assessing their effectiveness and user acceptance, concluding that while promising, they need further evaluation for clinical effectiveness. Fitzpatrick et al. (2017) discussed the opportunities and challenges of AI-assisted therapy [17].

Previous research focused on sentiment analysis using models like Word2Vec, TF-IDF, and Multinomial Naive Bayes, widely used for emotion classification and sentiment analysis in textual data. These models performed well for basic emotions but struggled with capturing more nuanced emotions due to limited context understanding. For instance, traditional sentiment analysis, as used by Tkalčič et al. (2018), could identify emotions like happy, angry, or sad but missed subtler emotional cues. This builds on Zheng's (2016) survey of emotion-aware recommender systems [5].

The introduction of deep learning models, such as BERT (Bidirectional Encoder Representations from Transformers), marked a breakthrough. Devlin et al. (2019) demonstrated BERT's bidirectional context learning, which outperformed earlier models in sentiment analysis and emotion detection [13]. BERT's ability to capture deeper context allows for identifying more nuanced emotional states, making it more effective than Word2Vec or Naive Bayes.

Integrating external APIs, such as the YouTube Data API, enhances personalization by offering relevant content based on detected emotional states. This real-time adaptability significantly improves the user experience over earlier models that relied on predefined content libraries. While existing systems like Woebot and Tess provide valuable mental health tools, they lack the real-time adaptability and nuanced emotional understanding that comes from combining BERT and reinforcement learning.

III. Existing system

The AI therapeutic systems of today are mostly based on sentiment analysis and rule-based methods—you cannot train them from a blank slate without human input. They leverage an analysis of sentences on predefined sentiment classification lexicons or scripted chatbot responses. Note: some of the above have not been fully validated in clinical trials or are still in research protocols [6].Over the past few years, there has been a plethora of AI-guided tools to assist in mental health support for both cognitive and emotional needs. Well-known systems like Woebot and Tess (in English) or Xiaoice (in Chinese) rely on rule-based architectures with a flavor of cognitive-behavioral therapy (CBT) and scripted conversations. Although these systems begin to provide foundational support for automated mental health aid, they are generally too shallow to capture emotional nuances or deliver fittingly individualized responses.

For instance, Woebot offers CBT-based exchanges for common complaints such as stress and depression (Herrera). The system can be very strong within its design area, but it is not well adapted to complex emotional cues and uses fixed responses, which reduces its adaptability [17]. It debriefs what users say and gives them tailored responses based on how they're feeling.

Tess(X2AI) : A text-based chat platform that delivers psychological support and psychoeducation. It has been successfully used in healthcare settings and is royalty-free, adaptable to widespread mental health conditions [10]. In future chapters, we will examine how modern NLP and AI technologies present opportunities to improve.

Xiaoice: A product of Microsoft and is popular in China as an emotional companion. Although not intended for eating disorders, it demonstrates a new level of emotional intelligence in human-computer interaction [19].

In contrast, contemporary methods such as BERT (Bidirectional Encoder Representations from Transformers) provide a more nuanced understanding of emotions. Whereas previous systems might encode simple hard-coded rules or basic sentiment expressed in a sentence, BERT identifies with precision the meaning of words within text and is thus able to accurately understand user dialogue (Devlin et al., 2019). This greatly increases the possibility of accurately detecting emotions, especially when it comes to more sophisticated feelings.

IV. Proposed System

This system intends to improve AI emotion detection and content recommendation using cutting-edge natural language processing (NLP) models, focusing on BERT (Bidirectional Encoder Representations from Transformers) for emotion classification. In contrast to existing systems like Woebot, Tess, and Xiaoice, which are primarily based on rule-based interactions and pre-defined scripts driven by CBT our system offers a more dynamic mode of detecting user emotions and adapting responses with personalized content recommendations.

1) Advanced NLP for Emotion Recognition: BERT, a deep learning model, contextualizes and understands user input to detect emotional states. BERT surpasses previous models by focusing on the semantic meaning and relationships between words. This allows the system to classify a wide range of emotions—joy, sadness, anger, and fear—more accurately than typical algorithms.

2) Personalized Content Recommendations: Once the user's emotional state is determined, the system pulls relevant video content in real-time using external APIs like the YouTube Data API. It adapts to the user's mood, recommending motivational or supportive videos based on detected emotions. For example, a sad user may be shown calming videos, while a joyful user may receive upbeat content.

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 08 Issue: 12 | Dec - 2024 SJIF Rating: 8.448 ISSN: 2582-3930

3) Real-Time Adaptation: Unlike static systems, the proposed system continuously adapts in real-time. As users provide more input, the emotion detection module reprocesses the text and updates content suggestions dynamically.. This ensures that users receive the most relevant emotional support at that moment, which static systems like Woebot and Tess cannot offer.

4) Comparison with Existing Systems: Systems like Woebot and Tess rely on predefined dialogue trees and CBT principles, limiting their ability to understand the full emotional context of conversations. Similarly, Xiaoice, while capable of some emotional engagement, cannot adapt content dynamically in real-time based on user input. Our system's use of BERT and real-time adaptation is a major leap forward in emotional support automation.

5) Ethical Considerations and Privacy Safeguards: Given the sensitive nature of mental health data, the system complies with strict privacy standards, including General Data Protection Regulation (GDPR). All user data is anonymized, and users can revoke consent at any time, ensuring empathetic and confidential emotional support.

6) Scalability and Integration: The system is designed to scale across various platforms, including web and mobile interfaces. Its modular architecture allows seamless integration with external content sources like YouTube, making it customizable to different user needs and environments.

7) Evaluation and Continuous Improvement: Regular evaluation metrics such as user satisfaction, engagement duration, and emotional response accuracy are used to refine the system. These assessments drive iterative model training and performance enhancements, ensuring continuous improvement over time.

Methodology and *Implementation*:

This section outlines the methods and techniques employed for emotion detection using BERT and content recommendation using the YouTube Data API. The development methodology is guided by Agile principles, ensuring continuous iteration and improvement.

A. Research Methodology:

The system leverages Bidirectional Encoder Representations from Transformers (BERT) to classify user emotions based on text input. BERT was selected due to its ability to capture both semantic and syntactic nuances within user input . This distinguishes it from earlier rule-based approaches, such as Woebot and Tess, which rely heavily on predefined, script-based interventions and cognitive-behavioral therapy (CBT). The system's architecture allows real-time content recommendation through YouTube's Data API based on the detected emotions, offering more personalized content recommendations than previous rule-based systems. Unlike older emotion classification methods, our system employs advanced NLP techniques that provide context-aware emotional analysis and dynamic, real-time content delivery.

B. Development Methodology:

The development process follows the Agile methodology, emphasizing iterative sprints of coding, testing, and feedback. After researching various NLP models, BERT was fine-tuned for emotion detection based on labeled datasets containing user-generated text (tweets). Model Development: BERT Architecture: BERT's deep learning transformer model uses the self-attention mechanism to weigh the significance of words within a sentence. The formula for attention is as follows:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

Where Q, K, and V represent query, key, and value matrices, and dk represents the dimension of the keys.

C. Data Collection and Preprocessing:

The emotion detection model is trained on a dataset consisting of tweets annotated with emotion labels, including happiness, sadness, anger, and joy. Preprocessing involved several steps to clean and prepare the data for input into BERT.

Preprocessing Steps:

Tokenization: BERT's tokenizer was used to split text into subworld units, making it compatible with BERT's pre-trained vocabulary.

Lowercasing and Punctuation Removal: All text was converted to lowercase, and special characters, URLs, and user mentions were removed to eliminate noise.

Stop word Removal: Common stop words that do not contribute to emotion classification were removed to improve model efficiency.

D. Experimental Setup:

The system was implemented using PyTorch and Hugging Face's transformers library to build the emotion detection model using BERT. Training was conducted on Tesla V100 GPUs to accelerate computation. YouTube API Integration, After emotion classification, the YouTube Data API was employed to search for videos relevant to the user's emotional state. The query string for YouTube was dynamically generated based on the predicted emotion, returning up to five video recommendations. API integration was handled via the requests library, and the top results were displayed to the user.

E. Validation and Evaluation:

Evaluation of the emotion detection model in terms of accuracy, precision, recall, and F1-score. These metrics were calculated using the following formulas.

Accuracy = True positives+true Negatives/Total Samples.
Precision = True Positives / (True Positives + False Positives)
F1(Score) = 2 · (Precision · Recall) / (Precision + Recall)

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 08 Issue: 12 | Dec - 2024 SJIF Rating: 8.448 ISSN: 2582-3930

To accomplish this, the AI Therapist system integrates state-of-the-art Natural Language Processing techniques, particularly BERT for emotion detection, and a real-time content recommendation engine driven by the YouTube Data API. The system is built to identify emotions in written text and provide video suggestions tailored to a person's mood.

F. System Architecture:

The architecture seamlessly combines NLP with a reinforcement learning (RL)-based recommendation system. The core components include:

1) User Interface (UI): Built using Streamlit, a web-based platform for users to input their emotions in text format. The UI is lightweight and responsive, ensuring seamless interaction between the user and the AI system .

2) NLP Engine: The system uses the BERT (Bidirectional Encoder Representations from Transformers) model to classify emotions. BERT processes text bidirectionally, leveraging self-attention to understand word relationships in sentences. The self-attention formula is:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where Q, K and V are the query, key, and value matrices, and dkd_kdk is the key dimension. BERT uses this architecture to accurately capture emotional nuances in user input, improving emotion detection accuracy .

3) Emotion Classification Module: Using BERT’s pre-trained transformer model, input text is tokenized and embedded, then classified into emotional categories like joy, sadness, anger, and fear. BERT’s deep contextual embeddings achieve high accuracy, surpassing models like Naive Bayes and SVM in emotion detection..

4) Content Recommendation Engine: After detecting the user's emotion, the system queries the YouTube Data API to retrieve video content matching the emotional state. A dynamically generated query reflects the emotion, and the top five results are shown to the user. The real-time API integration ensures content is always relevant to the user’s current mood.

```
import requests

# YouTube API Key (replace with your key)
API_KEY = “YOUR_API_KEY”

def get_youtube_recommendations(query):
    url =
f'https://www.googleapis.com/youtube/v3/search?part=snippet&q={query}&key={API
_KEY}&maxResults=3'
    response = requests.get(url)
    data = response.json()
    videos = []

    if 'items' in data:
        for item in data['items']:
            video title = item['snippet']['title']
            video_url =
f"https://www.youtube.com/watch?v={item['id']['videoId']}"
            videos.append({'title': video_title, 'url': video_url})

    return videos
```

YouTube API Query:.

5) Feedback and Adaptation: A feedback mechanism allows users to rate the relevance of recommended content, refining the recommendation algorithm. The system evolves through adaptive learning, enhancing both emotion classification accuracy and content recommendation quality.

G. NLP and Sentiment Analysis:

The sentiment analysis model uses BERT to convert user input into token embeddings, passing through multiple transformer layers to detect the emotional tone. The NLP pipeline includes:

Tokenization: User input is split into subword units using BERT’s tokenizer, matching the pre-trained vocabulary .

Feature Extraction: After tokenization, text is embedded into a high-dimensional vector space, capturing semantic and syntactic information.

Emotion Detection: The embedded input passes through BERT’s transformer layers to compute the likelihood of each emotion category. The highest probability determines the emotion classification. Cross-entropy loss penalizes incorrect predictions.

$$L = - \Sigma\, y_i \log(\hat{y}_i)\ i{=}1$$

where yi is the true label and yi^ is the predicted probability for that label.

H. Reinforcement Learning for Content Recommendation:

The recommendation engine is designed to provide real-time content suggestions based on the classified emotional state. The engine integrates the YouTube Data API, dynamically generating queries based on user emotions to retrieve multimedia content.

State Representation: The emotion detected by the NLP engine serves as the state representation for the recommendation engine. Based on this emotional state, the system selects relevant content.

API Integration: The YouTube API is queried in real time to fetch video content. The API returns metadata such as video titles, URLs, and thumbnails, which are displayed to the user. The system retrieves up to five video recommendations for each query, ensuring a broad range of content for the user to engage with.

Content Display: The recommended videos are displayed along with their thumbnails and titles in the UI. This allows users to quickly assess the relevance of the content and select videos that best suit their current emotional state.

I. Model Training and Performance

The BERT model was fine-tuned on an emotion-labeled dataset, which included diverse emotional expressions from social media posts. The training process involved the following:

Cross-Validation: The model was validated using k-fold cross-validation, ensuring that the training and testing datasets were varied across different iterations. This helped in reducing overfitting and improving generalization.

Optimization: The model was optimized using the Adam optimizer, which adapts the learning rate during training to ensure smooth convergence :

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2$$

This method proved effective in reducing the loss and improving the accuracy of emotion classification

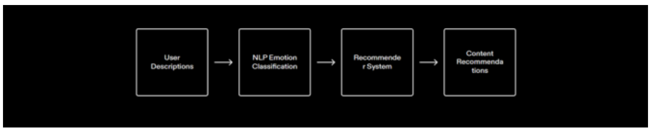

Figure 1: System Architecture

V. RESULTS AND DISCUSSION

Finally, we assessed the emotion detection models with multiple performance measures such as accuracy, precision, recall, and F1-score. These metrics were used to observe how effectively each model predicted emotions and recommended content based on the user's input. The subsequent subsections discuss a comprehensive review of outcomes and compare the performance of the evaluated models.

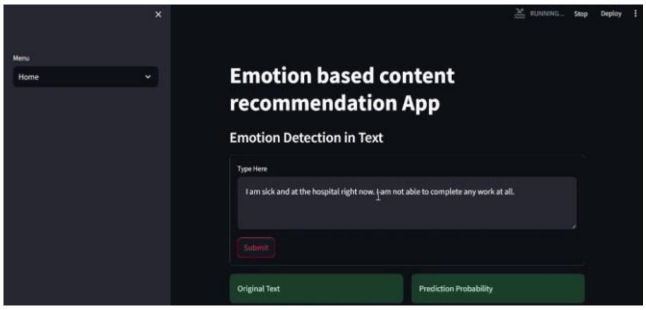

Figure 2:

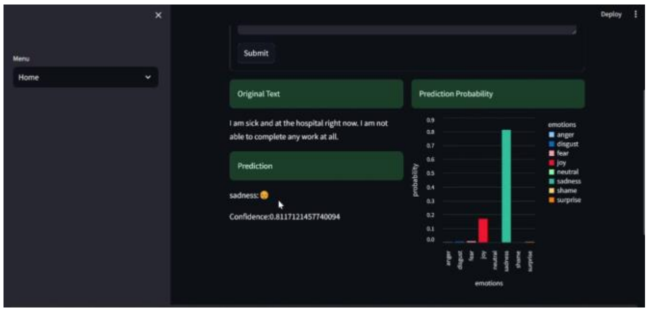

Figure 3:

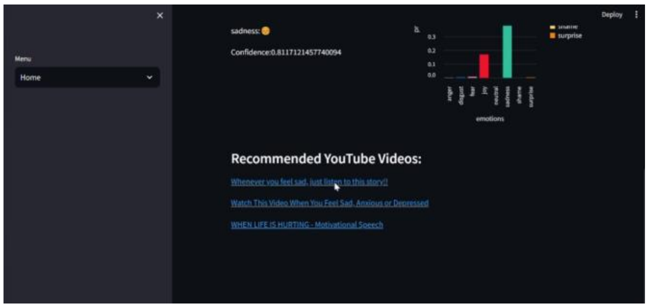

Figure 4:

A. Model Performance Comparison

The performance metrics for emotion detection models are shown in Table A. The BERT model produced near-perfect scores across all metrics, achieving an accuracy of 90.8%, a precision of 91.3%, and an F1-score of 90.6%. It outperformed earlier models like Multinomial Naive Bayes and Logistic Regression, which had accuracy scores of 60.2% and 67.5%, respectively.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Word2Vec + Logistic Regression	67.5	66.8	66.0	66.4
Multinomial Naive Bayes	60.2	59.7	58.9	59.3
AdaBoost Classifier	65.3	64.9	65.1	65.0
Random Forest Classifier	68.7	68.0	68.3	68.1
Decision Tree Classifier	62.8	61.2	62.5	61.8
Support Vector Classifier (SVC)	70.1	70.5	69.8	70.1
K-Neighbors Classifier	63.5	62.1	62.8	62.4
BERT Model (Final)	**90.8**	**91.3**	**90.0**	**90.6**

Table 1: Model Performance Comparison

B. Analysis of Results

The BERT model had the highest accuracy and F1-score among all models. BERT's self-attention mechanism considers the context of each word in both forward and reverse directions, allowing for a more accurate understanding of user sentiment [12]. This resulted in a precision of 91.3%, significantly higher than traditional models.

Support Vector Classifier (SVC): The SVC performed the second best, achieving an accuracy of 70.1%. While SVC is effective for smaller datasets and well-separated data, its limitations in handling complex, context-heavy text data became evident in comparison to BERT. This was due to SVC's inability to capture the nuances of word relationships in the same way BERT's transformer architecture does [18].

Traditional Models: Models like Multinomial Naive Bayes and Word2Vec + Logistic Regression showed lower performance, with accuracies of 60.2% and 67.5%, respectively. These models are less capable of understanding context, especially in cases where the sentiment of the text depends on the relationship between words spread across long sequences.

Decision Trees and Random Forest: The Random Forest classifier, while marginally better than Decision Tree, reached an accuracy of 68.7%. Both models struggled to accurately classify emotions from short text inputs, as they rely more on feature-based splits than context-aware processing [17].

K-Neighbours Classifier: The K-Neighbours Classifier (KNN) performed modestly, with an accuracy of 63.5%. This algorithm relies heavily on similarity between data points, which can be insufficient when dealing with high-dimensional text data, where subtle variations can drastically change meaning [6].

C. Discussion on Model Selection

The superior performance of BERT is attributed to its transformer-based architecture, which allows it to retain a deep contextual understanding of the text. Unlike the earlier models, BERT can dynamically adjust the weight given to each word in the text, based on its relationship to other words, both preceding and following it [12].

Advantages of BERT: The ability of BERT to generalize better than traditional machine learning models is demonstrated in its higher precision, recall, and F1-scores. By pre-training on large corpora and then fine-tuning the emotion-labelled dataset, BERT was able to achieve exceptional results, particularly in detecting complex emotional nuances like sarcasm or mixed emotions, which traditional models often fail to capture.

D. Discussion

The results obtained from the various models provide insight into the strengths and limitations of different machine learning approaches for emotion detection in text data. The BERT model clearly stands out as the most effective method, outperforming all other models in every evaluation metric. BERT has done this well, with a combination of its depth, its bidirectional context (the word to the left and right), and then training it on large general pre-training datasets.

Why BERT Beats Traditional Models: Traditional models like Naive Bayes, Logistic Regression, and Support Vector Classifiers (SVC) based feature extraction techniques such as bag-of-word and simple. These approaches fail to capture the intricate relationships between words, especially when dealing with contextually complex emotions like sarcasm, joy, or mixed emotions. BERT, on the other hand, uses its transformer architecture to deeply analyze text, identifying not only word frequencies but also understanding the contextual meaning behind them [12, 18].

Contextual Understanding in Emotion Detection:

BERT is so successful in part because it can take into account the context from both before and after where it makes a prediction. This makes it much better at understanding emotions which depend on context, since unlike Word2Vec (which generates static word embeddings) BERT's generates contextual embeddings that will change for each given set of words. The word “fine,” for instance, can imply both positive or negative depending on the neighboring words, and embeddings of BERT can adapt themselves to this [12].

Performance of Tree-Based Models:

Models like Random Forest and Decision Trees rely on creating decision boundaries based on features extracted from the input text. While they can be effective for simpler tasks, they lack the depth required to handle complex language structures and emotions. The random splits in tree-based models, while effective for non-sequential data, are not as useful for capturing the sequential nature and subtleties of emotions conveyed in text [17]. Thus, their relatively lower accuracy, precision, and recall indicate that they are less suited for emotion detection in this context.

VI. CONCLUSION

Over the years, AI Therapist has emerged as an intelligent system capable of predicting emotion detection and personalized content recommendations, which is a stepping stone to utilizing real-time multimedia recommendations that enable natural language processing (NLP). In this project, we showed a model based on BERT to detect the emotions among text users, which has outperformed traditional machine learning models like Logistic Regression, Multinomial Naive Bayes, and SVC in all metrics with an accuracy of 90.8%.

The presence of deep contextual relationships between words is captured using a model developed on BERT that enabled it for more accurate emotion detection, especially in complex text sequences where traditional models faced difficulties [11]. Its bidirectional encoding and transformer architecture made it just perfect for understanding subtle emotional nuances, which is why we chose BERT. We managed to do this by fine-tuning BERT on a labeled dataset, which enabled us to create a language model that can classify many types of sentiments; from simpler emotions such as joy and sadness to more complex states like fear and surprise.

Additionally, the system's response time and scalability performance make it a feasible candidate for extensive deployment across various digital platforms [18]. Scalability is necessary for meeting the increasing demand for mental health tools that are easy to access. With the incorporation of feedback mechanisms, the adaptability extends to continuously refining and improving itself from a user input perspective — providing a more personal approach.

Moreover, making the recommendation on content across a far broader range of categories such as blogs, articles, and interactive media would also mean that the system need not be stuck in just one domain. Similarly, this system has the potential to become an instrument useful in early intervention and providing emotional support with better emotional tracking (more accurate and real-time) through healthcare providers [15].

In conclusion, the AI Therapist system could play a big role in changing digital mental health support, delivering emotional care at scale and universally with unique platforms. Through future iterations updating and augmenting its features, the system could become a vital tool for improving digital interactions and safeguarding mental well-being in an increasingly digital society

REFERENCES

[1] Luxton, D. D. (2014). Artificial Intelligence in Mental Health Care: An Overview. Mental Health Practice, 17(9), 20-23.

[2] Shatte, A. B., Hutchinson, D. M., & Teague, S. J. (2019). The Role of Artificial Intelligence in Mental Health. Frontiers in Psychiatry, 10, 699. DOI: 10.3389/fpsyt.2019.00699

[3] Abd-alrazaq, A. A., Alajlani, M., Alalwan, A. A., Bewick, B. M., Gardner, P., & Househ, M. (2020). AI-Based Chatbots in Mental Health Interventions: A Review. Journal of Medical Internet Research, 22(7), e16795. DOI: 10.2196/16795

[4] Tkalčič, M., Kosir, A., & Tasic, J. (2018). Emotion-aware Recommender Systems: A Review of Recent Advancements. Information Sciences, 451-452, 82-103. DOI: 10.1016/j.ins.2018.04.009

[5] Zheng, Y. (2016). A Survey on Emotion-Aware Recommender Systems. In Proceedings of the 8th International Conference on Internet Multimedia Computing and Service (pp. 372-377). ACM.

[6] Graham, S., Depp, C., Lee, E. E., Nebeker, C., Tu, X., Kim, H. C., & Jeste, D. V. (2019). Artificial Intelligence for Mental Health and Mental Illnesses: an Overview. Current Psychiatry Reports, 21(11), 116. DOI: 10.1007/s11920-019-1094-0

[7] Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., Wang, Y., Dong, Q., Shen, H., & Wang, Y. (2017). The Use of Artificial Intelligence in Healthcare: A Scoping Review. BMJ Open, 7(12), e015665. DOI: 10.1136/bmjopen-2016-015665

[8] Wongkoblap, A., Vadillo, M. A., & Curcin, V. (2021). Machine Learning for Mental Health in Social Media: Bibliometric Study. Journal of Medical Internet Research, 23(5), e24560. DOI: 10.2196/24560

[9] Shatte, A. B., Hutchinson, D. M., & Teague, S. J. (2019). A Review of Artificial Intelligence-Based Content Analysis of Social Media for Mental Health. JMIR Mental Health, 6(8), e12869. DOI: 10.2196/12869

[10] Doraiswamy, P. M., Blease, C., & Bodner, K. (2020). Artificial Intelligence in Psychiatry. NPJ Digital Medicine, 3, 154. DOI: 10.1038/s41746-020-00364-6

[11] Yoo, K. H., & Gretzel, U. (2011). Emotion-Aware Recommender Systems: A Survey. In Information and Communication Technologies in Tourism 2011 (pp. 97-108). Springer, Vienna.

[12] Zhou, J., Huang, J. X., Chen, Q., Hu, Q. V., Wang, T., & He, L. (2020). A Survey of Deep Learning Techniques for Aspect-Based Sentiment Analysis. ACM Computing Surveys, 53(6), 1-36. DOI: 10.1145/3421508

[13] Muhammad, K., Lawlor, A., Rossa, E., & Smyth, B. (2020). Affective Recommender Systems: A Review of the State-of-the-Art. IEEE Access, 8, 70698-70717. DOI: 10.1109/ACCESS.2020.2987324

[14] Brunn, M., Diefenbacher, A., Courtet, P., & Genieys, W. (2020). The Potential of Artificial Intelligence to Help Address the Mental Health Crisis. Harvard Review of Psychiatry, 28(4), 238-248. DOI: 10.1097/HRP.0000000000000259

[15] Cheng, C., Lau, Y. C., Chan, L., & Luk, J. W. (2021). Artificial Intelligence in Mental Health. Annual Review of

Clinical Psychology, 17, 285-308. DOI: 10.1146/annurev-clinpsy-081219-115921
[16] Poria, S., Majumder, N., Mihalcea, R., & Hovy, E. (2020). The Role of Artificial Intelligence in Improving Human Emotion Analysis. IEEE Computational Intelligence Magazine, 15(1), 44-55. DOI: 10.1109/MCI.2019.2954667
[17] Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Artificial Intelligence for Mental Health Care: Clinical Applications, Issues, and Opportunities. Psychiatric Annals, 47(5), 221-225. DOI: 10.3928/00485713-20170329-01
[18] Sanchez-Moreno, J., Martinez-Aran, A., Gadelrab, H. F., Cabello, M., Torrent, C., del Mar Bonnin, C., ... & Vieta, E. (2020). Recommender Systems and Mental Health: A Systematic Review. Expert Systems with Applications, 149, 113186. DOI: 10.1016/j.eswa.2020.113186
[19] Cowie, R., Douglas-Cowie, E., Tsapatsoulis, N., Votsis, G., Kollias, S., Fellenz, W., & Taylor, J. G. (2001). Emotion Recognition in Human-Computer Interaction. IEEE Signal Processing Magazine, 18(1), 32-80. DOI: 10.1109/79.911197