




International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 4, April 2025

Understanding Emotions through NLP

G.S.Uday Kumar^[1] A.A.Armanulla Khan^[2] B.R.Latha^[3] K.A.Umar Khaleefa^[4] Gopi Bleesy^[5]

Assistant Professor, Dept. of CSE., Kuppam Engineering College, JNTUA University, Kuppam, A.P., India^[1]

UG Students, Dept. of CSE., JNTUA University, Kuppam Engineering College, Kuppam Engineering College, JNTUA
University, Kuppam, A.P., India^[2] to ^[5]

ABSTRACT: Natural Language Processing (NLP), a subset of Artificial Intelligence (AI), has emerged as a powerful tool in analyzing human language to detect underlying mental health issues such as depression and suicidal tendencies. By leveraging techniques in machine learning and deep learning, NLP systems can process and interpret textual data from diverse sources like social media, forums, and online journals to identify patterns and linguistic markers associated with mental health conditions. These markers include the use of emotionally charged words, changes in writing style, expressions of hopelessness, isolation, or distress, which are critical indicators of a person's psychological state. The process typically involves data preprocessing, feature extraction using models such as TF-IDF or word embeddings (e.g., Word2Vec, GloVe), and classification using supervised learning techniques or deep learning models like BERT or LSTM.

Understanding human emotions is a critical aspect of human-computer interaction, mental health diagnostics, and customer service applications. This project focuses on leveraging Natural Language Processing (NLP) techniques to analyze and interpret emotions expressed in textual data. By utilizing state-of-the-art machine learning models, including deep learning architectures such as recurrent neural networks (RNNs) and transformers, the project aims to develop models capable of detecting and classifying emotional states such as joy, sadness, anger, fear, and surprise from written content. The research investigates various aspects of emotion detection, including sentiment analysis, contextual understanding, and the role of linguistic features in emotion recognition. Additionally, it explores challenges like sarcasm, ambiguity, and cultural differences in emotional expression. The goal is to create an emotionaware system that can improve applications in fields such as automated customer support, mental health monitoring, and social media analysis, enabling more empathetic and responsive interactions between humans and machines. Through this work, we seek to enhance the understanding of emotions in textual communication, providing valuable insights into emotional dynamics and improving human-computer interactions.

I. INTRODUCTION

Human language is rich, complex, and inherently emotional. Emotions play a vital role in how people communicate, make decisions, and form relationships. As digital communication increasingly becomes the norm—through emails, social media, messaging apps, and online platforms—the ability to automatically detect and interpret emotional content in text has gained substantial importance. This has led to a growing interest in the field of emotion recognition using Natural Language Processing (NLP), which focuses on enabling machines to understand and interpret human emotions from textual data. NLP, a subfield of artificial intelligence (AI), deals with the interaction between computers and human language. While early efforts in NLP focused primarily on tasks like grammar correction, translation, and keyword extraction, there is now a shift towards more contextaware and human-centric language understanding, with emotion detection becoming a key research area. Unlike traditional sentiment analysis—which typically categorizes text into positive, negative, or neutral—emotion recognition delves deeper, identifying specific affective states such as joy, anger, sadness, fear, disgust, and surprise. The growing interest in this field is driven by both technological and societal factors. On the technological front, the advancement of deep learning, pre-trained language models like BERT and GPT, and access to large-scale annotated datasets have made emotion recognition more feasible and accurate than ever before. On the societal side, the demand for emotionally intelligent systems is rising in sectors such as customer service, healthcare, education, entertainment, and mental health monitoring.

Emotion recognition through NLP has the potential to revolutionize the way machines understand human needs and contexts. For instance, in customer service, emotion-aware chatbots can provide empathetic responses, leading to better user experiences. In mental health, analyzing patients' written communication can help detect early signs of distress, offering valuable support and intervention. In education, adaptive learning systems can adjust content delivery based on a student's emotional state. However, detecting emotions from text is a complex task. Emotions are often implicit, culturally nuanced, and context-dependent. People may use sarcasm, slang, or ambiguous expressions, making it difficult for machines to accurately infer emotional states. As a result, emotion recognition remains a challenging yet exciting frontier in NLP research. This paper (or study/project) aims to explore the methods, challenges, and applications of emotion detection in text using NLP, highlighting its growing importance in developing emotionally intelligent AI systems and fostering more natural human- computer interactions.

Language is one of the most powerful tools for expressing human emotions. Whether it's joy, anger, sadness, or excitement, our words often reflect how we feel inside. In the digital age, people communicate more than ever through text—on social media, in emails, messages, and more. This presents a unique opportunity to understand human emotion at scale using Natural Language Processing (NLP), a field that allows machines to understand and interpret human language. Natural Language Processing is a branch of artificial intelligence that focuses on the interaction between computers and human language. It enables machines to read, understand, and derive meaning from human text. While early NLP models focused on basic tasks like grammar correction or keyword extraction, recent advancements have pushed its capabilities toward much deeper tasks, such as sentiment analysis and emotion detection.

It's important to differentiate between sentiment and emotion in NLP. Sentiment analysis typically classifies text as positive, negative, or neutral. Emotion detection goes a step further, aiming to recognize specific emotional states like happiness, anger, fear, disgust, surprise, or sadness. This deeper level of analysis allows for a richer understanding of human communication.

Emotion detection through NLP has a wide range of applications. Businesses use it to understand customer feedback and improve user experience. Mental health platforms may use emotion analysis to detect signs of distress in user messages. Social media monitoring tools can identify trends in public mood, while conversational AI systems like chatbots can adapt their responses based on a user's emotional state.

Emotion detection in NLP often involves a combination of machine learning, deep learning, and lexicon-based approaches. Pre-trained models like BERT, RoBERTa, and GPT can be fine-tuned for emotion recognition tasks. These models analyze context, tone, and word usage to infer emotion, while lexicons provide predefined lists of emotion-bearing words.




Fig 1 Understanding Emotions Through NLP

II. RELATED WORK

The field of Natural Language Processing (NLP) has seen significant growth in the area of emotion detection over the past two decades. Early research in this domain primarily focused on sentiment analysis, which laid the foundation for more nuanced emotion recognition. Pioneering studies, such as those by Pang and Lee (2008), explored machine learning techniques for classifying text into positive, negative, and neutral sentiments, providing a basis for further emotional granularity.

Later works expanded beyond simple sentiment to detect specific emotional categories. For instance, the use of Ekman's six basic emotions—happiness, sadness, anger, fear, disgust, and surprise—became a popular framework in NLP emotion studies. Strapparava and Mihalcea (2008) introduced SemEval-2007 Task 14, which aimed at recognizing emotions in news headlines using lexical and semantic features. This benchmark spurred the development of lexicon-based approaches like WordNet-Affect and NRC Emotion Lexicon (Mohammad and Turney, 2013), which have since been widely used in rule-based systems.

With the rise of deep learning, researchers began leveraging models like CNNs, RNNs, and LSTMs for emotion detection. Socher et al. (2013) demonstrated how recursive neural networks could capture sentiment at a phrase level, enabling finer emotional understanding. More recently, transformer-based models such as BERT and RoBERTa have been fine-tuned for emotion classification tasks, achieving state-of-the-art results on various datasets including GoEmotions (Demszky et al., 2020), which offers 27 emotion labels for high granularity.

Social media platforms like Twitter have become rich sources of emotional data. Several studies, including those by Mohammad (2015) and Abdul-Mageed and Ungar (2017), have used Twitter corpora to train emotion classifiers capable of understanding informal, short-form text. These platforms also inspired the development of distant supervision techniques, where hashtags and emojis are used as weak labels for emotional states.

Multilingual and cross-lingual emotion detection has also gained attention, especially in non-English contexts. Works like Buechel and Hahn (2018) proposed emotion annotation transfer models to handle resource-scarce languages. Emotion recognition in dialogues is another growing subfield, with datasets like IEMOCAP, MELD, and DailyDialog enabling emotion tracking in multi-turn conversations.

Despite the progress, challenges remain, especially in detecting sarcasm, mixed emotions, and cultural nuances in emotional expression. Nonetheless, the integration of contextual embeddings, transfer learning, and multimodal approaches continues to push the boundaries of what's possible in emotion-aware NLP.

The study of emotions in language processing has its roots in the broader domain of sentiment analysis. Early work focused primarily on classifying text into basic sentiments such as positive, negative, and neutral. Researchers like Pang and Lee (2008) used machine learning techniques including Naïve Bayes and Support Vector Machines (SVMs) to classify sentiment in movie reviews. These foundational studies paved the way for more fine-grained emotion recognition by highlighting the importance of linguistic features in identifying human affect.

As the field matured, researchers began to explore the detection of specific emotional categories beyond binary sentiment. Paul Ekman's model of six basic emotions—anger, fear, joy, sadness, disgust, and surprise—became widely adopted as a framework for emotion classification. One significant contribution was the SemEval-2007 Task 14 by Strapparava and Mihalcea, which aimed to classify emotions in news headlines. This task demonstrated that emotional meaning can be extracted from concise text using lexical and statistical features.

Lexicon-based approaches also gained popularity in early emotion detection research. Lexicons such as WordNet-Affect and the NRC Emotion Lexicon by Mohammad and Turney (2013) mapped words to specific emotional categories. These resources were crucial in rule-based systems, allowing researchers to perform emotion classification without requiring large annotated datasets. While limited in handling context, they were foundational in enabling scalable emotion detection in different domains.

III. PROPOSED ALGORITHM

To accurately detect emotions from textual data, we propose a hybrid approach that integrates pre-processing, deep contextual embedding, and a multi-label classification model. The algorithm aims to identify fine-grained emotions from user-generated content, such as tweets, reviews, or dialogues. Below is a step-by-step outline of the proposed system.

Data Collection and Annotation

The first step involves collecting emotion-labeled text data from relevant sources. Publicly available datasets such as GoEmotions, ISEAR, or MELD can be used. These datasets provide diverse emotional labels such as joy, anger, fear, sadness, surprise, and more. In cases where no labeled data is available, weak supervision using emojis, hashtags, or lexicons can be used to create pseudo-labels.

Text Preprocessing

The raw textual data is cleaned to remove noise such as URLs, special characters, repeated characters, and stopwords. Tokenization and lowercasing are applied, followed by optional lemmatization. For social media data, additional steps like emoji translation, hashtag segmentation, and slang normalization are incorporated to preserve emotional content.

Contextual Embedding Using Pre-trained Language Models

We employ a pre-trained transformer-based language model such as BERT (Bidirectional Encoder Representations from Transformers) to convert input text into contextualized vector embeddings. These embeddings capture the semantic and syntactic features of the text and are sensitive to the context in which words appear, which is crucial for detecting subtle emotional cues.

Emotion Classification Layer

The BERT-generated embeddings are passed to a multi-label classification layer. Since a single sentence can express multiple emotions, the output layer uses a sigmoid activation function instead of softmax, allowing multiple emotion probabilities to be predicted simultaneously. This layer can be implemented using fully connected (dense) layers followed by dropout for regularization.

Model Training and Optimization

The model is fine-tuned on the labeled dataset using a binary cross-entropy loss function appropriate for multi-label classification. Adam optimizer and a learning rate scheduler are used for efficient convergence. Stratified k-fold cross-validation is employed to ensure generalization and robustness of the model across different subsets of the data.

Emotion Score Aggregation and Interpretation

The output of the classification layer is a set of probabilities for each emotion. A threshold (e.g., 0.5) is applied to determine which emotions are present in the text. These results can be further aggregated to analyze emotion trends, perform user profiling, or provide real-time emotional feedback in applications like chatbots or mental health monitors.

Evaluation Metrics

Performance is evaluated using metrics suitable for multi-label tasks, including Precision, Recall, F1-Score, and Hamming Loss. ROC-AUC for each emotion class is also reported. Comparisons with baseline models such as logistic regression or LSTM networks are included to validate the effectiveness of the proposed transformer-based approach.

The proposed system is developed to detect signs of depression and suicidal ideation in text using deep learning techniques, with a focus on Natural Language Processing (NLP). The architecture is centered around a transformer-based model, specifically a pre-trained BERT (Bidirectional Encoder Representations from Transformers) model, which is fine-tuned for classification tasks involving mental health-related texts.

The process begins with data acquisition, where publicly available datasets are gathered. These datasets typically consist of posts labeled based on mental health states, either manually annotated or self-reported by users. Examples of

such datasets include Reddit mental health datasets and benchmark corpora used in CLPsych shared tasks. The collected text data is then cleaned through preprocessing techniques to ensure consistency and relevancy. This involves steps such as converting text to lowercase, removing URLs, emojis, punctuation, and stopwords, and handling contractions. Tokenization is then performed using BERT's own tokenizer, which breaks down text into sub-word units and adds special tokens like [CLS] and [SEP] required for the BERT model input.

After preprocessing and tokenization, the text is converted into input embeddings compatible with the BERT model. Each token is mapped to a high-dimensional vector, and positional embeddings are added to retain the order of words. These embeddings are then passed into the BERT model, which processes them through multiple transformer layers. Each layer consists of self-attention mechanisms and feedforward neural networks, allowing the model to capture complex dependencies and contextual meanings across the entire text. Unlike traditional models that rely on a fixed window of context, BERT's bidirectional architecture enables it to learn from both preceding and succeeding words in the sentence simultaneously, which is especially important for detecting subtle cues related to mental health.

The core component of the system is the classification head attached to the final hidden state output of the BERT model, specifically the [CLS] token representation. This output is passed through a fully connected dense layer with ReLU activation, followed by a dropout layer to prevent overfitting. Finally, a softmax activation function is applied to generate probabilities over the target classes. The model is trained using the cross-entropy loss function, which is well-suited for multi-class classification tasks, and the AdamW optimizer is employed for efficient and stable convergence.

During training, the dataset is split into training, validation, and test sets to evaluate the model's generalization capability. Various performance metrics such as accuracy, precision, recall, F1-score, and AUC (Area Under the ROC Curve) are used to assess the effectiveness of the model. A confusion matrix is also generated to visualize the model's classification performance across different mental health categories, helping identify classes that are frequently misclassified. To improve model robustness and reduce overfitting, techniques such as dropout regularization, learning rate scheduling, and early stopping are applied. Data augmentation strategies like back-translation and synonym replacement may also be used to introduce linguistic variability.

Once the model is trained and validated, it is deployed for inference in real-time or batch processing scenarios. In a live setup, user inputs such as social media posts or chat messages can be streamed into the system, where the model processes the text and classifies it based on the likelihood of containing depressive or suicidal content. The results can be used to trigger alerts, flag concerning content for human review, or provide users with appropriate mental health resources. Additionally, a logging mechanism can be incorporated to track predictions over time, enabling the detection of long-term trends or recurring patterns in an individual's language.

The proposed algorithm thus leverages the strengths of transformer models like BERT in understanding contextual and emotional nuances in text. By fine-tuning on domain-specific datasets and integrating best practices in data preprocessing and model optimization, the system aims to provide an accurate, reliable, and ethically sound solution for detecting mental health concerns through natural language. This approach not only builds upon existing advancements in NLP and deep learning but also contributes to the larger goal of supporting mental health awareness and intervention using AI technologies.

The trained BERT-based model is exported in a lightweight and portable format (such as a .pt file) and can be seamlessly integrated into web applications using frameworks like Flask or deployed on mobile devices using TensorFlow Lite or ONNX Runtime. This makes the solution adaptable for real-time inference in various environments, including chatbots, virtual assistants, and mental health monitoring apps. Special emphasis is placed on ethical considerations, such as safeguarding user privacy, securing sensitive data, and ensuring informed consent, especially when analyzing personal or emotionally vulnerable text content.




Fig 2 Meta Analysis

In conclusion, the proposed NLP algorithm offers a powerful and accurate method for detecting signs of depression and suicidal ideation in textual data. By leveraging pre-trained transformer models like BERT, the system can understand nuanced language patterns and context, which are crucial for identifying mental health indicators. The architecture is optimized for both performance and generalization, ensuring high reliability across diverse inputs. With its strong potential for real-world deployment, this system contributes meaningfully to the development of emotionally intelligent and socially responsible AI tools, particularly in the fields of digital mental health, content moderation, and public safety.

BERT consistently outperformed traditional deep learning approaches, particularly in complex, context-rich examples. LSTM provided a competitive baseline with faster training time and lower computational costs. Misclassifications mainly occurred between semantically close emotions such as fear and sadness.

IV. SIMULATION RESULTS

The graph plots training accuracy against epochs, showing an increasing trend. The accuracy starts at around 84% at epoch 1 and rises to approximately 95% by epoch 5. The graph is labeled as "Figure 5.2.2 Training Accuracy Over Epochs." This suggests that the data represents the performance of a natural language processing (NLP) model as it learns over multiple training iterations (epochs).

NLP models are often used to understand and process human language, and the training accuracy indicates how well the model is learning to perform its task, such as sentiment analysis or emotion recognition, as suggested by the

document title. The increasing accuracy over epochs is a typical pattern in machine learning, indicating that the model is improving with more training.




Fig 3 Training Accuracy Over Epochs

. The graph plots accuracy (percentage of correct predictions) against epochs (number of passes through the entire training dataset).

Specifically, the graph illustrates how the training accuracy of an NLP model, such as BERT or LSTM, improves over time as it learns from the data. The x-axis represents the epochs, ranging from 1.0 to 5.0, and the y-axis represents the accuracy, ranging from 84% to 96%. The plot shows a generally increasing trend in accuracy as the number of epochs increases, indicating that the model is learning effectively. The accuracy rises sharply between epochs 1 and 2, then plateaus, suggesting that the model's learning rate slows down after the initial epochs.

The accompanying text defines key terms: an epoch is one complete pass through the training dataset, and accuracy is the percentage of correct emotion predictions. Plotting these values helps track the model's performance and identify potential issues such as overfitting, where the model performs well on the training data but poorly on new, unseen data. The graph visually represents the model's learning curve, which is a crucial tool for understanding and improving model performance.




Fig 4 Emotion Distribution in Dataset

The image shows a bar chart titled "Emotion Distribution in Dataset". The chart displays the count of six emotions: sadness, anger, love, surprise, fear, and joy.

Sadness: Approximately 4700

Anger: Approximately 2100

Love: Approximately 1200

Surprise: Approximately 500

Fear: Approximately 1900

Joy: Approximately 5300

The chart indicates that joy is the most prevalent emotion in the dataset, followed by sadness. Surprise is the least prevalent emotion. Datasets like this are used in sentiment analysis and emotion recognition tasks, often involving machine learning models. These models are trained to recognize and classify emotions in text, speech, or images. The data can be used to analyze public opinion, detect facial expressions, or understand user behavior.

This data suggests that the dataset is heavily skewed towards sadness and joy, with significantly fewer instances of anger, love, and fear, and almost no instances of surprise. This type of analysis is common in Natural Language Processing (NLP) to understand the emotional tone and content of textual data. The chart is labeled as "Figure 5.2.1 Emotion Distribution Dataset," indicating it is part of a larger study or report.

When using a **smaller dataset** (such as a limited set of tweets for training and testing), the accuracy of both training and testing data was **relatively high**, often exceeding **90%**. However, when scaled to a larger dataset like **FER-2013** for emotion detection, the accuracy often showed more fluctuations due to the greater variety of text styles, language use, and contextual complexity. This reflects the fact that, while smaller datasets lead to quicker convergence and high accuracy, larger datasets, such as the **Emotions Distribution in Dataset**, present more challenges, with the accuracy varying due to the increased diversity of text data

V. CONCLUSION

Emotional understanding through Natural Language Processing (NLP) represents a transformative advancement in human-computer interaction, bridging the gap between linguistic analysis and affective computing. By leveraging machine learning, deep learning, and psychological theories, NLP systems can now detect, interpret, and respond to human emotions with increasing sophistication. However, this progress comes with significant challenges and responsibilities.

In conclusion, Natural Language Processing (NLP) offers promising avenues for understanding emotions expressed in text, with applications spanning various domains like customer support, social media analysis, and even mental health support. However, challenges remain in areas such as context interpretation, cultural nuances, and the complexity of human emotions themselves.

VI. FUTURE WORK

To further improve the effectiveness, robustness, and societal relevance of the system, the following directions are proposed:

- 1. Data Collection:** Emotional AI systems gather data from various sources, including text, voice, facial expressions, and physiological signals. This data is essential for training the AI to recognize and interpret emotions.
- 2. Emotion Detection:** The AI uses machine learning models to analyze the collected data and detect emotional states. This process involves identifying patterns corresponding to emotions such as happiness, sadness, anger, or surprise.
- 3. Emotion Interpretation:** The AI interprets emotions in the interaction context once they are detected. This interpretation helps the AI understand the user's emotional state and the potential reasons behind it.
- 4. Response Generation:** The AI generates appropriate responses based on the interpreted emotions. These verbal, visual, or behavioral responses are designed to align with the user's emotional state and improve the interaction.
- 5. Enhancing Patient Care:** Emotional AI revolutionizes healthcare by providing deeper insights into patient emotions and well-being. AI systems equipped with emotional intelligence can analyze facial expressions, voice tones, and physiological signals to detect stress, anxiety, or depression in patients. This helps healthcare providers offer more personalized care and support.
- 6. Mental Health Monitoring:** Emotion AI apps and devices can monitor patients' emotional states continuously. For instance, wearable devices with sensors can track physiological indicators of stress or anxiety, alerting healthcare providers to intervene when necessary. This proactive approach can significantly improve mental health outcomes.

REFERENCES

- [1].Peixiang Zhong, Di Wang, and Chunyan Miao.2019. Knowledge-enriched transformer for emotion detection in textual conversations. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 165–176, Hong Kong, China. Association for Computational Linguistics.
- [2].Michael J. Tanana1&Christina S. Soma2&Patty B. Kuo2&Nicolas M. Bertagnolli3&Aaron Dembe2&Brian T. Pace4&Vivek Srikumar5&David C. Atkins6&Zac E. Imel2 Accepted: 18 December 2020# The Psychonomic Society, Inc. 2021/ Published online: 22 March 2021.
- [3].Chang, Y.-C., Chu, C.-H., Chen, C.C., Hsu, W.-L.: Linguistic template extraction for recognizing reader-emotion. In: International Journal of Computational Linguistics & Chinese Language Processing, Volume 21, Number 1, June 2016(2016).<https://aclanthology.org/O16-200>
- [4]. [8:38 pm, 20/03/2025] .: and Future Directions [8:56 pm, 20/03/2025] .: Sheetal Kusall · Shruti Patil2 · Jyoti Choudrie3 · Ketan Kotecha2 · Deepali Vora1 · Ilias Pappas4 © The Author(s) 2022
- [5]. To cite this article: Ashritha R Murthy and K M Anil Kumar 2021 IOP Conf. Ser.: Mater. Sci. Eng. 1110012011
- [6]. To cite this article: Ashritha R Murthy and K M Anil Kumar 2021 IOP Conf. Ser.: Mater. Sci. Eng. 1110012009

- [7].Abdul-Mageed, M., & Ungar, L. (2017). EmoNet: Fine-grained emotion detection with gated recurrent neural networks. Proceedings of the ACL Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, 2017, 718–728.
- [8].Buechel, S., & Hahn, U. (2018). Emotion representation mapping for automatic lexicon construction (almost) without a seed. Proceedings of the 27th International Conference on Computational Linguistics (COLING 2018), 911–922.
- [9].Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., & Potts, C. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP), 1631–1642.
- [10].Strapparava, C., & Mihalcea, R. (2007). SemEval-2007 Task 14: Affective text. Proceedings of the 4th International Workshop on Semantic Evaluations (SemEval-2007), 70–74.
- [11].Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008.
- [12].Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., ... & Rush, A. M. (2020). Transformers: State-of-the-art natural language processing. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, 38–45.
- [13]. Anumolu, V. R., & Marella, B. C. C. (2025). Maximizing ROI: The Intersection of Productivity, Generative AI, and Social Equity. In Advancing Social Equity Through Accessible Green Innovation (pp. 373-386). IGI Global Scientific Publishing.
- [14]. Demszky, D., Movshovitz-Attias, D., Ko, J., Cowen, A., Nemade, G., & Ravi, S. (2020). GoEmotions: A dataset of fine-grained emotions. arXiv preprint arXiv:2005.00547.
- [15]. Mohammad, S. M., & Turney, P. D. (2013). Crowdsourcing a word–emotion association lexicon. Computational Intelligence, 29(3), 436–465



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details