



Leveraging NLP and Machine Learning for Fine-Grained Emotion Detection from Textual Data

OPEN ACCESS

Volume: 13

Special Issue: 2

Month: January

Year: 2026

E-ISSN: 2582-0397

P-ISSN: 2321-788X

Citation:

Yadav, Anita Kamlesh, and Pushpa Susant Mahapatro. "Leveraging NLP and Machine Learning for Fine-Grained Emotion Detection from Textual Data." *Shanlax International Journal of Arts, Science and Humanities*, vol. 13, no. 2, 2026, pp. 101–08.

DOI:

<https://doi.org/10.34293/sijash.v13iS2-i4-Jan.10590>

Anita Kamlesh Yadav

Assistant Professor

Rajiv Gandhi College of Arts Commerce and Science, Vashi Navi Mumbai, Maharashtra, India

Dr. Pushpa Susant Mahapatro

Associate Professor

Vidyalankar School of Information Technology Wadala Mumbai, Maharashtra, India

Abstract

New research into emotion detection from text is a quickly growing field of Natural Language Processing (NLP) that has many different applications, including mental health, customer service and social media analytics. This study looks at both traditional machine learning models and newer, more advanced deep learning models to help identify emotions, such as happiness, sadness, anger, fear and surprise. Baseline models (Logistic Regression, Naive Bayes, and Support Vector Machines (SVM)) are compared with three different deep learning architectures (recurrent neural networks (RNNs), long short-term memory (LSTM) networks, and the transformer BERT model). Labeled datasets (e.g., GoEmotions, CrowdFlower) are used to evaluate the accuracy, precision, recall and F1 score of the models. The findings indicate that transformer-based deep learning models (e.g., BERT) outperform baseline models and will likely play a crucial role in developing emotionally intelligent artificial intelligence systems.

Keywords: Emotion Detection, Natural Language Processing (Nlp), Logistic Regression, Naive Bayes, And Support Vector Machines (SVM)

Introduction

Through rapid digitization, the coupling of text as a communication method and connection to one another has resulted in a wealth of human emotion stored within text itself. Emotion detection within Natural Language Processing (NLP), once viewed as only a computational concern, is now a fundamental area of growth for modern Human Computer Interaction (HCI). Previously, text processing focused mainly on the syntactic and literal meanings of the text. The current state of artificial intelligence (AI) seeks to uncover the psychological state of the user—feeling fear, exhilaration, frustration, etc.—within the unstructured data of the text. This study will focus on the use of advanced NLP models including the evolution of machine learning from static machine learning systems to dynamic systems such as Deep Learning models and BERT (Bidirectional Encoder Representations from Transformers). We will combine labelled emotional datasets with contextual embeddings so that this work can move from very

limited binary sentiment analysis to a more nuanced understanding of human feeling thereby enabling AI systems to possess “digital empathy” towards their users [1,2].

Historically, emotional content within language has only been studied computationally in terms of Sentiment Analysis — a broad-brush way of categorizing language by assigning it to one of three categories: Positive, Negative or Neutral. Although Sentiment Analysis can provide a high-level overview of market trends, its use of polarity is a highly reductive way of analyzing emotions. It does not allow for discrimination between the different types of negative emotions (anger vs. sadness). The rise of web-based communities on social platforms such as Reddit and X (formerly known as Twitter) and encrypted messaging has created an urgent need for a more refined approach to identifying emotional content (aspects) of language [3, 4]. Therefore, we need to create an additional tool or a “chromatic scale” for identifying emotional content. The primary challenge is the inherent ambiguity of human language since many factors can affect the way we interpret what someone says, including their cultural background, use of sarcasm, and use of metaphors that may not be obvious from context alone. For example, the same phrase, “that’s just great,” may express happiness or disgust depending on the context of the preceding statement. To address this disparity, the field has evolved over three distinct technological epochs [5].

Statistics-based Machine Learning epoch: Early machine learning methods used features and possibly combinations of multiple features defined through handcrafted and developed dictionary entries, being innovative at that time because they interpreted the tokens independently or separate from each other (aka, Bag-of-Words) [6].

The emergence of Deep Learning: The invention of Recurrent Neural Networks (RNN) and Long-Short Memory Networks (LSTM) transformed the approach taken with processing language by allowing the network to utilize prior input to inform its current input. This approach provides the network with the ability to see how different statements flow together within the same sentence emotionally. However, the LSTM networks were not very adept at remembering through multiple layers of abstraction used to create a time sequence but rather only remember the previous layer of abstraction [7].

The Transformational Approach to Technology: Currently almost all systems are based on transformer architectures such as BERT, which use self-attention mechanisms to correlate words within a sentence regardless of how far apart those words are; and it does this in parallel by looking at relationships between all words in a sentence simultaneously [8]. Secondly, BERT has been trained on large training sets to develop a pre-trained model of how language works prior to utilizing it for the purpose of identifying certain “emotional” characteristics (e.g. Joy, Sadness, Anger, Fear, Surprise, Disgust).

Finally, the Transformational Approach to Technology provides significant ethical and clinical implications. For example, we can use NLP technologies as diagnostic tools to discover subtle language patterns in patients’ written diaries indicating the onset of depression or general anxiety long before physical symptoms appear in individuals who require mental health treatment (e.g. psychologist’s office visits) [9,10]. By leveraging “emotionally aware” algorithms and integrating them into our digital infrastructure, we can move closer to creating an AI that is designed around human experience—not just as a data processor but rather as someone who truly understands the human condition beyond the computer screen.

Literature Review

There’s been a change from simply looking at the words people use (what they say) to understanding how they feel about what they are saying (the way they express their feelings) in Natural Language Processing. Traditional sentiment analysis used to work on either a binary (yes/no) or trinary (positive/neutral/negative) model for analysing people’s emotions; however, in this project, we will use an approach to give ‘affective granularity’ of individually identifiable emotional states (i.e., euphoric, indignant, anxious, grieving). With an increase in communication occurring through digital text messaging, being able to identify these more subtle emotions and how to interpret them from textual format is no longer an optional ‘nice to have’; it has

become a ‘need to have’ technology. This research moves beyond the crude forms of polarity analysis toward creating an emotionally multi-dimensional approach which provides for a new generation of systems capable of engaging users with a genuine emotional presence while still being aware of the context in which they are providing input [11].

In this study, we have a comprehensive process that starts with the intentional collection of trustworthy, labelled datasets from public platforms (such as Hugging Face) and many other different types of social media sources. The datasets will be used to compare simpler models like Naïve Bayes with more advanced architectures that use attention-based techniques such as the Transformer/BERT family of models. Once the models have been built and compared based on overall accuracy, the evaluation will be based on additional metrics such as precision, recall, and F1 to capture how well a model performs regarding the inherent imbalances in emotion in real-world datasets. At the end of this technical development, a working prototype will be created to translate what has been learned about the performance of each algorithm into a real-world application that can provide real-time emotional assessments.

There is reason to believe that the combination of powerful cloud-computing systems and pre-trained language models will lower the technical difficulty of executing large-scale Deep Learning projects. This development will have an especially positive impact on businesses and organizations that must consider human emotions as an important consideration in the way they provide products and services or conduct business with customers (in fields including but not limited to Health Care). In the Health Care Industry, this combination will facilitate a non-invasive diagnostic capability, which enables health care providers to monitor the health of their patients. In many commercial settings, organizations will now have the opportunity to become more proactive about how they communicate with their customers by providing them with timely and relevant information. This project highlights how the use of “transfer learning” to develop AI systems that understand human emotion can, in fact, be accomplished today and will also be a critical component in advancing the development of AI systems for the future use in areas such as education, mental health advocacy, and more responsive governance on behalf of the public.

Objectives of the Study

The aim of this project is to create a comprehensive emotion detection system, which goes beyond traditional polarity-based methods. The objective is to establish a system that will enable users to have a complete understanding of the emotional state of their audience. The focus will be on developing superior machine learning algorithms. The study will consist of a comparison between traditional machine-learning algorithms (i.e., logistic regression and naive Bayes’ method), and newer, high-parameter, deep-learning models (i.e., BERT). In addition to developing the algorithms, the study will also include empirical testing and validation of the developed algorithms using multiple metrics such as F1 scores and precision. Finally, these objectives will serve as a roadmap for the development of future human-centered AI applications that incorporate genuine digital empathy.

Methodology

The computational foundation of this project is based on two complementary strategies - classical statistical machine learning techniques and the high capacity of neural architectures - such that the combination is more powerful than either would be alone. To establish an objective benchmark for subsequent performance improvements, the project will first employ traditional machine learning methods that are representative of state-of-the-art approaches to the analysis of text data (such as Multinomial Naive Bayes or Support Vector Machine). These statistical models will utilize handcrafted feature sets (i.e. n-gram distributions, TF-IDF) for their input to produce classifications based on the text. Despite producing classifications quickly, they typically do not provide the accuracy needed to interpret the intricate details of an individual’s linguistic context.

The project uses transformer models, such as BERT, to address some of the problems that have been experienced when training models to recognize emotion. In contrast to LSTMs, which utilize sequential processing of input, the BERT architecture allows simultaneously comparing relationships among tokens or words on all sides of one another. In particular, it capitalizes upon multi-head self-attention (MHA), a technology that allows for capturing subtle differences in emotion based on the position of words within a sentence. Thus, BERT enables greater accuracy in recognizing complex emotions such as sarcasm and fear than do other types of model architectures. Additionally, BERT enhances the previous models' effectiveness by having been pre-trained on large datasets representing different emotions (e.g, happy, sad, etc.), moreso than previous models, which have relied upon supervised learning techniques.

Data Collection and Description

The foundation of this research is based on high-quality, publicly available emotional corpora, which provide an essential source of supervisory signals for the fine-grained multi-class classification task. These datasets have been created with a careful compilation of hundreds of different types of written linguistic snippets, from short messages in various types of social media to complex written documents. Each of the data pieces has been assigned its own specific label based on the emotion captured, such as joy, anger, sadness, or fear. The training pipeline for the model leverages the diversity of these labeled pairs to produce a high level of lexical variability; this high level of lexical variability is critical for training Transformer-based models like BERT by allowing them to go beyond simple word-matching performance to comprehend and recognize contextual depth within the data. For example, the samples provided below depict the unique and existing variability of human expression. The statistical density of these cases allows the model to create meaningful and relevant representation of unique human expression patterns and to generalize effectively to other types of language data in different real-world contexts.

Table 1 Representative Samples of Annotated Emotional Data

Textual Input	Target Emotion Label
"I am absolutely thrilled about the prospects of my new career path!"	Joy
"This is undeniably the most devastating day I have ever experienced."	Sadness
"I fail to understand why you would commit such an inconsiderate act."	Anger
"The uncertainty of the upcoming results is causing significant distress."	Fear

Preprocessing Techniques

The steps below outline a Text Normalization and Preprocessing Pipeline. The goal of this multi-stage normalization process is to prepare raw, unstructured text data for statistical and neural classifiers by minimizing linguistic "noise" that may interfere with detecting emotional signals. This enables models to focus on semantic intent rather than typographic variability.

The First Step: Discrete Semantic Tokenization - The tokenization process is the first step in converting large continuous character streams into smaller, more easily digestible semantic units. Tokenization in traditional models utilizes word-level tokenization. In contrast to this approach, our project employs the WordPiece or Byte-Pair Encoding (BPE) types of tokenization often used in Transformer architectures. Emotion detection is a case where working with subword units for OOV words improves the capacity of the model to derive emotional context for morphologically complex and slang words that may otherwise be ignored.

Lexical Standardization (Lowercasing) - To alleviate the problem of feature sparsity, we utilized case folding on the corpus. Case folding means transforming text to a common format in the form of lowercase. Thus, the model sees "Angry," "ANGRY," and "angry" as one feature. The standardization of text in this way is especially important in traditional machine learning models to avoid increased dimensionality in the

feature space. However, in some deep learning implementations, models with case preferences will be kept if capitalization is shown to be a strong predictor of emotional intensity, such as in examples where one is shouting or expressing anger.

Selective Pruning and Stopword Heuristics - Stopword removal refers to the elimination of high-frequency tokens (e.g., “the,” “is,” or “at”) that typically carry no affective meaning. The process of stopword pruning with Logistic Regression or Naive Bayes is necessary for reducing dimensionality and improving computational throughput. Conversely, with BERT, this step is taken with caution or avoided. As Transformers depend on knowing where each word is located in relation to all other words, removing even small function words can interfere with the attention mechanism, diminishing the model’s ability to interpret the subtle distinctions of emotional expression.

Project Synthesis: Problem and Proposed Framework

The existing research addresses a critical aspect of the “semantic gap” of previous affective computing research. With traditional sentiment analysis methods focusing only on three emotions (positive, negative, and neutral); they miss the many nuances of human emotion as shown in complex sensations or feelings. The primary limitation of most traditional sentiment analysis research is that they oversimplify the linguistic nuances. A “negative” sentiment will not provide distinctions or differentiations between feelings of discomfort (i.e., clinical) or dissatisfaction (i.e., frustration) that a customer may experience.

The gap created by oversimplified sentiment analysis is being addressed through a hybridized computational architecture supporting multi-class, rather than binary, sentiment classifications. Therefore, the proposed pipeline has been organized to provide a comprehensive comparison of traditional statistical models, such as Support Vector Machines (SVM) and Naive Bayes Classifiers, with state-of-the-art Transformer architectures (BERT) that utilize bidirectional and high dimensional representations. The new pipeline creates an efficient way to decode subtle emotional cues. Additionally, this dual approach increases the classification performance by providing a better means of implementing empathy in real-world applications, such as predictive mental health monitoring and high-sensitivity consumer relationship management.

Data Integration and Structural Synthesis

The empirical validity of this research is composed of two established emotional datasets, which are the GoEmotions and Crowdfower emotion datasets. Providing a densely concentrated dataset of approximately 58,000 Reddit annotations with a finely grained 27-category taxonomy enables the lexical diversity needed for training very large Transformer models. The Crowdfower dataset serves as a point of cross-validation to assess the general emotional performance of model performance based on less contextually specific emotion labels. Thus, the two datasets combined provide a complete overview of digital emotion (affect) from informal social media-based interactions to formal structuring of emotions.

In preparing the raw corpora for product use, a highly disciplined and systematic data engineering process was followed. The primary step in this engineered process was Tokenization, where the raw text was decomposed into smaller semantic subunits using WordPiece in order to facilitate compatibility with Transformer. Second, non-contributing linguistic information was removed by selectively applying the process of Stopword Removal. In order to align with deep learning architecture requirements of fixed input sizes, Padding and Truncation techniques were applied to achieve uniform sequence lengths across all batches in the training set. Normalizing sequence lengths through structural standardization contributes to enhanced computational efficiency, which allows for the resulting feature vectors to accurately represent the emotional intention of the text, thus eliminating the negative influence of typographical noise and discrepancies in the length of the input sequence.

Results and Discussion

The results of our experimental study exhibit a measurable hierarchy between statistical methods of natural language processing and more modern neural approaches to natural language processing. Traditional machine learning approaches have provided a functional baseline for certain types of emotional categories, such as joy and anger; however, they exhibit the limitation of being less sensitive to the lexical shifts present in emotion-laden language. The benefits of the BERT architecture arise because of the ability of BERT to provide a significant increase in classification accuracy, which can be attributed to its bi-directional self-attention architecture. That is, BERT uses its bidirectional structure to form a better understanding of the meanings of words in the context of their surrounding words. Thus, BERT is capable of performing “contextual disambiguation,” or understanding how a positive token such as “happy” can have its psychological meaning inverted through the use of distal negation (e.g. “not wholly happy”). The transition from local word-analysis to global sentence-context has contributed to an increase in the F1-score for BERT, thus confirming the need for Transformer models to capture and describe fine distinctions between emotions.

Three principal technical and conceptual barriers were made clear through the application of an emotion-aware system:

1. **Subjectivity and Linguistic Ambiguity** - Human emotion is constantly changing, which creates a major hurdle for emotion classification systems. Since each individual experiences emotions differently, as well as the contextual effects on emotion classification, there will always be inter-annotator differences on what the true positive labels of texts are.
2. **Dataset Imbalance and Representation** - An additional technical challenge was created because there are far fewer instances of some emotions classified in emotion-annotated corpora than others (i.e., a high imbalance among the positive, neutral, and negative classifications). Large numbers of emotion-instance/classification imbalances ultimately result in classifiers having a bias toward a few labels (or emotion classes); thus, they require additional advanced techniques (e.g., SMOTE and/or class-weighted loss functions) to correct.
3. **Hardware Limitations and Scalability** - Training of state-of-the-art (SOTA) models (e.g., BERT) creates an exorbitant increase in computations required. The Memory required to store the parameters of these large models and the Memory needed while fine-tuning a large number of parameters necessitate extremely powerful GPU environments if the fine-tuning is going to occur in a timely manner. Therefore, the trade-off between the amount of complexity and detail an L3 (low-resource, low-complexity) emotion classification algorithm can produce (based on the emotions that can be classified) will need to be reconciled in order for emotion-aware systems to function on low-end devices/low-resource environments.

```
# Example usage: Predict emotion for a single input text
sample_text = "I am Thrilled about my new job!"
result = predict_emotion(sample_text)
print(f"Input text: {sample_text}")
for res in result:
    print(f"Predicted emotion: {res['label']}, Confidence: {res['score']:.4f}")

# Evaluate the model on the test dataset
def evaluate_model():
    test_texts = dataset['test']['text']
    test_labels = dataset['test']['label']

    # Make predictions
    predicted_emotions = [emotion_classifier(text)[0]['label'] for text in test_texts]
    actual_emotions = [dataset['train'].features['label'].int2str(label) for label in test_labels]

    # Calculate accuracy, precision, recall, and F1-score
    accuracy = accuracy_score(actual_emotions, predicted_emotions)
    precision, recall, f1, _ = precision_recall_fscore_support(actual_emotions, predicted_emotions, average='weighted')

    print("\nModel Evaluation Metrics:")
    print(f"Accuracy: {accuracy:.4f}")
```

Figure 1 [Figure caption not provided in source document.]

```

print(f"Precision: {precision:.4f}")
print(f"Recall: {recall:.4f}")
print(f"F1-Score: {f1:.4f}")

# Generate confusion matrix
conf_matrix = confusion_matrix(actual_emotions,
                               predicted_emotions)
disp = ConfusionMatrixDisplay(confusion_matrix=conf_matrix,
                              display_labels=dataset['train'].features['label'].names)
disp.plot(cmap=plt.cm.Blues)
plt.title("Confusion Matrix for Emotion Detection")
plt.show()

# Call the evaluation function
evaluate_model()

```

Figure 2 [Figure caption not provided in source document.]

Conclusion and Future Scope

Through a detailed investigation into the structure of this research effort, it is clear that the research team has implemented a well-defined multi-level architecture for an ML-based emotion recognition system. More importantly, this research has highlighted the journey that emotion recognition has taken through their use of a comparative analysis between two different ML approaches: traditional statistical learners and an extremely high-parameter, transformer-based approach (i.e. BERT). BERT is demonstrated to be the best option for processing human emotions, as it can decode the high-level semantic density of emotions and is able to resolve the aforementioned structural ambiguities that are created when using language. The researcher concludes that, by using the Go Emotion dataset to fine-tune a pre-trained architecture like BERT (using the most accurate emotional data available), the researcher can create an AI-based system that can process emotional information and also be able to interpret what exactly the emotional responses of the user mean.

The benefits that this system provides to the end user extend well beyond its technical capabilities, as there are practical applications that represent a fundamental shift in HCI (human-computer interaction) for the better. By enabling companies to incorporate these types of models into their customer service systems and mental health monitoring applications and real-time social analytics, companies can begin to create a more empathetic digital infrastructure. However, there are still challenges that need to be overcome (e.g., emotional labels' subjective nature and the high demand for computation of SOTA models). The results provided from this project present an opportunity for the next generation of "emotion-aware" AI to build on as a scalable foundation. Future studies should focus on optimizing these types of models for low-resource environments and evaluating multimodal input types for machine learning models (e.g., text with audio or visual input) to improve how well machines understand what it means to be human.

References

1. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. <https://doi.org/10.1145/3442188.3445922>
2. Demszky, D., Movshovitz-Attias, D., Ko, J., Cowen, A., Rashkin, H., & Leviathan, N. (2020). GoEmotions: A dataset of fine-grained emotions. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4040–4054. <https://doi.org/10.18653/v1/2020.acl-main.372>
3. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, 4171–4186. <https://arxiv.org/abs/1810.04805>
4. Hutto, C., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of

- social media text. Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media, 216–225.
5. Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis Lectures on Human Language Technologies*, 5(1), 1–167. <https://doi.org/10.2200/S00416ED1V01Y201204HLT016>
 6. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint. <https://arxiv.org/abs/1907.11692>
 7. Mohammad, S. M., & Turney, P. D. (2013). Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3), 436–465. <https://doi.org/10.1111/j.1467-8640.2012.00460.x>
 8. Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval*, 2(1–2), 1–135. <https://doi.org/10.1561/1500000011>
 9. Poria, S., Majumder, N., Mihalcea, R., & Hazarika, D. (2019). Emotion recognition in conversation: Research challenges, datasets, and recent advances. *IEEE Access*, 7, 100943–100953. <https://doi.org/10.1109/ACCESS.2019.2929050>
 10. Saravia, E., Liu, H. C. T., Huang, Y. H., Wu, J., & Chen, Y. S. (2018). CARER: Contextualized affect representations for emotion recognition. Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, 3687–3697. <https://doi.org/10.18653/v1/D18-1404>
 11. 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, 1631–1642.
 12. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008. <https://arxiv.org/abs/1706.03762>
 13. Zhang, L., Wang, S., & Liu, B. (2018). Deep learning for sentiment analysis: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4), e1253. <https://doi.org/10.1002/widm.1253>
 14. Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Edunov, S., Stoyanov, V., & Zettlemoyer, L. (2020). Unsupervised cross-lingual representation learning at scale. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 8440–8451. <https://doi.org/10.18653/v1/2020.acl-main.747>
 15. Plutchik, R. (2001). The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *American Scientist*, 89(4), 344–350. <https://www.americanscientist.org/article/the-nature-of-emotions>
 16. Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. OpenAI. <https://openai.com/research/language-unsupervised>
 17. Sun, C., Huang, L., & Qiu, X. (2019). Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentences. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics, 380–385. <https://doi.org/10.18653/v1/N19-1035>
 18. Yadav, R. K., Jha, S., Sah, A., & Maurya, A. K. (2021). Deep learning-based emotion detection from text: A review. 2021 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS), 546–551. <https://doi.org/10.1109/ICCCIS51004.2021.9397121>