

# An empirical study on how emotion affects the probability of replies based BERT

Zhijie Gan<sup>1</sup>, Zicheng Zhu<sup>1,2</sup>

<sup>1</sup>Computational Mathematics and Cybernetics, Shenzhen MSU-BIT University, Shenzhen, China

<sup>2</sup>1120200072@smbu.edu.cn

**Abstract.** The rapid development of the internet, social media, and online forums have become crucial platforms for people to express their views and emotions. Comments are not only a way for users to express their opinions but also play a vital role in promoting discussions and interactions between users, significantly influencing public opinion. This paper aims to explore the impact of emotions on the likelihood of comments receiving replies, deepening the understanding of the role of emotional factors in interactions on social media and online forums. Through large-scale model training and pipeline parallel computation, this paper employs the Bidirectional Encoder Representations from Transformers (BERT) model for learning and prediction, enhancing accuracy and efficiency. The experimental results show that the response rate of negative emotional comments is about 27%, while the response rate of positive emotional comments is about 18%. It means that the comments with negative emotions are more likely to receive replies than those with positive emotions.

**Keywords:** Large-Scale Model Training, Pipeline Parallel, BERT Model, Sentiment Recognition.

## 1. Introduction

Social media and online forums have become indispensable to people's daily lives in the digital age. Individuals can share their views, post comments, and interact with others on these platforms, facilitating broader and faster dissemination of information. Comments are not only a way for users to express their opinions but also play a vital role in promoting discussions and interactions between users, significantly influencing public opinion. Past research has begun to focus on the role of emotional factors in social media and online forum interactions. The expression of emotions may influence users' attention and response, but the extent of their influence is not yet clear. In past research, text sentiment analysis has garnered widespread attention. Researchers have tried identifying and comprehending emotional information in text using various methods, including rule-based approaches, traditional machine learning techniques, and deep learning methods. Among these, the engineering of sentiment lexicons and features for emotion expression has been a focal point, with researchers constructing sentiment dictionaries and developing emotion-specific features to capture sentiment polarity in text.

For instance, Bo Pang et al. introduced an approach for subjectivity summarization using the minimum cut method to perform sentiment analysis. This method pioneered a novel approach to sentiment analysis [1]. Additionally, traditional machine learning algorithms such as Support Vector

Machines (SVMs) and Naïve Bayes classifiers have found extensive application in sentiment analysis tasks. These methods typically involve training and classification using manually extracted features, but their performance may be limited when dealing with complex natural language data. For example, Andrew L. et al. presented a word embedding learning approach for sentiment analysis, acquiring representations of vocabulary from extensive textual data through unsupervised learning [2].

The development of deep learning technology in recent years has changed the sentiment analysis scene. Deep neural networks have been widely used for natural language processing applications, including sentiment analysis, particularly Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs). Through end-to-end learning, these models can automatically extract valuable information from text, leading to significant success in sentiment analysis tasks. For instance, Ashish Vaswani et al. introduced the Transformer model, a representative deep learning model in text sentiment analysis. It has achieved outstanding performance in machine translation and sentiment analysis tasks. The Transformer model has demonstrated remarkable capabilities and significant advancements, particularly in tasks related to natural language understanding and emotional sentiment analysis [3].

However, despite making some progress, text sentiment analysis still faces challenges. The diversity of emotions, the complexity of context, and the ambiguity in texts continue to make sentiment analysis an active research area. Additionally, the impact of emotional sentiment in texts on social movements requires extensive case studies. For instance, Johan Bollen et al. examined how emotional information on Twitter is correlated with stock market fluctuations, revealing the influence of social media sentiment on financial markets [4]. Online comments have always been a research hotspot, as exemplified by Andrei Oghina et al., who investigated sentiment analysis across multiple languages and proposed a feature selection method for sentiment classification from online forums [5].

However, this paper notices that specific comments have a more pronounced impact among numerous replies. Some comments with many replies even attract more attention than the videos themselves. Therefore, this study aims to explore the reasons behind this phenomenon. This research seeks to expand upon previous scholarly accomplishments and incorporate the most recent advancements in deep learning methodologies to investigate the influence of emotional perceptions on responses to comments. This will offer more insight into the significance of affective elements in social media platforms and digital discussion boards. To improve the accuracy of predicting comment responses, this study leverages the power of the Bidirectional Encoder Representations from Transformers (BERT) model. As introduced in the seminal work by Devlin et al., BERT has demonstrated remarkable capabilities in various natural language understanding tasks [6]. This paper harnesses BERT through large-scale model training and pipelined parallel computing to expedite the model learning process and enhance prediction accuracy.

The observations from this paper indicate that comments with a negative emotional undertone have a response rate of approximately 27%, in contrast to those with a positive sentiment, which stands at roughly 18%. This underscores the inclination for negative comments to garner more responses. The insights derived from this research are instrumental for strategizing and overseeing comment systems on platforms such as social media and digital discussion boards.

## 2. Methods

### 2.1. *The overview of Model*

To effectively manage and analyze extensive datasets while enhancing prediction accuracy, this research strategically employs the Bidirectional Encoder Representations from Transformer (BERT) model. BERT represents a pioneering advancement in natural language processing, characterized by its ability to capture contextual relationships within text data. The BERT model is grounded in the Transformer architecture, a structural framework celebrated for its impeccable parallelism and capacity to sustain operational efficiency despite vast datasets. This architectural choice aligns with the overarching goal of this study: to navigate large-scale data and provide precise predictions, particularly within the context of comment responses on online platforms. This research capitalizes on the power of multiple Graphics

Processing Units (GPUs) to maximize the utility of available computational resources. The study further optimizes the model's performance by processing data concurrently across these GPUs. This parallel processing strategy enhances both the efficiency and speed of the model's learning process. In summary, the amalgamation of BERT, renowned for its contextual awareness and the efficient parallelism of the Transformer architecture, establishes a robust foundation for this research's predictive model. This model, developed through an intricate combination of advanced machine learning techniques and parallel computing, is a pivotal tool for assessing the likelihood of responses to comments in online environments.

### *2.2. Data Acquisition and pre-processing*

This paper meticulously constructs its foundation by proficiently collecting substantial volumes of real-world internet comments. These harvested comments encapsulate a vast spectrum of emotions, opinions, and subjects, rendering them invaluable information for the ensuing investigation. Within this diverse corpus of comments, the study endeavours to uncover the intricate interplay between emotions and engagement, shedding light on the factors influencing online interactions. After the data collection phase, a meticulous categorization process ensues. Comments are systematically tagged according to their associated sentiment. Those comments attracting responses are denoted as "1," while those left unanswered are labelled "0." This sentiment-based labelling schema furnishes the dataset with a structured foundation and enables the BERT model to make nuanced and precise predictions. The study aggregates a dataset comprising 8,000 comments, underpinning the subsequent analyses. Quality control measures play a pivotal role in safeguarding the integrity of the dataset. In this regard, the research initiates preliminary pre-processing procedures. These preparatory steps encompass removing special characters, punctuation marks, and extraneous spaces. The objective is to distil the raw text content, purging it of superfluous elements that might introduce noise and impede subsequent analyses. Furthermore, the dataset is judiciously partitioned into training and testing subsets in anticipation of model evaluation. This partitioning adheres to a 7:3 ratio, balancing training data volume and the need for a robust evaluation framework. This meticulous approach ensures that the model's performance can be systematically assessed post-training, bolstering the credibility and rigor of the study's findings.

## **3. Experimental results and analysis**

### *3.1. Experimental setting*

This paper employs the AdamW optimizer with a learning rate scheduler, which is pivotal in optimizing our model's performance. This paper caps the maximum number of training rounds to 10 to balance training effectiveness and computational resources. Moreover, this paper introduces a warm-up step and an early stop mechanism to fine-tune the training process further. Recognizing the substantial time investment required for model training, this paper aims to optimize GPU utilization. This paper involves adopting a distributed training strategy, splitting the training process into two parallel processes. This innovation maximizes the GPU's processing power while maintaining training stability. This paper employs various optimization techniques to further enhance training efficiency and model performance. These include dynamic adjustments to learning rates and batch sizes, the introduction of a learning rate scheduler to manage learning rates adaptively, increasing model complexity for better representation of underlying patterns, and implementing an early stopping mechanism to prevent overfitting. Another noteworthy addition to this paper is introducing a weight decay term within the loss function. This term is proportionate to the square of the model weights, serving as a powerful tool to combat overfitting and enhance generalization. The culmination of these optimizations yields a significant reduction in training time, a marked improvement in model fitting performance, and a substantial enhancement in the overall quality of our experimental results. Notably, the experimental setup utilizes two A4000 GPUs to execute the research, demonstrating the feasibility and effectiveness of strategies in a practical GPU configuration.

### 3.2. Analysis

The Long Short-Term Memory (LSTM) model represents a specialized modification of recurrent neural networks (RNNs). This design was conceptualized and developed to confront the pervasive challenges of diminishing and surging gradients, which are frequently faced by traditional RNN architectures when they grapple with extensive sequences [7]. LSTM has intricate internal apparatuses tailored to seize and sustain extended temporal correlations, rendering it particularly adept for tasks within natural language processing, such as linguistic pattern identification and emotional tone discernment. Its architectural design incorporates a trifold system encompassing input regulation mechanisms, memory erasure functions, and information release modules, collectively guiding data progression within the model [8].

In a vein parallel to the LSTM, the Gated Recurrent Unit (GRU) stands as another refined adaptation of the recurrent neural network architecture, with its primary objective being to surmount the inherent obstacles presented by the conventional RNN framework. Regarding structural complexity, the GRU exhibits a more streamlined design than the LSTM, characterized by merely two regulatory components: a data revision mechanism and a memory reinitialization module [9]. This relative structural simplicity translates to computational agility without compromising the capability to latch onto and preserve extended sequential interrelations. The GRU exhibits commendable performance within natural language processing, often demonstrating swifter model training and prediction velocities [10].

This study compared BERT with traditional RNN models, including LSTM and GRU, to explore their performance differences in text classification tasks, as shown in Table 1. Experimental results demonstrate that the BERT model excels in text classification tasks, particularly tasks requiring a global understanding of context, such as sentiment analysis. BERT's advantage is its ability to automatically learn language representations without manual feature engineering, making the model easier to train and transfer across different tasks. Traditional RNN models, despite having some capabilities in modeling sequential data, need help dealing with natural language tasks. LSTM, GRU, and other traditional RNN models need help to capture long-term dependencies, which may limit their performance in handling long texts or complex contexts. Compared to traditional RNN models, BERT stands out with its revolutionary pretraining capabilities and contextual sensitivity. Through extensive pretraining as a language model, BERT acquires a deep understanding of vocabulary, syntax, and semantics, enabling it to capture information and relationships within text better. Additionally, BERT is bidirectional, allowing it to consider context information simultaneously, whereas traditional RNN models process text sequentially.

**Table 1.** Experimental results.

	Accuracy	F1
BERT	0.798	0.76
LSTM	0.672	0.65
GRU	0.472	0.51

To ensure the datasets were as representative as possible, a filtering process was conducted to remove comments with insufficiently expressed emotions, resulting in a curated dataset of 2,000 comments. This filtered dataset was used as the foundation for subsequent analysis and prediction.

Using the BERT model, the study aimed to predict whether a given comment would receive replies or not. The experimental results revealed intriguing patterns. Specifically, it was observed that comments characterized by negative emotions had a reply rate of approximately 27%, significantly higher than the reply rate of comments with positive emotions, which stood at approximately 18%.

These findings suggest that comments expressing negative emotions tend to garner more responses compared to their positively toned counterparts. This insight provides valuable information regarding user engagement dynamics on online platforms and may be instrumental in shaping online community management strategies. Furthermore, it underscores the potential impact of emotional sentiment on user interactions in digital spaces.

#### 4. Conclusion

The revelations from this paper's efforts brought to light a somewhat paradoxical understanding of virtual interchanges. Disputing the conventional wisdom, this paper's observations discerned that commentaries imbued with more somber or negative emotional tonalities appear to magnetize a greater volume of reactions than those radiating more uplifting sentiments. This unexpected discovery alludes to a curious dynamic — melancholic or adverse sentiments seemingly wielding a beguiling magnetism, prompting individuals to participate more actively in conversational exchanges. Fundamentally, somber tones manifest a conspicuous presence within the expansive digital landscape, culminating in an augmented incidence of responsive engagements. However, it is essential to approach these results cautiously, as they are subject to various influencing factors. The nuances of data collection may influence our conclusions' accuracy, our sample distribution, and the inherent complexities involved in training machine learning models like BERT. Despite its robustness, the BERT model is not infallible, and there may be instances where its predictions deviate from reality. For comments laden with negative emotions, platforms could implement strategies to steer discussions towards positivity and constructive solutions, thereby mitigating controversy and inappropriate remarks. Simultaneously, platforms could provide additional attention and support to affirmative comments, fostering an environment where users are motivated to share positive sentiments and opinions. Future research endeavours should expand their horizons to further our understanding of the intricate dynamics at play. This could involve increasing the sample size, enhancing data quality and representativeness, and investigating the influence of emotions on comment responses across diverse online communities and scenarios. Moreover, researchers could explore more complex models and algorithms to refine the accuracy and stability of response rate predictions.

#### Authors Contribution

All the authors contributed equally and their names were listed in alphabetical order.

#### References

- [1] Pang, B., Lee, L., & Vaithyanathan, S. (2002). "A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts." *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 271-278.
- [2] Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., & Potts, C. (2011). "Learning Word Vectors for Sentiment Analysis." *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT)*, 142-150.
- [3] Systems (NeurIPS) Bollen, J., Mao, H., & Zeng, X. (2011). "Twitter Mood Predicts the Stock Market." *Journal of Computational Science*, 2(1), 1-8.
- [4] Oghina, Andrei, Ciprian Breazu, Bogdan Butoi, and Teodor Stratulat. "Sentiment Analysis in Multiple Languages: Feature Selection for Opinion Classification in Web Forums."
- [5] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." *arXiv preprint arXiv:1810.04805*.
- [6] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention Is All You Need. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [7] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735-1780.
- [8] Graves, A., Wayne, G., & Danihelka, I. (2014). Neural Turing machines. *arXiv preprint arXiv:1410.5401*.
- [9] Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*
- [10] Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*