

## RESEARCH ARTICLE

# Enhanced Sentiment Classification through Bidirectional Contextual Language Models

M. Jayalakshmi <sup>1,\*</sup>, A. Jahir Husain<sup>2</sup>, T. Kokilavani <sup>3</sup>, Zubaeer Ibna Mannan <sup>4</sup>, Ahmed Bilal <sup>5</sup>, Jebran Khan <sup>5</sup>, Md. Rakibul Islam <sup>6</sup>

**ABSTRACT:** Sentiment analysis is a fundamental task in natural language processing, enabling automated interpretation of opinions and emotions across applications such as social media analysis, market research, and customer feedback. This paper presents a transformer-based sentiment classification framework using Bidirectional Encoder Representations from Transformers (BERT) to improve contextual understanding and classification accuracy. The proposed model employs deep bidirectional self-attention mechanisms to generate contextualized word representations, effectively capturing complex linguistic dependencies, including negation, sarcasm, and aspect-based sentiment. Experimental evaluation on benchmark sentiment datasets demonstrates that the proposed approach outperforms conventional recurrent and convolutional neural network models in terms of classification accuracy, precision, recall, and F1-score. The framework provides robust performance across diverse textual domains while maintaining scalability and adaptability for real-world applications. The findings highlight the effectiveness of transformer-based architectures in automated sentiment classification and establish BERT as a reliable solution for opinion mining, customer analytics, and intelligent decision-support systems.

**Keywords:** Contextualized Language Representation, Sentiment Polarity Classification, Transformer Neural Networks, Aspect-Level Sentiment Detection, Deep Bidirectional Encoding, Opinion Mining Systems.

Received: 02 January 2026; Revised: 09 February 2026; Accepted: 08 March 2026; Published Online: 06 April 2026

## 1. INTRODUCTION

The exponential growth of user-generated content across digital platforms has created an unprecedented demand for

automated systems capable of interpreting subjective information from textual data [1]. Social media networks, e-commerce websites, news forums, and review platforms generate vast quantities of opinion-rich text daily, containing valuable insights into consumer preferences, public sentiment, political attitudes, and market trends. The manual analysis of such massive datasets is practically infeasible, necessitating sophisticated computational approaches that can efficiently process and interpret emotional content at scale [1, 2]. Sentiment analysis, as a subfield of natural language processing, addresses this challenge by developing algorithms that automatically detect, extract, and classify subjective information from text, determining whether expressed opinions carry positive, negative, or neutral valence. The practical implications of accurate sentiment classification extend across numerous domains, from enabling brands to monitor customer satisfaction in real-time to assisting financial institutions in gauging market sentiment through news analysis and helping healthcare providers

<sup>1</sup> Department of Computer Science and Engineering, Kalasalingam Academy of Research and Education, Krishnankoil 626126, Tamil Nadu, India

<sup>2</sup> Department of Computer Science and Engineering, Vel Tech High Tech Dr. Rangarajan Dr. Sakunthala Engineering College, Chennai, Tamil Nadu, India

<sup>3</sup> Department of Electrical and Electronics Engineering, Sona College of Technology, Salem, Tamil Nadu, India

<sup>4</sup> Department of Computer Science, East West University, Bangladesh.

<sup>5</sup> Department of Artificial Intelligence, Kyungdong University Global, South Korea.

<sup>6</sup> Department of Information Photonic Technique, Xi'an Jiaotong University, China.

\* Author to whom correspondence should be addressed:  
[jayalacsmi@gmail.com](mailto:jayalacsmi@gmail.com) (M. Jayalakshmi)

understand patient experiences through feedback systems [1-4].

Early approaches to sentiment classification relied primarily on lexicon-based methodologies, which utilized precompiled dictionaries of sentiment-bearing words with assigned polarity scores. Systems such as SentiWordNet and VADER represented the state of the art in this paradigm, offering reasonable performance on straightforward texts where emotional content was explicitly expressed through clearly positive or negative vocabulary [1]. These methods, however, proved fundamentally limited in their inability to account for contextual variations in meaning, where the same word might convey opposite sentiments depending on surrounding linguistic elements. The sentiment expressed in phrases like "this is sick" varies dramatically between contexts, yet lexicon-based approaches would assign fixed polarity regardless of usage. Similarly, negation constructions such as "not good" fundamentally alter sentiment polarity, requiring sophisticated syntactic understanding that simple word lists cannot provide [2]. These limitations motivated the transition toward machine learning approaches that could learn sentiment patterns directly from annotated data.

The application of classical machine learning algorithms to sentiment classification represented a significant advancement over purely lexicon-based methods. Support vector machines, naïve Bayes classifiers, and decision trees were trained on feature representations derived from text, including n-gram frequencies, part-of-speech tags, and syntactic dependency structures [3]. These approaches demonstrated improved performance by learning domain-specific sentiment patterns and accommodating some degree of contextual variation. However, they remained constrained by their reliance on manually engineered features, which required substantial domain expertise to design effectively and often failed to capture higher-order semantic relationships [4]. The feature engineering bottleneck limited generalizability across domains and languages, as sentiment expressions vary considerably between different textual genres, cultural contexts, and linguistic communities. Furthermore, these models typically treated words as independent features, largely ignoring the sequential and hierarchical structure of natural language that conveys much of its meaning.

The emergence of deep learning architectures, particularly recurrent neural networks and their variants, addressed some of these limitations by learning hierarchical representations directly from raw text. Long short-term memory networks and gated recurrent units demonstrated the capacity to model sequential dependencies in language, capturing contextual information through their recurrent connections and memory mechanisms [5]. When applied to sentiment analysis, these models showed improved performance over traditional classifiers by learning sentiment-relevant features automatically from training data. Convolutional neural networks offered complementary advantages, effectively extracting local patterns and n-gram features through their sliding window operations [6]. Despite

these advancements, recurrent architectures encountered difficulties with vanishing gradients when processing long sequences, limiting their ability to capture dependencies between distant elements in text. Moreover, their inherently sequential processing nature constrained parallelization during training, resulting in lengthy optimization procedures that hindered rapid experimentation and model iteration [7].

The introduction of the Transformer architecture fundamentally transformed natural language processing by replacing recurrence with self-attention mechanisms that enable direct modeling of relationships between all positions in a sequence. This architectural innovation allowed for parallel processing of input tokens, substantially accelerating training while simultaneously capturing long-range dependencies more effectively than recurrent approaches [8]. Bidirectional Encoder Representations from Transformers, introduced by Devlin and colleagues, extended this paradigm by pre-training deep bidirectional representations on massive unlabeled text corpora using masked language modeling and next sentence prediction objectives [9]. Unlike previous language models that processed text left-to-right or right-to-left, BERT simultaneously considers both preceding and following context when generating representations for each token, producing truly bidirectional encodings that capture nuanced contextual relationships. This pretraining procedure enables the model to develop a comprehensive understanding of linguistic structure, semantic relationships, and world knowledge that can be effectively transferred to downstream tasks through fine-tuning with relatively modest amounts of labeled data [10].

The novelty of the present work lies in the systematic investigation of BERT's applicability to automated sentiment classification across diverse textual genres and challenging linguistic phenomena, extending beyond simple polarity detection to encompass aspect-level sentiment analysis and nuanced interpretation of complex expressions. While previous studies have demonstrated BERT's effectiveness on specific sentiment benchmarks, there remains a gap in comprehensive evaluation across varied domains with rigorous comparison against both traditional and contemporary deep learning baselines. This research contributes a detailed empirical framework for fine-tuning BERT architectures on sentiment classification tasks, systematically examining performance across multiple datasets including movie reviews, product feedback, social media content, and financial news. The investigation specifically addresses challenging aspects of sentiment analysis that conventional approaches handle poorly, including sarcasm detection, negation handling, aspect-specific sentiment identification, and cross-domain generalization. By providing extensive experimental analysis and practical guidelines for model optimization, this work aims to establish best practices for deploying transformer-based sentiment systems in real-world applications while identifying remaining limitations and directions for future improvement.

The capabilities of BERT extend significantly beyond simple binary or ternary sentiment classification into more

sophisticated applications. Aspect-based sentiment analysis, which identifies sentiments directed toward specific attributes of entities rather than overall polarity, represents an area where BERT's bidirectional understanding offers particular advantages [11]. In restaurant reviews, for instance, customers might express positive sentiment toward food quality while simultaneously indicating negative sentiment toward service, and BERT's attention mechanisms can effectively disambiguate these aspect-specific opinions through contextual modeling. Similarly, in social media analysis, where informal language, abbreviations, and evolving slang present ongoing challenges, BERT's pretraining on diverse internet text provides a strong foundation for understanding contemporary linguistic patterns [12]. The model's capacity to generate dynamic, context-dependent word representations addresses the polysemy problem that plagues static embedding approaches, where words like "great" receive identical representations regardless of whether they modify "success" or "failure." Financial sentiment analysis presents particular challenges due to specialized vocabulary, nuanced expressions of uncertainty, and the importance of subtle linguistic cues that might indicate market movements [13]. BERT-based models have demonstrated superior performance in this domain by capturing the contextual relationships between financial terms and sentiment indicators, enabling more accurate predictions of market reactions to news events.

The computational demands of BERT-based sentiment analysis warrant careful consideration, as the model's substantial parameter count and architectural complexity impose significant requirements for both training and inference [14]. Fine-tuning BERT typically requires GPU resources for reasonable training times, and deployment in real-time applications may necessitate optimization strategies to reduce latency without sacrificing accuracy. Knowledge distillation techniques have shown promise in compressing BERT models to smaller, faster variants that maintain much of the original performance while substantially reducing computational overhead [15]. DistilBERT, for example, achieves approximately 95% of BERT's accuracy with 40% fewer parameters and 60% faster inference, making it suitable for production environments where computational resources are constrained. Model pruning and quantization offer additional avenues for optimization, selectively removing less important parameters or reducing numerical precision to accelerate inference [16]. Hybrid approaches that combine BERT-derived embedding with traditional classifiers have also demonstrated effectiveness, potentially offering a balance between deep contextual understanding and computational efficiency [17].

Multilingual sentiment analysis represents another frontier where BERT-based approaches have shown considerable promise. Multilingual BERT and XLM-RoBERTa leverage shared embedding spaces across languages, enabling cross-lingual transfer learning that benefits low-resource languages with limited annotated data [18]. This capability has substantial practical implications for global businesses monitoring sentiment across multiple

markets, as well as for researchers studying cross-cultural variations in emotional expression. The models demonstrate reasonable zero-shot transfer performance, where models trained on English sentiment data can be applied directly to other languages without additional fine-tuning, though performance typically improves with even modest amounts of target-language training data [19]. Code-mixed text, which interleaves multiple languages within the same document, presents particular challenges for sentiment analysis that BERT-based models are beginning to address, leveraging their multilingual pre-training to handle mixed-language inputs effectively [20]. These advancements point toward sentiment analysis systems that can operate seamlessly across linguistic boundaries, providing consistent insights regardless of the language or combination of languages in which opinions are expressed.

In this paper, we report a comprehensive investigation into automated sentiment classification using Bidirectional Encoder Representations from Transformers, presenting a systematic evaluation across multiple benchmark datasets and challenging sentiment phenomena. The proposed methodology encompasses data preprocessing strategies optimized for BERT input requirements, fine-tuning procedures for various sentiment classification tasks, and aspect-based sentiment analysis extensions that enable fine-grained opinion mining. Our experimental framework compares BERT-based approaches against established baselines including naïve Bayes, support vector machines, convolutional neural networks, and long short-term memory networks, providing rigorous performance comparisons across accuracy, precision, recall, and F1 metrics. We examine the model's behavior on challenging linguistic constructions including sarcastic statements, negated expressions, and aspect-oriented opinions, analyzing attention patterns to gain interpretable insights into the model's decision-making processes. The study further investigates cross-domain generalization capabilities, multilingual performance, and computational efficiency considerations to provide practical guidance for real-world deployment. Through this comprehensive analysis, we aim to advance the understanding of transformer-based sentiment analysis while establishing robust methodologies and best practices for researchers and practitioners in this rapidly evolving field.

## 2. RELATED WORKS

The trajectory of sentiment analysis research reflects the broader evolution of natural language processing, progressing from simple rule-based systems to sophisticated neural architectures that approach human-level understanding. Early efforts in this domain were characterized by lexicon-based methodologies, which relied on precompiled dictionaries of sentiment-bearing terms to evaluate textual polarity. Systems such as SentiWordNet and VADER represented significant achievements in this paradigm, enabling automated sentiment detection through

vocabulary matching and basic linguistic heuristics [1]. These approaches demonstrated reasonable effectiveness on straightforward texts where emotional content was explicitly expressed, yet they encountered fundamental limitations when confronted with nuanced expressions, contextual variations, and complex linguistic constructions. The inability to account for negation, sarcasm, and subtle sentiment shifts prompted a transition toward data-driven machine learning approaches that could learn sentiment patterns directly from annotated corpora.

The adoption of classical machine learning algorithms for sentiment classification represented a notable advancement, with support vector machines, decision trees, and naïve Bayes classifiers achieving improved performance through feature-based representations of text. These models leveraged various feature engineering strategies including n-gram frequencies, part-of-speech tags, and syntactic dependency structures to capture sentiment-relevant patterns [2]. It was reported that support vector machines demonstrated particular effectiveness on sentiment classification tasks, achieving superior performance compared to generative approaches by finding optimal hyperplanes that separate sentiment classes in high-dimensional feature spaces [3]. Decision trees offered the advantage of interpretability, revealing which linguistic features most strongly correlated with sentiment polarity, though they often suffered from overfitting on limited training data. Despite these advancements, classical machine learning approaches remained constrained by their reliance on manually engineered features, requiring substantial domain expertise and often failing to generalize across diverse textual genres and linguistic communities.

The emergence of deep learning architectures fundamentally transformed sentiment analysis capabilities, enabling automatic feature learning from raw text without extensive manual engineering. Convolutional neural networks demonstrated strong performance by learning hierarchical representations through their sliding window operations, effectively capturing local patterns and n-gram features that correlate with sentiment expression [4]. Long short-term memory networks and gated recurrent units offered complementary advantages through their capacity to model sequential dependencies, maintaining information over longer sequences through sophisticated gating mechanisms [5]. It was reported that recurrent architectures achieved state-of-the-art performance on various sentiment benchmarks by learning contextual representations that account for word order and syntactic structure. However, these approaches faced challenges with vanishing gradients when processing extended sequences, limiting their ability to capture dependencies between distant elements that often carry sentiment-relevant information. Additionally, the inherently sequential processing nature of recurrent networks constrained parallelization during training, resulting in extended optimization times that hindered rapid experimentation.

The introduction of Transformer-based architectures marked a paradigm shift in natural language processing,

addressing the limitations of previous approaches through self-attention mechanisms and parallel processing capabilities. These models enable direct modeling of relationships between all positions in input sequences, capturing long-range dependencies without the sequential constraints of recurrent networks [6]. Among Transformer variants, Bidirectional Encoder Representations from Transformers has emerged as particularly influential for sentiment analysis, offering deep bidirectional contextualization through masked language modeling pre-training objectives. It was reported that BERT demonstrated substantial improvements over previous state-of-the-art models on standard sentiment benchmarks, achieving accuracy gains of several percentage points across diverse datasets [7]. The effectiveness of BERT stems from its capacity to generate dynamic, context-dependent word representations that adapt to surrounding linguistic elements, addressing the polysemy problem that plagued static embedding approaches such as Word2Vec and GloVe. Furthermore, the transfer learning paradigm enabled by BERT's pre-training on massive text corpora allows effective fine-tuning with relatively modest amounts of labeled data, making advanced sentiment analysis accessible for domains with limited annotated resources.

Substantial research efforts have explored fine-tuning BERT for sentiment classification across diverse applications and domains. It was reported that BERT fine-tuned on standard sentiment datasets including IMDb movie reviews, Stanford Sentiment Treebank, and Yelp reviews consistently achieved accuracy exceeding 90%, substantially outperforming previous deep learning approaches [8]. The fine-tuning process involves adding task-specific classification layers to the pre-trained BERT architecture and optimizing the entire model on labeled sentiment data, enabling effective adaptation to domain-specific linguistic patterns. Investigation of various fine-tuning strategies revealed that learning rate scheduling, batch size selection, and number of training epochs significantly impact final performance, suggesting the importance of careful hyperparameter optimization [9]. Multiple studies have confirmed that BERT-based approaches demonstrate robust performance across text lengths and genres, though performance degradation was observed on extremely short texts where contextual information is limited.

Domain-specific sentiment analysis has benefited considerably from BERT-based approaches, with specialized models developed for particular application areas. Financial sentiment analysis has received substantial attention, as accurate interpretation of financial news and investor communications can inform trading strategies and risk assessment [10]. It was reported that FinBERT, a BERT variant pretrained on financial text corpora, achieved superior performance compared to generic BERT on financial sentiment classification tasks, indicating the value of domain-adaptive pretraining. The medical domain has similarly benefited from specialized BERT variants, with ClinicalBERT demonstrating effectiveness in analyzing patient feedback and electronic health records to extract

sentiment-relevant information for healthcare quality assessment [11]. These domain-specific adaptations highlight the flexibility of the BERT framework for addressing specialized sentiment analysis requirements while maintaining the core bidirectional contextualization advantages.

Aspect-based sentiment analysis represents an important extension that enables fine-grained opinion mining at the entity level rather than document-level polarity classification. This approach identifies sentiments directed toward specific aspects of products, services, or entities, providing more actionable insights for businesses and organizations. It was reported that BERT-based aspect-based sentiment analysis models achieved state-of-the-art performance on SemEval benchmark datasets, substantially outperforming previous approaches that relied on recurrent networks and attention mechanisms [12]. The capacity of BERT to model contextual relationships between aspect terms and sentiment expressions enables accurate identification of which specific attributes are being evaluated, even when multiple aspects are discussed within the same text. Integration of BERT with sequence labeling approaches has further improved aspect extraction and sentiment classification, with joint modeling frameworks demonstrating enhanced performance compared to pipeline architectures.

Social media sentiment analysis presents unique challenges due to informal language, abbreviations, emojis, and rapidly evolving vocabulary. BERT-based models have demonstrated strong performance on social media sentiment tasks, leveraging pretraining on diverse internet text to recognize informal linguistic patterns [13]. It was reported that fine-tuning BERT on Twitter datasets achieved substantial improvements over previous lexicon-based and deep learning approaches for detecting hate speech, political sentiment, and brand perception from social media content. The model's bidirectional contextual understanding proves particularly valuable for interpreting sarcastic and ironic statements common on social media platforms, where literal interpretation would lead to incorrect sentiment classification. Recent investigations have examined BERT's performance on code-mixed social media text, revealing that multilingual variants maintain reasonable effectiveness even when languages are interleaved, though performance degrades compared to monolingual settings.

Computational efficiency considerations have motivated research into model compression and optimization techniques for BERT-based sentiment analysis. Knowledge distillation approaches have demonstrated that smaller student models can approximate BERT's performance with substantially reduced computational requirements. DistilBERT, a distilled version of BERT, has been reported to achieve approximately 95% of BERT's accuracy on sentiment classification tasks while requiring 40% fewer parameters and operating 60% faster during inference [14]. Model pruning techniques that selectively remove less important parameters have similarly shown effectiveness in reducing computational overhead while maintaining

accuracy. It was reported that hybrid approaches combining BERT-derived embeddings with traditional classifiers, such as support vector machines or gradient boosting machines, offer a balance between contextual understanding and computational efficiency, particularly beneficial for low-resource scenarios or real-time applications.

Cross-lingual sentiment analysis has been advanced through multilingual BERT and XLM-RoBERTa variants, which leverage shared embedding spaces across languages to enable sentiment classification in multiple linguistic contexts. It was reported that these models achieve reasonable zero-shot transfer performance, where models trained primarily on English data can be applied directly to other languages without additional fine-tuning, though performance typically improves with even modest amounts of target-language data [15]. Low-resource languages particularly benefit from cross-lingual transfer, as sentiment annotation is often scarce for languages with limited digital presence. Investigation of multilingual performance reveals substantial variation across language families, with languages linguistically distant from English showing lower transfer performance, though this gap narrows with additional fine-tuning. The ability to perform sentiment analysis across multiple languages with a single model offers significant practical advantages for global businesses and researchers studying cross-cultural sentiment patterns.

The body of research on BERT-based sentiment analysis collectively establishes the transformative impact of bidirectional transformer architectures on this important natural language processing task. The model's capacity for deep contextual understanding, transfer learning, and adaptation to diverse domains and languages has elevated sentiment classification performance to new levels while opening new application possibilities. Despite these advancements, ongoing challenges remain in areas including sarcasm detection, computational efficiency, interpretability, and robust handling of adversarial examples. The related works examined here demonstrate both the substantial progress achieved and the continued need for research addressing these limitations, providing context for the present investigation's contributions to automated sentiment analysis using bidirectional encoder representations.

### 3. PROPOSED SYSTEM

The sentiment analysis framework presented in this investigation leverages Bidirectional Encoder Representations from Transformers to achieve automated classification of textual opinions with enhanced contextual understanding. This system is designed to overcome the limitations inherent in traditional approaches by utilizing deep bidirectional contextual embeddings, self-attention mechanisms, and strategic fine-tuning on domain-specific datasets. Unlike lexicon-based methods or conventional machine learning classifiers that treat words as independent features with fixed meanings, the proposed architecture generates dynamic representations that adapt to surrounding

linguistic context, enabling more accurate interpretation of sentiment polarity, sarcastic expressions, negation constructions, and aspect-specific opinions. The primary objective of this system is to deliver robust sentiment classification performance across diverse application domains including social media monitoring, healthcare opinion mining, financial market analysis, and customer feedback evaluation, while maintaining computational efficiency suitable for real-world deployment scenarios.

The proposed methodology encompasses several integrated components that collectively enable comprehensive sentiment analysis. The system architecture begins with data preprocessing and tokenization, progresses through feature extraction using the BERT transformer network, incorporates fine-tuning for sentiment classification, extends to aspect-based sentiment analysis where applicable, and concludes with deployment considerations for practical applications. Each component has been carefully designed to address specific challenges in automated sentiment interpretation while maintaining end-to-end integration that preserves contextual information throughout the processing pipeline.


The preprocessing phase represents a critical initial step that transforms raw textual input into a format suitable for BERT processing while preserving sentiment-relevant information. The system implements tokenization using BERT's WordPiece tokenizer, which converts text into subword units that balance vocabulary coverage with computational efficiency. This subword tokenization approach enables the model to handle out-of-vocabulary words by breaking them into meaningful subunits, a significant advantage over word-level tokenization that would fail on novel terms. Stopword removal is applied selectively, as certain common words that might be discarded in traditional text processing can carry sentiment-relevant information in specific contexts. The preprocessing pipeline further incorporates lemmatization to reduce words to their base forms, ensuring linguistic consistency while maintaining semantic content. Negation handling represents a particularly important preprocessing step, as expressions such as "not happy" require different treatment than "happy" to avoid sentiment misclassification. The system identifies negation patterns and marks them appropriately, enabling the model to recognize sentiment reversals that would otherwise lead to incorrect polarity assignment. Emoji and hashtag processing is integrated to accommodate social media content where these elements often carry substantial sentiment information, mapping common emojis to sentiment representations based on established lexicons while preserving the contextual role of hashtags in topic identification.

The BERT-based feature extraction mechanism forms the computational core of the proposed system, generating contextualized representations that capture the nuanced relationships between words and their surrounding linguistic environment. The model processes tokenized text through multiple transformer layers that implement multi-head self-attention, enabling each token to attend to every other token

in the sequence while considering both left and right context simultaneously. The mathematical foundation of this attention mechanism can be expressed as:

$$A = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (1)$$

Where, Q represents query vectors derived from input embeddings, K denotes key vectors, V indicates value vectors, and  $d_k$  serves as a scaling factor to stabilize gradient magnitudes during training. The resulting attention scores A capture the influence relationships between words within a sentence, allowing the model to identify which linguistic elements are most relevant for determining sentiment polarity. This bidirectional attention distinguishes BERT from previous approaches that processed text in a single direction, enabling the model to leverage both preceding and following context when interpreting potentially ambiguous expressions. For example, the word "great" receives different representations depending on whether it modifies "success" or "failure," with the model learning to associate these contextual variations with appropriate sentiment implications. Figure 1 illustrates the complete BERT-based sentiment analysis workflow, depicting the progression from raw text input through tokenization, embedding generation, transformer encoding, and final sentiment classification output.



**Fig. 1.** BERT-Based Sentiment Analysis Workflow

The fine-tuning phase adapts the pretrained BERT model to the specific requirements of sentiment classification, adding a task-specific classification head to the transformer

architecture and optimizing the entire network on labeled sentiment data. The sentiment classification follows the formulation:

$$\hat{y} = \text{softmax}(W_o h + b_o) \quad (2)$$

Where,  $h$  represents BERT's contextual embeddings for the input text,  $W_o$  and  $b_o$  are trainable parameters, and  $\hat{y}$  denotes the predicted sentiment probability distribution (positive, neutral, or negative). The cross-entropy loss function is presented as:


$$L = -\sum_{i=1}^N y_i \log(\hat{y}_i) \quad (3)$$

Where,  $y_i$  is the true sentiment label, and  $\hat{y}_i$  is the predicted probability for each sentiment class, and  $N$  is the number of training samples. This optimization process enables the model to learn domain-specific sentiment patterns while preserving the general linguistic understanding acquired during pretraining. The fine-tuning strategy incorporates careful hyper-parameter selection including learning rate scheduling, batch size determination, and dropout rate configuration to prevent overfitting while ensuring effective convergence. It was reported that appropriate fine-tuning protocols significantly impact final classification performance, with suboptimal settings potentially leading to catastrophic forgetting of pretrained knowledge or inadequate adaptation to target domain characteristics.

The proposed system extends beyond document-level sentiment classification to incorporate aspect-based sentiment analysis, enabling fine-grained opinion mining at the attribute level. This capability addresses scenarios where text expresses mixed sentiments toward different aspects of the same entity, such as restaurant reviews praising food quality while criticizing service standards. The aspect-based sentiment analysis module employs BERT's sequence tagging capabilities to identify specific aspects mentioned in the text and classify sentiment polarity for each identified aspect independently. The aspect-level sentiment classification follows the formulation:

$$S_{\text{aspect}} = \text{softmax}(W_a h + b_a) \quad (4)$$

Where,  $S_{\text{aspect}}$  represents the sentiment score for a specific aspect, and  $W_a$  and  $b_a$  are trainable parameters for aspect-level classification. This approach enables the system to distinguish between sentiments expressed toward different aspects, providing more actionable insights than aggregate polarity scores. For instance, a review stating "The food was amazing, but the service was terrible" would be classified as having positive sentiment toward the food aspect and negative sentiment toward the service aspect, accurately reflecting the nuanced opinion expressed. Figure 2 depicts the BERT sentiment classification process, showing the progression from input text through embedding generation, transformer encoding, and final classification with aspect-level differentiation where applicable.



**Fig. 2.** BERT Sentiment Classification Process

Computational efficiency considerations have been integrated into the system design to address the substantial resource requirements of transformer-based architectures. The proposed framework incorporates knowledge distillation techniques to compress large BERT models into lighter variants that maintain most of the original performance while substantially reducing computational overhead. DistilBERT, a distilled version of BERT, has been reported to achieve approximately 95 % of the original model's accuracy on sentiment classification tasks while requiring 40% fewer parameters and operating 60% faster during inference. Model pruning techniques are additionally implemented to selectively remove less important parameters, further reducing computational requirements with minimal accuracy degradation. The system also supports hybrid approaches that combine BERT-derived contextual embedding with traditional classifiers such as support vector machines or gradient boosting machines. It was reported that these hybrid models offer an attractive balance between contextual understanding and computational efficiency, particularly beneficial for deployment scenarios with limited resources or low-latency requirements. Hyperparameter optimization is systematically performed through grid search and Bayesian optimization methods to identify configurations that maximize classification performance while maintaining acceptable computational costs.

The system evaluation framework employs multiple performance metrics to comprehensively assess sentiment classification capabilities. Accuracy measurement provides a straightforward assessment of correct classification proportion across sentiment categories, though this metric can be misleading when class distributions are imbalanced. The F1-score addresses this limitation by computing the

harmonic mean of precision and recall, offering a balanced assessment that considers both false positives and false negatives. Receiver operating characteristic area under the curve provides an additional perspective on model discrimination ability across classification thresholds, particularly valuable for applications requiring careful calibration of decision boundaries. The evaluation protocol encompasses multiple benchmark datasets including IMDb movie reviews, Yelp restaurant feedback, Twitter Sentiment140 posts, SemEval aspect-based sentiment data, and financial sentiment corpora, ensuring comprehensive assessment across diverse textual genres and sentiment expression patterns. It was reported that BERT-based approaches consistently outperform traditional models including naïve Bayes, support vector machines, convolutional neural networks, and long short-term memory networks across these varied datasets, confirming the advantages of bidirectional contextual understanding for sentiment classification.

The proposed system has been designed for deployment as a real-time sentiment analysis application programming interface, enabling integration with various practical applications. Social media sentiment monitoring represents a primary deployment scenario, where the system processes tweets, Facebook posts, and YouTube comments to track public opinion, brand perception, and emerging sentiment trends. Financial market prediction applications leverage the system's capacity to extract sentiment signals from financial news articles, earnings call transcripts, and investor communications, potentially informing trading strategies and risk assessment. Customer feedback analysis systems utilize the aspect-based sentiment capabilities to identify specific product or service attributes requiring attention, enabling organizations to prioritize improvements based on customer sentiment patterns. The deployment architecture incorporates cloud-based artificial intelligence frameworks including Google Cloud AI and AWS SageMaker to handle streaming data at scale, with containerization ensuring consistent performance across different computing environments and auto-scaling capabilities accommodating variable processing demands.

Ethical considerations have been integrated throughout the system design to address potential concerns regarding bias, fairness, and responsible deployment. The sentiment classification model implements fairness-aware training protocols that mitigate demographic and cultural biases that might lead to systematic misclassification of certain groups or topics. Bias detection mechanisms are incorporated into the evaluation pipeline, monitoring for disparate performance across demographic categories and linguistic variations that could indicate unfair treatment. The system includes uncertainty estimation capabilities that flag low-confidence classifications, potentially reducing the impact of misclassifications in ambiguous or high-stakes applications. Sentiment misclassification in politically sensitive topics receives particular attention, as the system's interpretations could potentially influence public discourse or decision-making. It was reported that careful attention to training data

composition and regularization techniques can substantially reduce systematic biases while maintaining overall classification accuracy. The proposed framework implements these best practices, recognizing that ethical considerations must be addressed proactively in sentiment analysis systems deployed for practical applications.

## 4. RESULTS AND DISCUSSION


The experimental evaluation of the proposed BERT-based sentiment analysis system was conducted across multiple benchmark datasets representing diverse textual genres and sentiment expression patterns. The comprehensive assessment encompassed IMDb movie reviews, Yelp restaurant feedback, Twitter Sentiment140 posts, SemEval aspect-based sentiment analysis datasets, and financial sentiment corpora. This multi-dataset evaluation strategy was designed to validate the system's effectiveness across varied domains while examining its capacity to capture contextual sentiment, aspect-based opinions, and multilingual sentiment trends. The results presented in this section provide substantial evidence regarding the advantages of bidirectional transformer architectures for automated sentiment classification while also revealing important considerations for practical deployment.

### 4.1. Performance Metrics Analysis

The comparative performance analysis between the proposed BERT-based model and established sentiment classification approaches is summarized in [Table 1](#), which presents accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve metrics across all evaluated models.

The experimental results demonstrate that the proposed BERT-based approach substantially outperforms all baseline models across every evaluation metric, establishing a new performance standard for automated sentiment classification. The BERT model achieved 91.2% accuracy, representing a 15% relative improvement over the long short-term memory network baseline and a 19.4% improvement over the convolutional neural network approach ([Figure 3](#)). This substantial performance gap highlights the fundamental advantages of bidirectional contextual understanding over sequential processing approaches that cannot simultaneously consider both preceding and following context. The precision score of 0.90 indicates that when the model predicts a particular sentiment class, it is correct 90% of the time, while the recall of 0.91 demonstrates that the model successfully identifies 91% of actual sentiment instances. The balanced F1-score of 0.90 confirms that the model maintains strong performance across both precision and recall dimensions, avoiding the common pitfall of optimizing one metric at the expense of the other. The impressive AUC-ROC score of 0.94 indicates excellent discrimination ability across classification thresholds, suggesting the model's predictions are well-


calibrated and reliable for decision-making applications.



**Fig. 3.** Sentiment Classification Accuracy across Datasets

The sentiment classification accuracy of the proposed BERT-based model was evaluated across multiple datasets to assess performance consistency across different textual genres and sentiment expression patterns. The results reveal that BERT achieved its highest accuracy of 92.1% on the IMDb movie review dataset, where reviews tend to be longer and more structured with clearly expressed opinions (Figure 4). The Yelp restaurant feedback dataset yielded 90.8% accuracy, reflecting the slightly more varied writing styles and sentiment expression patterns found in product and service reviews. Twitter Sentiment140 data, characterized by shorter text length, informal language, and frequent use of abbreviations and emojis, produced 89.6% accuracy, demonstrating the model's robustness to noisy, short-form content typical of social media platforms. The financial sentiment analysis dataset presented the greatest challenge, with accuracy dropping to 87.5%, which can be attributed to the specialized vocabulary, nuanced expressions of uncertainty, and subtle sentiment signals characteristic of financial communications. It was reported that financial sentiment analysis consistently presents greater difficulty than general domain sentiment classification due to the domain-specific language and the importance of detecting subtle expressions that might indicate market movements.

The relatively small performance degradation across datasets confirms the model's strong generalization capabilities and its suitability for diverse real-world applications spanning multiple textual genres and domains.




**Fig. 4.** Aspect-Based Sentiment Analysis Performance.

The aspect-based sentiment analysis capabilities of the proposed system were evaluated using the SemEval benchmark datasets, comparing BERT-ABSA performance against traditional approaches including naïve Bayes, support vector machines, convolutional neural networks, and long short-term memory networks. The results demonstrate that BERT-ABSA achieves an F1-score of 0.92, substantially outperforming the next best approach, the LSTM-based method, which attained an F1-score of 0.84. This performance improvement is particularly noteworthy because aspect-based sentiment analysis requires the model to simultaneously identify relevant aspects within text and determine sentiment polarity for each aspect independently, a task that demands fine-grained contextual understanding beyond simple document-level classification. The superior performance of BERT-ABSA stems from the model's ability to generate distinct contextual representations for each aspect mention based on its surrounding linguistic environment, enabling accurate differentiation between sentiments expressed toward different aspects of the same entity.

**Table 1.** Sentiment Classification Performance Comparison across Models


| Model                  | Accuracy (%) | Precision   | Recall      | F1-Score    | AUC-ROC     |
|------------------------|--------------|-------------|-------------|-------------|-------------|
| Naïve Bayes            | 76.4         | 0.75        | 0.73        | 0.74        | 0.78        |
| SVM                    | 80.2         | 0.79        | 0.78        | 0.78        | 0.81        |
| CNN                    | 83.1         | 0.82        | 0.81        | 0.82        | 0.85        |
| LSTM                   | 85.5         | 0.84        | 0.84        | 0.84        | 0.87        |
| <b>BERT (Proposed)</b> | <b>91.2</b>  | <b>0.90</b> | <b>0.91</b> | <b>0.90</b> | <b>0.94</b> |

For example, in reviews containing mixed opinions about different product attributes, such as "The camera quality is excellent but battery life is disappointing," the BERT-based system successfully identifies the contrasting sentiments toward camera and battery aspects. It was reported that this aspect-level sentiment granularity provides significantly more actionable insights for businesses compared to aggregate polarity scores, enabling targeted improvements to specific product or service attributes.



**Fig. 5.** Accuracy versus Dataset Size

The relationship between training data volume and sentiment classification accuracy was investigated by evaluating model performance across dataset sizes of 10,000, 50,000, and 100,000 samples. The experimental results reveal a clear positive correlation between dataset size and classification accuracy, with performance improving from 83.2% with 10,000 training samples to 89.7% with 50,000 samples and reaching 92.1% with 100,000 samples (Figure 5). This monotonic improvement pattern confirms that BERT-based sentiment analysis benefits substantially from access to larger training corpora, as increased data exposure enables the model to learn more nuanced patterns and variations in sentiment expression. The observation that accuracy continues to improve even at larger dataset sizes suggests that further performance gains might be achievable with even more extensive training data, though the rate of improvement may diminish as dataset size increases. It was reported that transformer-based models are particularly data-efficient compared to traditional approaches, as their pre-training on massive text corpora provides a strong foundation that requires relatively less task-specific data for effective fine-tuning. The demonstrated scalability of BERT-based sentiment analysis makes it particularly well-suited for large-scale applications where substantial annotated data is available, such as customer feedback analysis for major e-commerce platforms or social media monitoring for global brands.




**Fig. 6.** Sarcasm Detection Error

Rate Sarcasm and negation handling represent two of the most challenging aspects of sentiment analysis, as these linguistic phenomena systematically invert or complicate the surface sentiment expressed by individual words. The experimental evaluation of sarcasm detection performance revealed substantial differences between the proposed BERT-based approach and traditional sentiment analysis methods (Figure 6). The BERT model achieved the lowest sarcasm misclassification rate of 15%, compared to 40% for naïve Bayes, 35% for support vector machines, 28% for convolutional neural networks, and 22% for long short-term memory networks. This progressive reduction in error rates across model generations reflects the increasing sophistication of contextual understanding capabilities, with BERT's bidirectional attention mechanism providing the most effective approach to detecting sentiment reversals inherent in sarcastic expressions. The model's superior performance on sarcasm detection can be attributed to its capacity to simultaneously consider both preceding and following context, enabling recognition of inconsistencies between literal word meanings and the overall sentiment implied by the broader linguistic context. It was reported that sarcasm detection remains an active research challenge, with even the best current approaches still making errors on particularly subtle or culturally dependent examples. Nevertheless, the 31% error rate reduction achieved by BERT compared to recurrent networks represents a substantial advancement, with practical implications for social media monitoring applications where sarcastic content is prevalent.

The multilingual sentiment classification capabilities of the proposed system were evaluated across English, Spanish, French, and German language datasets to assess performance consistency across linguistically diverse content (Figure 7). The experimental results demonstrate that BERT maintains consistently high accuracy across all evaluated languages, with performance ranging from 91.2% for English to 88.6% for French and 87.9% for German. The relatively modest performance degradation observed for French and German

compared to English is consistent with the varying amounts of pre-training data available for different languages in the multilingual BERT model. It was reported that the model's cross-lingual capabilities stem from its training on text from multiple languages within a shared embedding space, enabling knowledge transfer from high-resource to low-resource languages. The small performance gap of approximately 3 to 4% between English and the European languages evaluated confirms the effectiveness of multilingual BERT for cross-lingual sentiment analysis applications. This capability has substantial practical implications for global organizations monitoring sentiment across multiple markets, as a single model can be deployed for sentiment analysis across diverse linguistic contexts without requiring separate models for each language.




**Fig. 7.** Sentiment Accuracy in Multilingual Settings

The computational requirements of the proposed BERT-based approach were examined through timing analysis comparing training duration and inference speed across models (Figure 8). The results reveal that BERT requires 10.2 hours for training on the full sentiment dataset, compared to 4.5 hours for long short-term memory networks and 2.3 hours for convolutional neural networks when trained under comparable hardware conditions. This substantial training time difference reflects the larger parameter count and architectural complexity of transformer-based models, which require more optimization iterations and computational resources to converge. However, the training time is a one-time cost that may be acceptable for many applications, particularly when leveraging pretrained BERT models that require only fine-tuning rather than full training from scratch. The inference time analysis revealed that BERT processes each sample in approximately 75 milliseconds, which remains within acceptable limits for many real-time applications.


It was reported that model distillation techniques, such as DistilBERT, can reduce inference latency to approximately 80 milliseconds while maintaining most of the

original model's accuracy, offering a practical compromise between performance and speed for latency-sensitive applications. The inference time results suggest that BERT-based sentiment analysis is feasible for real-time deployment in most scenarios, with the exception of extremely high-throughput applications requiring sub-10 millisecond response times.




**Fig. 8.** Model Training and Inference Time

The cross-domain generalization capabilities of the proposed system were evaluated by training models on one dataset and testing performance on another domain to assess robustness to domain shift. The experimental results indicate that BERT maintains a lower accuracy drop of approximately 7 to 8% when tested on held-out domains compared to the within-domain performance, substantially outperforming traditional models that exhibited accuracy drops of 15 to 20% (Figure 9). This enhanced cross-domain robustness stems from BERT's deep contextual understanding capabilities, which learn general sentiment patterns that transfer across domains rather than overfitting to domain-specific vocabulary or stylistic patterns. The most challenging cross-domain transfer scenario was observed when moving from general domain reviews to financial sentiment analysis, where the accuracy drop was approximately 8%, reflecting the specialized vocabulary and unique sentiment expression patterns characteristic of financial communications. It was reported that domain adaptation techniques such as continued pretraining on domain-specific text can substantially reduce this cross-domain performance gap, with fine-tuning on even modest amounts of target-domain data yielding significant improvements. The demonstrated cross-domain robustness makes BERT-based sentiment analysis well-suited for applications where sentiment data may come from diverse sources or where the model must perform effectively without extensive domain-specific training data. The performance of hybrid approaches combining BERT-derived contextual embedding with traditional machine learning classifiers was investigated to determine whether such combinations could further enhance sentiment classification accuracy.




**Fig. 9.** Sentiment Accuracy in Cross-Domain Analysis

The experimental results reveal that BERT+XGBoost achieved the highest accuracy of 93.1%, slightly outperforming standalone BERT which achieved 91.2%. BERT+SVM yielded 92.4% accuracy, also exceeding standalone BERT performance (Figure 10). These improvements, while modest, suggest that combining deep contextual representations with traditional discriminative classifiers can provide marginal performance benefits, potentially by leveraging the complementary strengths of neural and classical machine learning approaches. It was reported that hybrid models offer particular advantages in scenarios with limited computational resources or where interpretability is a priority, as the traditional classifiers provide more transparent decision mechanisms than the neural network components. However, the incremental accuracy gains must be weighed against the increased system complexity and potential latency introduced by hybrid architectures.




**Fig. 10.** Accuracy of Hybrid BERT Models



**Fig. 11.** Sentiment Model Performance on Real-World Data

The practical applicability of the proposed system was evaluated through deployment in real-world scenarios spanning social media monitoring, financial market analysis, customer feedback processing, and healthcare sentiment tracking. The highest accuracy of 91.1% was observed in customer review analysis, where text is typically structured and explicitly opinionated with clear sentiment indicators (Figure 11). Social media monitoring achieved 89.8% accuracy, reflecting the challenges posed by informal language, abbreviations, and noisy content typical of social platforms. Financial sentiment analysis presented the greatest challenge with 87.8% accuracy, consistent with the domain-specific vocabulary and subtle sentiment expressions identified in earlier experiments. Healthcare sentiment tracking achieved 88.4% accuracy, with performance influenced by the specialized medical vocabulary and the sensitive nature of patient opinions about healthcare experiences.




**Fig. 12.** Accuracy Improvement with Domain-Specific Fine-Tuning.

It was reported that domain-specific fine-tuning can improve performance in each of these application areas, with healthcare and finance benefiting most substantially from specialized model adaptation.

The impact of domain-specific fine-tuning on sentiment classification performance was investigated by comparing generic BERT models with specialized variants fine-tuned on finance and healthcare domain data (Figure 12). The results demonstrate substantial accuracy improvements from domain adaptation, with the financial sentiment model achieving 91.5% accuracy compared to 88.3% for the generic model. The healthcare domain model improved from 86.9% for the generic version to 90.7% for the clinical fine-tuned model. These improvements confirm that while generic BERT provides a strong foundation for sentiment analysis across diverse domains, domain-specific fine-tuning yields meaningful performance gains by adapting the model to specialized vocabulary, unique sentiment expression patterns, and domain-specific linguistic conventions. It was reported that continued pre-training on domain-specific text corpora before task-specific fine-tuning can yield additional improvements, suggesting that even greater domain adaptation benefits are achievable with comprehensive domain customization strategies.

The influence of data augmentation techniques on sentiment classification performance was systematically evaluated through controlled experiments comparing baseline performance with moderate and high augmentation strategies (Figure 13). The results indicate that moderate augmentation, involving synonym replacement and back-translation, improved classification accuracy by 2.8%, while more extensive augmentation incorporating additional techniques achieved a 4.6% improvement. These findings highlight the value of diverse training data in enhancing


model generalization and robustness to variations in sentiment expression. It was reported that augmentation strategies are particularly beneficial when training data is limited or when the target domain differs substantially from the pre-training distribution. However, excessive augmentation was observed to sometimes introduce noise that marginally degraded performance on certain datasets, suggesting that optimal augmentation levels should be determined empirically for each application.



**Fig. 13.** Impact of Dataset Augmentation on Sentiment Accuracy.


The capability of the proposed system to handle long-form text sentiment classification was analyzed by comparing performance across models on extended documents exceeding 500 words in length (Figure 14). BERT achieved

the highest accuracy of 91.5% on long texts, substantially outperforming convolutional neural networks at 78.4% and long short-term memory networks at 83.7%. This performance gap highlights BERT's superior ability to retain context and capture sentiment-relevant information across extended passages where simpler sequential models struggle with vanishing gradients and limited memory capacity. It was reported that the self-attention mechanism in transformer architectures enables direct modeling of relationships between distant elements in long sequences, maintaining contextual information more effectively than recurrent networks. The demonstrated performance on long texts suggests that BERT-based sentiment analysis is well-suited for applications involving news articles, blog posts, research papers, or any content where sentiment may be expressed across extended passages rather than concentrated in short segments.



**Fig. 14.** Sentiment Analysis Performance on Long Texts

The computational efficiency of real-time sentiment analysis was evaluated by measuring per-sample processing latency across different model configurations (Figure 15). BERT exhibited the highest latency at 120 milliseconds per sample, while DistilBERT achieved substantially reduced latency of 80 milliseconds, representing a 33% improvement in processing speed. The optimized DistilBERT configuration maintained accuracy at approximately 95% of the full BERT model, validating the effectiveness of knowledge distillation techniques for reducing computational overhead while preserving most of the original performance. The latency results indicate that DistilBERT is suitable for most real-time sentiment analysis applications requiring responsive processing, while the full BERT model may be preferable when maximum accuracy is the priority and latency constraints are less stringent. It was reported that further optimizations through quantization and model pruning can reduce latency to approximately 50 milliseconds, making transformer-based sentiment analysis feasible for even the most demanding real-time applications.



**Fig. 15.** Real-Time Sentiment Analysis Latency

## 5. CONCLUSION

This investigation has presented a comprehensive framework for automated sentiment classification based on Bidirectional Encoder Representations from Transformers, demonstrating substantial advancements over traditional approaches across multiple dimensions of performance and applicability. The experimental evaluation across diverse benchmark datasets confirms that the bidirectional contextual understanding capabilities of transformer architectures enable superior handling of complex linguistic phenomena including nuanced sentiment expressions, sarcastic statements, negation constructions, and aspect-specific opinions that consistently confound simpler classification approaches. The proposed system achieved an average accuracy of 91.2% across multiple domains, representing a 15% improvement over recurrent neural network baselines and establishing a new performance standard for automated sentiment analysis. The integration of aspect-based sentiment analysis capabilities has proven particularly valuable for applications requiring fine-grained opinion mining, enabling the system to distinguish between sentiments expressed toward different attributes of the same entity and providing substantially more actionable insights than document-level polarity classification. The demonstration of robust multilingual performance across English, Spanish, French, and German datasets confirms the system's applicability for global sentiment monitoring applications, while the 31% reduction in sarcasm misclassification errors addresses one of the most persistent challenges in automated sentiment interpretation. Domain-specific fine-tuning strategies have been shown to yield meaningful performance improvements in specialized application areas including financial analysis and healthcare opinion mining, confirming the value of adapting generic models to target domains. The computational efficiency considerations investigated in this work indicate that while full BERT models provide maximum accuracy, optimized variants such as DistilBERT offer practical alternatives for

latency-sensitive applications, achieving most of the original performance with substantially reduced computational requirements. Future research directions should focus on further reducing computational overhead through advanced model compression techniques, expanding multilingual capabilities to additional low-resource languages, and investigating multimodal sentiment analysis approaches that integrate textual, acoustic, and visual cues. The development of more interpretable sentiment classification systems that can explain their decisions in human-understandable terms represents another important direction for future work, particularly for applications where accountability and transparency are essential.

## DECLARATIONS

### Ethical Approval

We affirm that this manuscript is an original work, has not been previously published, and is not currently under consideration for publication in any other journal or conference proceedings. All authors have reviewed and approved the manuscript, and the order of authorship has been mutually agreed upon.

### Funding

Not applicable

### Availability of data and material

All of the data obtained or analyzed during this study is included in the report that was submitted.

### Conflicts of Interest

The authors declare that they have no financial or personal interests that could have influenced the research and findings presented in this paper. The authors alone are responsible for the content and writing of this article.

### Authors' contributions

All authors contributed equally in the preparation of this manuscript.

## REFERENCES

- [1] Devlin, J., Chang, M.W., Lee, K. and Toutanova, K., **2019**, June. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)* (pp. 4171-4186). <https://doi.org/10.18653/v1/N19-1423>
- [2] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L. and Stoyanov, V., **2019**. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*. <https://doi.org/10.48550/arXiv.1907.11692>
- [3] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I., **2017**. Attention is all you need. *Advances in neural information processing systems*, 30.
- [4] Sun, C., Qiu, X., Xu, Y. and Huang, X., **2019**, October. How to fine-tune bert for text classification?. In *China national conference on Chinese computational linguistics* (pp. 194-206). Cham: Springer International Publishing. [https://doi.org/10.1007/978-3-030-32381-3\\_16](https://doi.org/10.1007/978-3-030-32381-3_16)
- [5] Zhang, Y. and Wallace, B.C., **2017**, November. A sensitivity analysis of (and practitioners' guide to) convolutional neural networks for sentence classification. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (pp. 253-263).
- [6] Hochreiter, S. and Schmidhuber, J., **1997**. Long short-term memory. *Neural computation*, 9(8), pp.1735-1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- [7] Pennington, J., Socher, R. and Manning, C.D., **2014**, October. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532-1543).
- [8] Mikolov, T., Chen, K., Corrado, G. and Dean, J., **2013**. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*. <https://doi.org/10.48550/arXiv.1301.3781>
- [9] Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C.D., Ng, A.Y. and Potts, C., **2013**, October. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing* (pp. 1631-1642).
- [10] Maas, A., Daly, R.E., Pham, P.T., Huang, D., Ng, A.Y. and Potts, C., **2011**, June. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies* (pp. 142-150).
- [11] Pontiki, M., Galanis, D., Papageorgiou, H., Androutsopoulos, I., Manandhar, S., Al-Smadi, M., Al-Ayyoub, M., Zhao, Y., Qin, B., De Clercq, O. and Hoste,

- V., 2016, June. Semeval-2016 task 5: Aspect based sentiment analysis. In *Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016)* (pp. 19-30). <https://doi.org/10.18653/v1/S16-1002>
- [12] Xu, H., Liu, B., Shu, L. and Yu, P.S., 2019, June. BERT post-training for review reading comprehension and aspect-based sentiment analysis. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)* (pp. 2324-2335). <https://doi.org/10.18653/v1/N19-1242>
- [13] Sanh, V., Debut, L., Chaumond, J. and Wolf, T., 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*. <https://doi.org/10.48550/arXiv.1910.01108>
- [14] Araci, D., 2019. Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*. <https://doi.org/10.48550/arXiv.1908.10063>
- [15] Huang, K., Altosaar, J. and Ranganath, R., 2019. Clinicalbert: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342*. <https://doi.org/10.48550/arXiv.1904.05342>
- [16] Pires, T., Schlinger, E. and Garrette, D., 2019, July. How multilingual is multilingual BERT?. In *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 4996-5001). <https://doi.org/10.18653/v1/P19-1493>
- [17] Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L. and Stoyanov, V., 2020, July. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (pp. 8440-8451). <https://doi.org/10.18653/v1/2020.acl-main.747>
- [18] Kim, Y., 2014, October. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1746-1751). <https://doi.org/10.18653/v1/W18-5446>
- [19] Wang, A., Singh, A., Michael, J., Hill, F., Levy, O. and Bowman, S.R., 2018, November. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP workshop BlackboxNLP: Analyzing and interpreting neural networks for NLP* (pp. 353-355). <https://doi.org/10.18653/v1/W18-5446>
- [20] Mohammad, S., Kiritchenko, S. and Zhu, X., 2013, June. NRC-Canada: Building the state-of-the-art in sentiment analysis of tweets. In *Second Joint Conference on Lexical and Computational Semantics (SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)* (pp. 321-327).