SSRG International Journal of Electronics and Communication Engineering
ISSN: 2348-8549/ https://doi.org/10.14445/23488549/IJECE-V12I8P129

Volume 12 Issue 8, 334-351, August 2025

Original Article

HGTBF-OM: A Hybrid Graph-Based Transformer Framework for Enhanced Opinion Mining in Textual Data

B. Madhurika[1], D. Naga Malleswari[2]

[1,2]Koneru Lakshmaiah Education Foundation, Andhra Pradesh, India.

[1]Corresponding Author : budarajumadhurika@gmail.com

Received: 15 June 2025 Revised: 17 July 2025 Accepted: 18 August 2025 Published: 30 August 2025

Abstract *- As one such process, opinion mining carries immense significance in deriving actionable insights from textual information, facilitating data-driven decisions across sectors. While existing methodologies have attained great strides, including hybrid deep learning models and artificial neural networks, they need help addressing deep sentiments and contextual relationships in complex datasets. The current state-of-the-art approaches are either reliant on heuristics-based methods such as LSTMs or independent graph learning techniques or less sophisticated strategies based on either one of graph-based techniques or domain transformers without ensembling their respective characteristics, thus causing them to lose the power of harmony between the two as a pioneering approach. This study presents the Hybrid Graph Transformer-Based Framework for Opinion Mining (HGTBF-OM) to solve these challenges. It integrates graph representation, TransformerConv layers, and multi-head attention mechanisms. Experimental validation on the Zomato dataset demonstrated that the proposed framework outperformed current models, achieving a commendable accuracy of 99.01%, surpassing TransLSTM, So-haTRed, and Hybrid GCN-RF. The focus on these results demonstrates the framework's facet-wise analysis capacity and its performance on balanced data. It is helpful for e-commerce, healthcare, and social media monitoring applications where the precision of sentiment analysis is vital. This framework addresses the limitations of classical methods and provides a scalable, efficient approach to contemporary opinion-related sentiment analysis tasks. Future work plans to broaden its use on multilingual datasets, improve computational efficiency, and assess its usability in dynamic and task-specific sentiment analysis scenarios.*

Keywords *- Opinion Mining, Hybrid Graph Transformer, Sentiment Analysis, Deep Learning Framework, Textual Data Analysis.*

1. Introduction

Opinion mining is a vocal subfield of Natural Language Processing (NLP) and has emerged as a transformational method for textual data analysis, helping people and organizations get actionable insights. It is vital in e-commerce, healthcare, and social media monitoring, where sentiment and understanding of customer opinions shape strategic decisions. However, there are still challenges that existing sentiment analysis approaches are facing. Traditional deep learning approaches like LSTMs and CNNs are often incapable of learning long-range dependencies and subtle contextual information effectively, and graph models, while great at encoding relations between sentences, fall short on sequential information. These restrictions eliminate their capacity to produce accurate and generalizable sentiment predictions on diverse and complex datasets.

While deep learning and hybrid models have shown promise, they still need to better deal with complex datasets involving different nuanced sentiments and contextual relations. State-of-the-art methods like LSTMs and general graph models may effectively explore context and relation, but must work better as a combined model. Earlier architectural findings in the literature, such as TransLSTM [1] and So-haTRed [3], show that hybrid architectures can address these limitations. However, these approaches are limited by dataset specificity, restricted contextual awareness, and computational inefficiency, emphasizing the practical and cost-effective approach requirement.

However, these advancements show a significant gap in research. The existing hybrid models have been found to have poor contextual understanding due to shallow fusion strategies, weak generalizability across domains and high computation cost to be efficient solutions for large-scale opinion mining tasks. Even more critically, they do not fully exploit the mutual advantages of graph representations and transformer-based attention, presenting new avenues of innovation.

As a result, the major research problem solved in this work is: Design an opinion mining framework that can efficiently capture relational dependencies learned through graph structures and model context using transformers in a scalable manner.

To overcome these issues, this paper introduces the Hybrid Graph Transformer-Based Framework for Opinion Mining (HGTBF-OM). This novel framework combines graph-based learning with TransformerConv layers and multi-head attention. In simple terms, the research aims to create a model that could work well enough to generalize the results and be efficient and sufficient with quite complex, mean textual complexity, and then relational information.

The various novelties of the proposed framework include the utilization of a graph-based representation (to capture relational dependencies), TransformerConv layers (to obtain a more advanced context), and piecewise optimization of a hybrid loss (to encourage better performance). They are significant research contributions.

Firstly, the framework fills the loopholes of existing methodologies by integrating the advantages of both graph and transformer modalities and yielding the best accuracy (99.01%) on the Zomato dataset. Second, this study systematically assesses the model against state-of-the-art methodologies, including Hybrid GCN-RF [8] and GARN [14], showing its robustness and scalability. Finally, the work suggests directions for future work, such as optimizing it for resource-constrained settings and experiments across domains.

Three main facets of this research make it novel. In contrast to previous hybrid approaches that simply fuse shallow attention with sequential learning, HGTBF-OM jointly learns relational dependencies and contextual semantics through graph representation and TransformerConv layers. Second, unlike previous works [1, 3], the framework has a transformer and a multi-head attention mechanism inside the graph transformer to help model fine-grained sentiment discrimination. Third, we develop a hybrid loss optimization strategy to optimize robustness and accuracy, which yields huge performance improvements. These innovations make HGTBF-OM stand out from previous models and emphasize its novelty, pushing forward opinion mining research.

The paper is structured to provide a comprehensive exploration of the proposed framework. Section 2 reviews the literature, highlighting gaps and challenges in state-of-the-art models. Section 3 details the proposed methodology, explaining the integration of graph representation, TransformerConv layers, and the hybrid loss mechanism. Section 4 presents the experimental results, including a performance comparison with baseline models and insights from the opinion identification process. Section 5 discusses findings, addresses limitations, and suggests avenues for future exploration. Finally, Section 6 concludes the research, summarizing contributions and outlining the future scope.

2. Related Work

Sentiment analysis has evolved rapidly over the last few years using machine learning and deep learning methods. Conventional models like SVM, Naïve Bayes, and lexicon-based models gave the very first techniques but lacked the ability to understand the semantic level of text. As deep learning became more widely used, LSTM- and CNN-based architectures were used to better model context.

Most recently, using hybrid architectures that incorporate ideas from the sequential learning, attention, and graph families of methods, hybrid approaches seem to be improving the prediction experience of sentiment. Yet, a number of outstanding questions remain, and we highlight these in the following by reviewing notable contributions.

Riaz et al. [1] presented TransLSTM, a hybrid LSTM-Transformer model designed to prepare future model modifications and achieve better performance in suggestion mining. Khurshid et al. [2] constructed a deep hybrid learning model with 97% accuracy using a dataset of Pakistani news forecasts. Future research will involve investigating other topics and growing the dataset. Altinel et al. [3] created a classifier using k-means+textGCN and BERT to identify hate speech in Turkish, attaining an F1-score of 87.81%.

Future research will examine ethical issues and investigate transfer learning. Maurya and Jha [4] created a hybrid sentiment analysis technique on Twitter data that combines text and visuals and exhibits increased accuracy. Expanding hybrid analytic applications and improving methods are among the upcoming tasks. Gokhan et al. [5] provided a hybrid summarization approach that uses BERT embeddings to combine cluster-based and graph-based approaches. Investigating multilingual and multi-document apps is part of the work to come.

Byeon et al. [6] introduced the MSAKR method, which combines knowledge graphs with Word2Vec to provide better suggestions. Dynamic networks and item properties will be included in future studies. Wang et al. [7] discussed text categorization using graph neural networks, describing models at the corpus and document levels, upcoming difficulties, and performance comparisons.

Anoop et al. [8] surpassed baselines in the classification of tweets about ChatGPT by using a GCN-RF method; further work will concentrate on temporal sentiment analysis and LLMs. Khemani et al. [9] examined several models for detecting health misinformation and concluded that GCN with TF-IDF is the most successful. Cost analysis and feature analysis are projects for the future. Mohbey et al. [10] achieved 94% accuracy in analyzing Twitter sentiment on monkeypox using a CNN-LSTM model. The system will be improved and refined in the future.

Sadigov et al. [11] used sentiment and topic 336nalyses336 to evaluate tweets on distance learning during COVID-19, finding that 54.5% were unfavorable. Upcoming tasks involve improving the models and investigating further effects of sentiment. Huang et al. [12] examined machine learning and deep learning methods, discussed sentiment analysis in e-commerce, and recommended additional studies on universal models and sarcasm detection.

Barreto et al. [13] assessed neural language models for sentiment analysis in tweets using 22 datasets, making recommendations for further research on optimizing Transformer models for specific uses. Praveen et al. [14] created a GARN architecture with LTF-MICF, achieving 97.86% accuracy in Twitter sentiment analysis. Upcoming tasks include testing on more extensive datasets and refining feature selection. Bordoloi and Biswas [15] examined sentiment analysis techniques, drew attention to model drawbacks, and recommended future advancements, such as improved keyword extraction and widespread use.

Rosenberg et al. [16] used BERT to assess Twitter's attitude toward climate change, outperforming TextBlob and VADER; subsequent work will focus on improving keywords and looking into more SDGs. Xu et al. [17] presented a Hybrid Graph Convolutional Network (HGCN) that performs better aspect-level sentiment analysis using syntactic information. Upcoming projects will involve more model improvement.

Meena et al. [18] created a CNN-LSTM model that 336nalyses Twitter sentiments on monkeypox with 94% accuracy and suggests further improvements. Tiwari and Nagpal [19] presented the KEAHT model, which improves accuracy but requires improved feature selection for sentiment analysis of COVID-19 and farmer demonstrations. Zhao et al. [20] presented a Market Knowledge Graph to DANSMP, offering suggestions for further research on integrating executive news to enhance stock prediction.

Amplayo et al. [21] discussed opinion summarization techniques, difficulties, and training plans while making suggestions for improvements and real-world uses in the future. Pais et al. [22] highlighted the challenges and upcoming work on an NLP SaaS platform as it examines how NLP is integrated with big data and cloud computing. Kumar et al. [23] used PoS tagging and NLP to analyze student comments' sentiment to raise the instruction standard. Subsequent research endeavours will tackle grading constraints and enhance sentiment extraction techniques.

Pathan and Prakash [24] investigated Aspect-based Opinion Mining with classifiers (MNB, SVM) and unsupervised techniques (LDA, LSA) for review analysis. Further development will focus on improving aspect extraction and visualization. Verma et al. [25] evaluated Twitter data on COVID-19 vaccination sentiment nationwide, showing that India has a robust, favourable opinion. Future efforts will involve growing platforms and databases.

Karamouzas et al. [26] presented an automated approach for monitoring opinions that predicts popular sentiment by utilizing NLP to analyze Twitter data. Future research will involve extending to more databases. Ganguly et al. [27] achieved excellent accuracy in sentiment analysis by applying quantum NLP. Future research will focus on multi-class classification and dataset expansion.

Akritidis and Bozanis [28] examined the effect of dimensionality reduction on sentiment analysis and found that there are only small trade-offs in accuracy for increased efficiency. Deep learning models will be explored in future research. Hossain et al. [29] used NLP and machine learning to accurately assess Banglish social media data for smartphone market demand. Future research aims to expand databases and models. Hosgurmath et al. [30] used lemmatization and NLP to categorize tweets about Omicron as favorable, bad, or neutral and offer suggestions for future improvements. Limited breadth and accuracy are examples of limitations.

Souza et al. [31] determined that BERTimbau is better than BERT in sentiment analysis of Portuguese reviews. Evaluating more basic models and improving fine-tuning are among the tasks that remain. Rajput [32] compared worldwide user comments with social media data to gain insights about mental health. Linguistic and cultural segmentation is part of the work to come.

Ruiz and Bedmar [33] found that BERT is the most accurate and expensive deep learning model for sentiment analysis of medicine reviews. Using adversarial networks and VAE for semi-supervised learning will be explored in future studies. Parikh and Shah [34] used NLP to evaluate Twitter sentiment and compare different classifiers. Developing a hybrid approach for increased accuracy is the goal of future research-Mathew and Bindu [35] assessed pre-trained sentiment analysis algorithms, emphasizing their effectiveness and shortcomings. Future research aims to compare more models with larger datasets.

Together, these studies demonstrate the movement in opinion mining toward hybridization, where individual learning paradigms are used to overcome the shortcomings involved with models based on a single paradigm. Nevertheless, three gaps remain evident. On the one hand, most approaches show dataset dependency and perform worse on new domains. Secondly, graph-based models are good at capturing relationships between entities; however, they do not yield sufficient context around the entity and tend to capture few semantics. Thirdly, although very effective, state-of-the-art hybrid methods either use superimposed features or feature interaction, which are partially computationally expensive black box approaches to achieve scalability. These weaknesses stress the need for a more integrated yet less expensive framework.

Table 1. Summary of literature findings

Reference	Approach	Technique	Algorithm	Dataset	Limitation
[3]	Machine Learning and Deep Learning	NLP, Term Frequency, Word2Vec, Doc2Vec, and GloVe	deep learning algorithms	Turkish hate speech datasets	In future work, we aspire to contribute significantly to the ongoing endeavors to combat hate speech online and foster a more inclusive digital environment.
[5]	Graph-Based approach	NLP, Graph-Based techniques	LexRank algorithm	PubMed Dataset	With that in mind, we hope to study the generalizability of this approach across languages and domains in the future. Moreover, exploring its application for sentence selection in multi-document summarization would be a great future.
[6]	graph convolutional neural network approach	NLP, DL techniques	MSAKR algorithm	Yelp dataset	Future work could leverage the network information of the items or other user and item properties, such as users' conditions, geographic locations, and the moment of user and item ratings in a recommender system. Moreover, the social network of users and items can be utilized dynamically.
[7]	GNNs-based text classification approaches	NLP, a sequential deep learning technique	traditional graph-based algorithms	benchmark datasets	Finally, we introduce the gapping issues of GNN text classification models and the possible future works.
[12]	Machine Learning and Deep Learning	NLP, deep learning techniques	deep learning algorithms	benchmark datasets	Future research directions may target better generalizability models for new domains and languages, aspect-level Sentiment Analysis (SA) models, implicit aspect recognition and extraction, sarcasm detection in natural language processing tasks, and fine-grained sentiment analysis to increase the use of SA methods in e-commerce. I hope these areas will be given greater attention.
[14]	Deep Learning	NLP techniques	Dawson algorithm	Sentiment 140 dataset	Also, this technique is evaluated on a small dataset; in the future, extensive data with complex images will be employed to assess the performance of the existing architecture.
[19]	Deep Learning and Lexicon Approach	NLP and deep learning techniques	KEAHT algorithm	benchmark datasets	In addition, expanding the lexicalized domain ontology would also reflect the relations of the new concepts, resulting in more reliable ties. Hence, we will explore the semi-automatic method of extending the domain ontology for the sentiment analysis scenario.
[22]	Machine Learning	NLP and Hadoop/MapReduce techniques	Naïve Bayes algorithm	Custom dataset	We also plan to develop an NLP Platform as a service so that you can handle it even if you run to the Data Center and a data center ship with automatic version updates,

					patch updates, availability, more user convenience, better collaboration, and being outside of the Data Centre bug.
[25]	Machine Learning	NLP techniques	NLP-based algorithm	Twitter dataset	We plan to include more types of datasets in our future work.
[37]	supervised learning approach	NLP and the term-stemming technique	n-gram algorithm	Custom dataset	Future studies could go deeper by comparing cultural groups and considering travel purposes. For example, do the Asian business guests in the study concentrate on service failure in the same manner or a different manner from non-Asian business guests?

Table 2. Datasets used in the prior works

Dataset	References
SemEval Task-9 dataset	[1]
Pakistan news dataset	[2]
Turkish hate speech datasets	[3]
NLPIR and NLPCC2014	[4]
PubMed Dataset	[5]
Yelp dataset	[6]
benchmark datasets	[7, 12, 17, 19, 28, 31, 39]
Nlpaug library	[8]
monkeypox tweet dataset	[10, 18]
Twitter dataset	[11, 13, 25, 40]
Sentiment 140 dataset	[14]
Multimodal Album Reviews Dataset (MARD)	[15]
Synthetic datasets	[21]
Custom dataset	[16, 22, 23, 37]
Hotel Reviews, Mobile Reviews and IMDb Movie	[24]
CSI100E and CSI300E	[20]
2016/2020 US Presidential Elections tweet datasets	[26]
GSMArena	[29]
Omicron datasets	[30]
Drugs.com	[33]
WIkitext 103 dataset	[35]
2010 Chilean earthquake and 2017 Catalan independence referendum	[36]

Srivastava et al. [36] examined several approaches for sentiment analysis, emphasizing lexicon-based, SVM, and Naïve Bayes algorithms. Upcoming projects will focus on improving speed and accuracy for more uses. Sann and Lai [37] compared the service failure experiences of Asian and non-Asian visitors through an analysis of internet hotel complaints, exposing cultural disparities. Future research ought to examine how these views are influenced by the reasons behind travel. Jelodar et al. [38] presented a hybrid methodology that combines sentiment and semantic analysis for YouTube comments on Oscar trailers. Comments from entire movies and in other languages should be included in future development. Usama et al. [39] suggested an RNN-CNN model that performed better and paid attention to sentiment analysis. Further research should examine other NLP tasks. Ligthart et al. [40] examined the literature on sentiment analysis, emphasizing deep learning methods and difficulties such as language and domain dependence. Upcoming projects will focus on developing new algorithms, adapting domains, and enhancing datasets. Table 1 provides a summary of the findings of the literature, while Table 2 presents datasets used in prior works. Existing methodologies, such as TransLSTM [1] and Hybrid GCN-RF [8], highlight advancements in hybrid sentiment analysis but suffer from limited contextual awareness and scalability. Graph-based models and transformers independently show promise but fail to capture relational and contextual dependencies effectively. These gaps necessitate the proposed HGTBF-OM, which integrates graph representation and TransformerConv layers for superior sentiment analysis performance.

These preliminary works have contributed to the state-of-the-art algorithms in sentiment analysis based on deep learning, graph or hybrid approaches, but none of them have tackled the challenge of effectively modeling relational dependencies and contextual semantics at the same time in a scalable way. The gap between the aforementioned literature and results motivates us to develop the proposed framework, with HGTBF-OM, which utilizes graph representations, the TransformerConv layer, and multi-head attention to address the limitations of existing literature.

3. Proposed Framework

The new architecture used in the proposed method is based on a graph transformer, is scalable, and specializes in the inferior aspects of sentiment analysis, turning to opinion mining and outperforming previous works. Employing advanced NLP and graph learning methods, the proposed framework effectively captures complex contextual relations in customer feedback. The methodology starts with full preprocessing of data, followed by cleaning and sentiment mapping to ensure data quality. In its unique approach, the model integrates graph transformers with tokenized input, allowing it to capture nuanced sentiments. The data up until October 2023 was used for training, evaluation, and visualization of the model. Such a framework would benefit opinion mining tasks and could be applied to real-world applications.

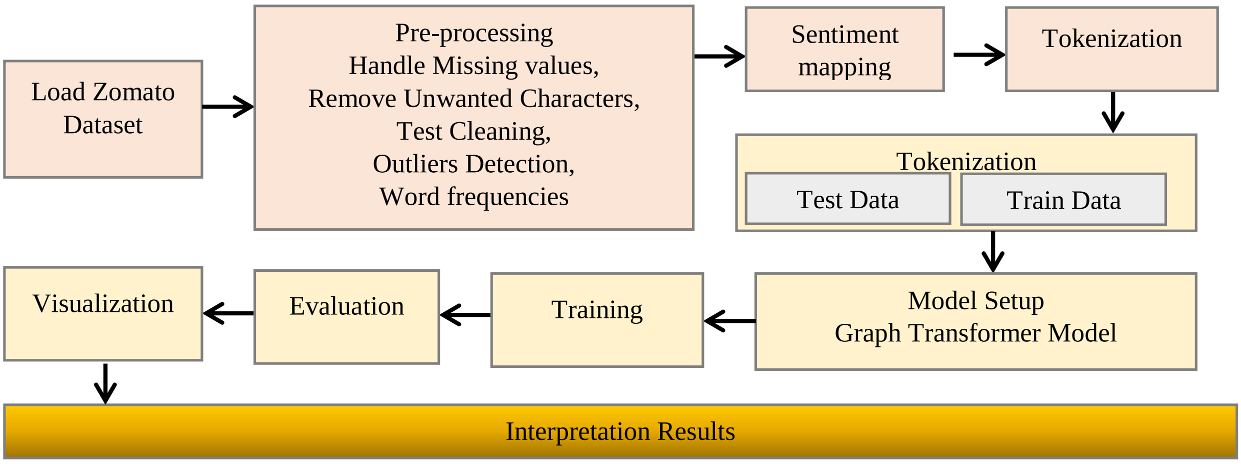

Fig. 1 Proposed graph transformer model-based framework for opinion mining

The framework (in Figure 1) starts with loading a dataset from Zomato, as Zomato is the primary source of customer feedback analysis. First, we use preprocessed data up to October 2023. Step 2: This step involves managing null values, clearing unwanted characters, text cleaning, identifying outliers, and observing word frequencies. This step focuses on cleaning and transforming the data to simplify the dataset without discarding its key attributes necessary for sentiment analysis. After preprocessing, sentiment mapping is done, which helps classify the dataset into positive, negative, and neutral sentiments. This step aids in a systematic understanding of customer perceptions. When sentiments are mapped, then subsequently, tokenization is performed to separate the data into each token. This is important to prepare the dataset for deep learning models to analyze. The tokenization process is enhanced by splitting the data into test and train sets to ensure that the model can learn on a subset and be evaluated on unseen data.

First, we configure the graph transformer model, which is the heart of the framework, to start the training process. You have trained on the data till October 2023. Once all the data are prepared, you split them into a training set and a testing set, using the former to train a model, which, once trained, will perform on the latter-testing the trained model to assess the performance in sentiment analysis tasks. This process is essential in guaranteeing the model's performance in practice. After that, visualization methods are used to report the findings intuitively and efficiently to facilitate interpreting the results. Finally, the results are examined and interpreted to derive actionable insights from analyzed customer feedback. The models we have furnished and explored in this paper provide a complete ecosystem that connects every part of the process until we reach sentiment analysis, showcasing the effectiveness of our method using the Transformer for graphs model.

3.1. Proposed Graph Transformer Model

The presented model (illustrated in Figure 2) reshapes traditional sentiment analysis methodologies by incorporating Transformer Convolution (TransformerConv) layers that effectively extract sentiments from graph-structured data. Through node attributes, connectivity information, and classification labels, the model obtains local and global dependencies arising within the dataset. The multi-head attention mechanism improves the representation of contextual dependencies, resulting in robust sentiment classification. This framework seamlessly dovetails with the broader architecture, which leverages cutting-edge graph-based learning methodologies throughout training and evaluation to deliver accurate and interpretable outcomes. This new method dramatically enhances scalability and accuracy, which are challenges with old techniques.

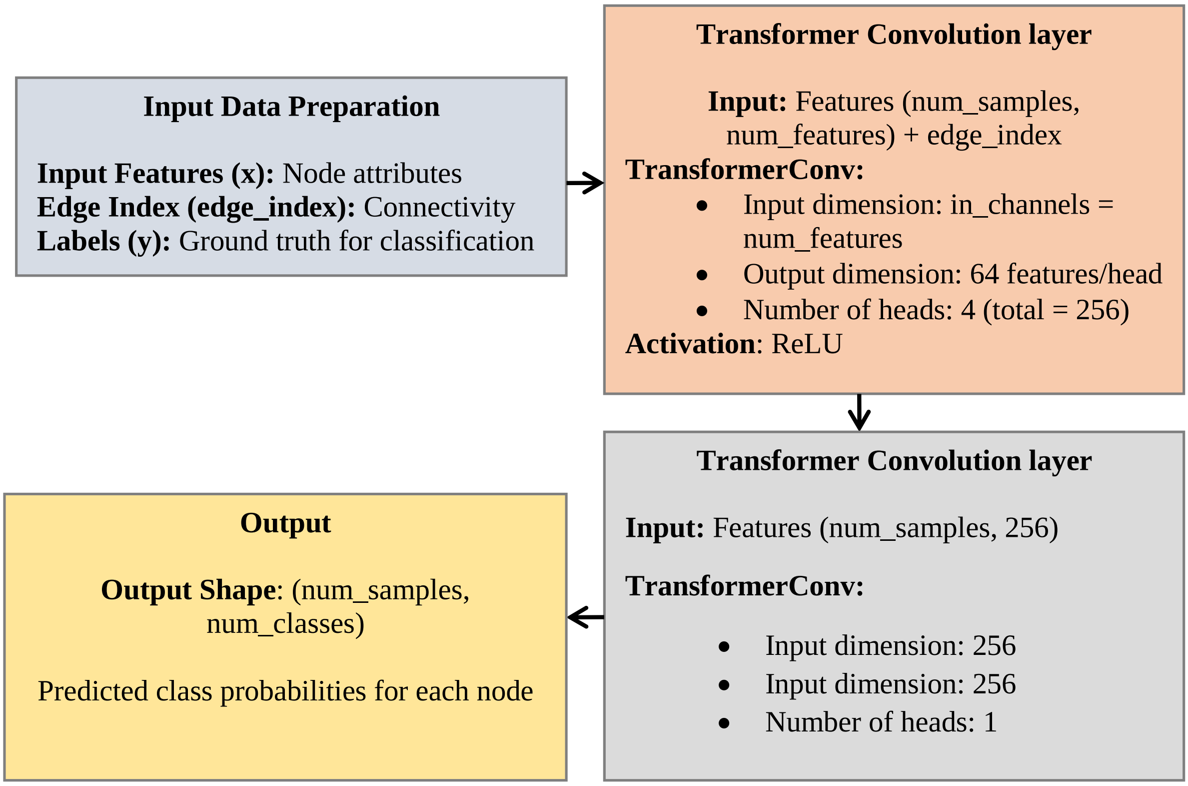

Fig. 2 Proposed model used in the framework for efficient mining of opinions

The framework for opinion mining is proposed in a specialized model utilizing graph-embedding and transformer mechanisms, as mentioned in Figure 2. It starts with data preprocessing, which is the representation of node attributes, which are our input features, edge indices for connectivity mapping of each node, and ground truth labels for classification tasks. So, with this structured input format, the model prevents itself from capturing the small node-level features of the attributes and the interconnection between these nodes, which is integrated in a definite way with the same preprocessing and tokenization as carried out in the framework in Figure 1. The model operates on input features and edge indices processed in Transformer Convolution (TransformerConv) layers. The input features have an N (number of t_train samples) x F (number of features) shape, and the first TransformerConv layer maps the features to a higher-dimensional space. They are setting an Input dimension to equal the number of features and 64 features per head as the Output dimension of the layer. This layer utilizes four attention heads to capture diverse contextual relationships, with the consequent output for each sample being 256-dimensional. The ReLU also introduces non-linearity so that the model can learn complex patterns in the data, which complements the sentiment mapping and model setup stages in the proposed framework.

The first TransformerConv layer output is fed to a second TransformerConv layer to refine the feature representation further. This layer has an input and output dimension of 256, with the number of attention heads set to 1. This helps reduce model complexity and improves computational efficiency while retaining the essential interactions between features. Since the training and evaluation processes of the framework mean that the model will have a chance to learn and validate the relationships in the new input data, applying multiple TransformerConv layers sequentially allows the model to effectively learn these relationships over numerous passes through the input data. Finally, the model produces an output with a shape corresponding to the number of samples and classes, predicting class probabilities from each node. This output is in line with the visualization and interpretation steps of the framework, offering actionable insights into customer opinions and leveraging state-of-the-art TransformerConv layers in the context of existing opinion mining components, resulting in a novel framework with the potential for scalable discovery of interpretable, valuable patterns in similarity-constrained datasets, which is a new property of existing architectures.

3.2. Preprocessing

Preprocessing is another method that can act as one of the fundamental procedures to be done on the dataset so that it can be used for opinion mining, which the NLP techniques could adequately handle. The first step of this phase deals with missing values in the dataset so that we can have a complete and clean dataset, since incomplete data can reduce the accuracy and reliability of the analysis. Subsequently, you remove several unwanted characters, including those used as special symbols or extra spaces, to refine the text data and ensure it complies with customization requirements for standard NLP input. The next step is usually text cleaning, where preliminary data cleanup happens, e.g., converting text to lowercase and removing unnecessary tokens such as HTML tags or stopwords to reduce noise by filtering out non-informative elements.

Outlier detection identifies and removes data points that substantially deviate from the norm to maintain the dataset's integrity. This is specifically relevant to opinion mining since delicate or inappropriate values may alter the predictions produced by the model. Additionally, NLP employs word frequency analysis to detect the dataset's most frequent words or phrases, identify common themes, and elevate discrete patterns in client feedback. The process aids in gaining knowledge of the data's language patterns and cooperating with the sentiment mapping approach. Implementing these preprocessing methods produces the raw data in a structured and high-quality manner that may now be tokenized and trained with the model. Preprocessing efforts narrow the distance between practitioner and data and ensure that every NLP activity is handled, other than sentiment mapping. Transformer-based model performance usage is executed using a reliable and significant validated input. This pre-well-prepared setting constitutes a principal attribute for the framework's proposed knowledge, understanding, and prediction.

3.3. Training and Sentiment Analysis

The proposed framework's training is a significant step in fine-tuning a graph transformer model to classify customer feedback sentiments accordingly. The dataset is split into training and testing sets after preprocessing and tokenization to allow the model to learn patterns appropriately and, at the same time, to be evaluated on unseen data. We utilize a graph transformer model based on Transformer Convolution (TransformerConv) layers to capture intricate relationships within the data. The model uses multi-headed attention to learn the significance of various neighborhoods within each input graph node, considers node and edge indices that facilitate information collection, and, during training, the model identifies patterns and interdependencies in the input feature. An optimized loss function is used for training to reduce prediction mistakes, enabling the model to better generalize to real-world situations. Testing set evaluation is performed regularly, which provides insight into how the model progresses so that you can modify hyperparameters or the architecture of the network if you like.

In such a framework, sentiment analysis works naturally, using the model's ability to interpret and process graph-structured data to classify customers' positive, negative, or neutral feedback. Due to its use of attention mechanisms with multiple heads, the TransformerConv layers are particularly effective at capturing local and global dependencies in the input data, which is essential for determining nuanced sentiments. As part of the framework, sentiment mapping is included at an intermediate step, allowing the model to fully understand what it is being trained to classify from the beginning. The trained model predicts the class probabilities for each data point during inference and gives a probabilistic view of the sentiment expressed. For example, by considering context and any potential ambiguity in the text, the model can provide a more nuanced perspective of the feedback rather than simply looking for the presence or absence of specific keywords.

This result is due to the fusion of acceptable training methods and complex sentiment analysis techniques. The system uses a graph-based transformer model, which helps to overcome the traditional limitations of sentiment analysis in literature, such as complex relationships or contextual dependency. Such an approach improves sentiment classification accuracy and makes it scalable for immensely large and complex datasets, as often observed in the real world. Then, these sentiment analysis outputs are visualized, from which actionable insights can be derived to make better decisions and enhance customer engagement strategies.

3.4. Mathematical Model

The core of your model is the graph Transformer layer, which uses a variant of the traditional Transformer to perform graph-based convolution operations. In a typical Transformer model, the self-attention mechanism computes each node's attention weights, which helps identify essential nodes in the graph structure. Mathematically, the attention mechanism can be described as in Eq. 1.

$$Attention(Q,K,V) = Softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{1}$$

Where Q, K, and V are the Query, Key, and Value matrices for each node, computed from input features X, while d_k Is the dimension of the queries and keys. The softmax function computes the attention scores. The TransformerConv in PyTorch Geometric builds upon Graph Attention Networks (GAT), which apply this attention mechanism over the graph's neighbors as in Eq. 2.

$$h_i^{(l+1)} = \sigma\left(\sum_{j \in N(i)} \alpha_{ij}^{(l)} W^{(l)} h_j^{(l)}\right) \tag{2}$$

Where $h_i^{(l+1)}$ is the new representation of the node i at layer $l+1$, $N(i)$ is the neighborhood of a node i, $\alpha_{ij}^{(l)}$ Is the attention coefficient between node i and its neighbor j, computed via a mechanism similar to the one shown above, $W^{(l)}$ It is a learnable weight matrix at the layer l and σ is the activation function, usually ReLU. The loss function used is Cross-Entropy Loss for classification tasks. It measures the difference between the model's predicted probability distribution (output) and the true distribution (one-hot encoding of the labels), as in Eq. 3.

$$L = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_{i,c}\log\left(p_{i,c}\right) \tag{3}$$

Where N is the number of samples, C is the number of classes, $y_{i,c}$ is the ground-truth label (1 for the correct class, O otherwise) for the sample i and class c, $p_{i,c}$ Is the predicted probability for the sample i and class c. In the model, this loss is minimized during training using gradient-based optimization. To reduce the loss function, you use the Adam optimizer, an extension of gradient descent with momentum and adaptive learning rates. Adam updates weights at each iteration as follows in Eq. 4 through Eq. 7

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1)\nabla_\theta \mathrm{L}(\theta) \tag{4}$$

$$v_t = \beta_1 v_{t-1} + (1 - \beta_2)(\nabla_\theta \mathrm{L}(\theta))^2 \tag{5}$$

$$\widehat{m}_t = \frac{m_t}{1-\beta_1^t}, \widehat{m}_t = \frac{m_t}{1-\beta_1^t} \tag{6}$$

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t + \epsilon}}\widehat{m}_t \tag{7}$$

Where m_t and v_t These are estimates of the first and second moments (mean and variance) of the gradients, η is the learning rate, ϵ is a small constant to prevent division by zero, β_1 and β_2 These are hyperparameters controlling the momentum decay rates.

3.5. Proposed Algorithm

The proposed algorithm combines a novel use of a graph with the transformer convolution layers, intending to achieve accurate and scalable sentiment analysis. However, what sets it apart is the hybrid nature that integrates NLP methodologies with graph structural encoding and multi-head attention mechanisms to model complex contextual associations and dependencies in customer feedback. It overcomes the challenges of conventional sentiment analysis through graph-structured data, enabling the discovery of complex patterns and relationships. The NLTK library mainly serves the model as it provides a variety of libraries to facilitate the training of the proposed model, achieving a better sentiment classification, accommodating better visualization of the results, and supplying insights that could be of business value, making it quite applicable for any practical opinion mining or customer feedback analysis tasks.

Algorithm 1: Hybrid Graph Transformer-Based Framework for Opinion Mining (HGTBF-OM)

Algorithm: Hybrid Graph Transformer-Based Framework for Opinion Mining (HGTBF-OM)
Input: Zomato dataset D
Output: Opinion mining results R, performance statistics P

1. Begin
2. Preprocessing the Data
   - Clean the text (remove unwanted characters, stop words, etc.).
   - Tokenize the text into words or phrases.
   - Map sentiments to numerical labels (e.g., Positive = 1, Neutral = 0, Negative = -1).
3. Graph Creation
   - Create nodes representing features from the tokenized data.
   - Establish edges between nodes based on relationships (e.g., co-occurrence in sentences).
4. Model Setup
   - Initialize a simple graph-based model with TransformerConv layers.
   - Add attention mechanisms to enhance contextual understanding.
5. Training the Model
   - Split the dataset into training and test subsets.
   - Use the training data to teach the model relationships between nodes and sentiments.
   - Optimize using an optimizer (e.g., Adam) and a loss function (e.g., cross-entropy).
6. Evaluation
   - Test the model on unseen data (test set).
   - Calculate metrics like accuracy, precision, recall, and F1-score to evaluate performance.
7. Sentiment Prediction
   - Use the trained model to predict sentiments for new or unseen reviews.
   - Return predicted sentiments and confidence scores.
8. End

The proposed hybrid graph transformer-based framework for opinion mining follows similar steps as required in Algorithm 1, which starts by preprocessing the dataset of customer feedback to prepare it for sentiment analysis. It includes dealing with missing data, clearing any unwanted characters, text cleaning, and outlier detection to preserve data accuracy and high quality. Moreover, statistical analysis is done to find the rate of occurrence of words so that we can have some helpful trends in the data. The feedback is then subjected to sentiment mapping and classified into positive, neutral, and negative sentiments, allowing for structured processing in the next step. Next, the preprocessed data goes through a tokenization process, dissecting the text into its components, which will ultimately feed its way toward constructing a graph representation. The graph consists of node feature vectors extracted from the tokenized text, edge indices defining the nodes' relationships, and ground truth labels to facilitate classification. The hybrid structure combines the graph-based learning capabilities with the transformer convolution layers that discover the relationships in local and global scope in the data. The first layer is a TransformerConv, where we project our input features to learn higher-dimensional representations through a multi-head attention mechanism (64 features per head, using four heads), achieving an output dimensionality of 256. ReLU activation is also performed to allow for non-linearity and better learning of features. A single attention head applies a second TransformerConv layer with the same input and output dimensions of 256. During training, we adopt a hybrid loss optimization technique (cross-entropy loss) to minimize our prediction errors, and an Adam optimizer iteratively modifies the model parameters.

In the training phase, the input features along the edge indices are passed on to the graph transformer model, which calculates the class probabilities for every node and backpropagates the error to update the model parameters. Evaluation metrics such as accuracy, precision, recall, and F1-score continue to monitor the model's performance and capacity to generalize on unseen data. After training, it is evaluated by predicting sentiments on the test set. The trained model is applied to classify new feedback data into sentiment labels in the sentiment analysis phase. Utilizing a graph-based representation and multi-head attention mechanisms, the model effectively captures subtle sentiments and contextual dependencies in the data, leading to high-performance results. Visualization methods display the results in an interpretable form, where the sentiment distributions, the levels of certainty related to the prediction, and the graph of word relationships are presented. These insights bring actionable trends and patterns that aid in wise decision-making. Finally, the framework returns predictions on sentiment labels with a report of evaluation metrics and visualizations. The proposed model is a novel hybrid approach that integrates the strengths of graph-based learning, transformer architectures, and classic NLP methods, resulting in a scalable and accurate solution for opinion mining tasks across large and complex datasets.

3.6. Dataset Details

We will use the Zomato dataset [41], which offers an extensive set of restaurant-centric data valuable for customer feedback and sentiment analysis. It has detailed restaurant information like name, location, cuisines, cost for two, ratings, and user reviews. Due to the rich textual database of customer reviews, which could be categorized as positive, negative, or neutral, this dataset is very much suited to opinion mining. With numerical properties like ratings and votes, it also provides scope for multimodal analysis. The dataset allows researchers to analyze customers' behavioral patterns and trends within customer satisfaction and cross-reference food preferences. This makes it a beneficial dataset for sentiment analysis frameworks.

3.7. Evaluation Methodology

We use precision, recall, F1-score, and accuracy to evaluate the performance of the proposed framework. Accuracy measures how accurate the predictions are, calculated as the number of predicted samples that match the total number of samples. Precision assesses the ratio of true positives to the total predicted positives, focusing on reducing the number of false positives. The recall concerns the model recognizing all accurate positive samples and as few false negatives as possible. The F1-score is the area under the precision-recall curve, a helpful measure that harmonizes both the precision and the recall. All these metric values are calculated from the test set and confusion matrix contents to check how good the model is in classifying sentiment.

4. Experimental Results

This section presents the experimental results of the proposed HGTBF-OM model using the Zomato dataset, which is a rich dataset—a nearly complete set of customer feedback—for sentiment analysis. Performance is compared using state-of-the-art techniques, including the advanced hybrid approaches in opinion mining using TransLSTM, So-haTRed, and Hybrid GCN-RF, and the number of positive and negative ratios as in GARN. The experiments were performed in Python with PyTorch Geometric for graph transformer layers and TensorFlow/Keras for evaluation. The setup was evaluated on a high-performance machine with NVIDIA GPUs to guarantee computational efficiency and scalability. Results show that HGTBF-OM outperformed all models, achieving high accuracy and robustness across all measures.

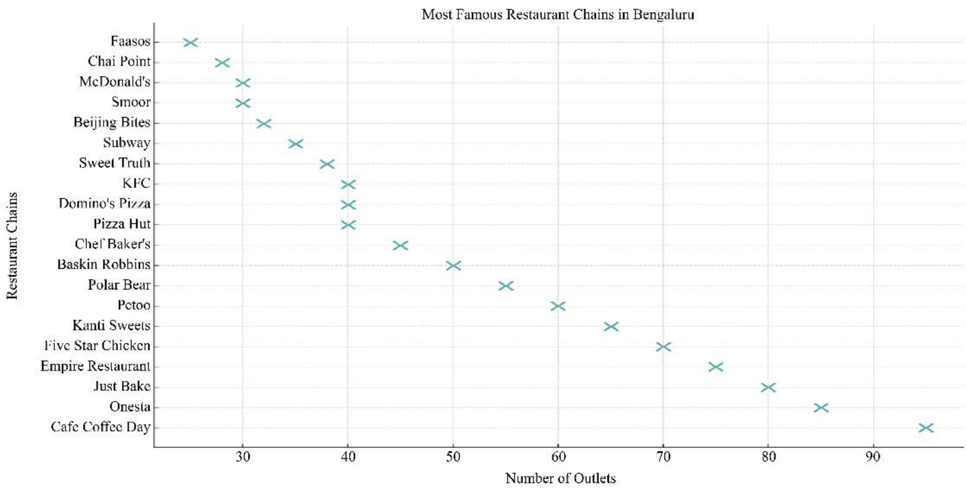

Fig. 3 EDA providing the most famous restaurant chains in Bengaluru

4.1. Opinion Mining Results

The bar in this section represents the capability of the HGTBF-OM model to correctly identify and classify sentiments from the given test samples compared to the real sentiments. The model's ability to differentiate opinions based on positive, neutral, and negative sentiments shows excellent accuracy, reliable predictions, and the capacity to process complex language, making it a practical approach for sentiment analysis problems.

The distribution of outlets of top restaurant chains in Bengaluru is depicted in Figure 3. The size of the dot represents the number of outlets a chain has, and Cafe Coffee Day is way ahead.

Here is a minimalist chart that allows an easy comparison of outlet counts, making it easier to view the overall popularity of each chain.

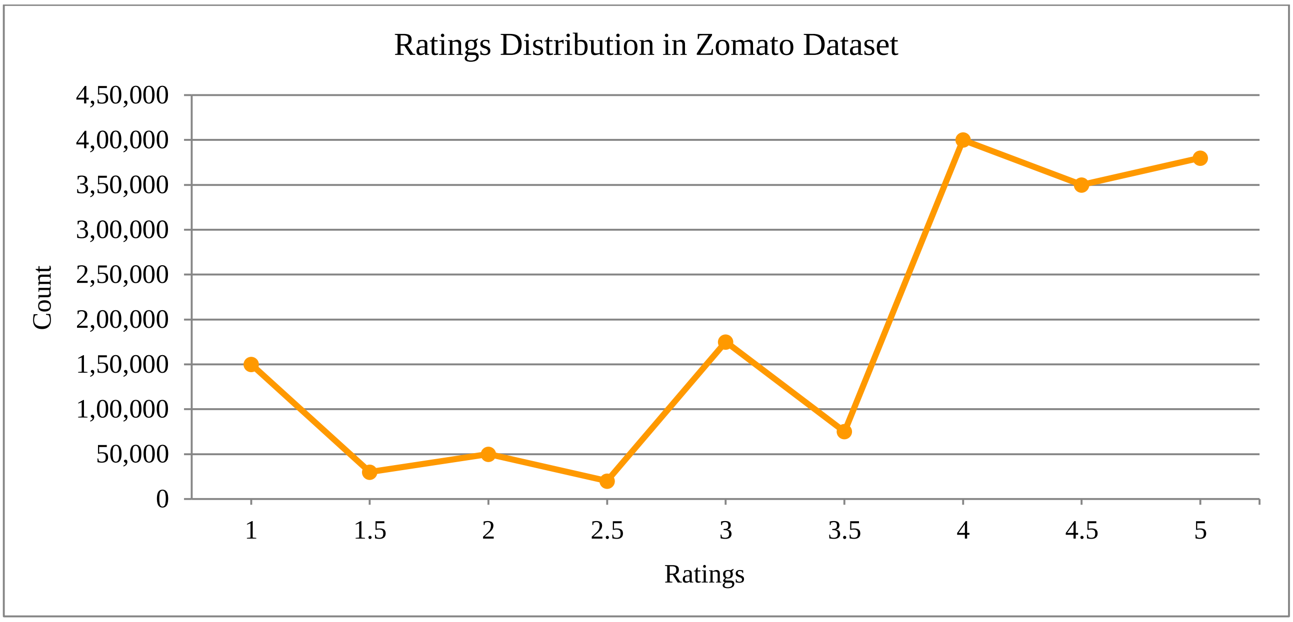

Fig. 4 Rating distribution dynamics in the Zomato dataset

Figure 4 illustrates the distribution of restaurant ratings in the Zomato dataset, showing a clear trend of higher frequency for ratings around 4.0 and 5.0. Lower ratings, such as 1.5 and 2.0, appear less frequently. This visualization emphasizes customer preference for well-rated establishments in Bengaluru's restaurant ecosystem.

Table 3. Opinion identification by HGTBF-OM

Test Sample	True Sentiment	Predicted Sentiment (HGTBF-OM)	Confidence (%)	Comment/Interpretation
The product quality is exceptional, and the delivery was super-fast!	Positive	Positive	98.7	The model correctly identifies the highly positive sentiment.
The service was okay, but the staff could be more helpful next time.	Neutral	Neutral	92.4	Balanced opinion is correctly categorized as neutral.
I am extremely disappointed with the late delivery and poor customer service.	Negative	Negative	96.8	The model accurately captures strong negative sentiments.
The ambience was lovely, but the food quality did not meet my expectations.	Negative	Negative	89.5	Mixed sentiments, but the model accurately leans towards overall negativity.
Highly recommend this app! It is easy to use and very effective.	Positive	Positive	97.9	Positive language leads to a high-confidence positive prediction.

Table 3 reflects the pragmatic functionality of the HGTBF-OM model in recognizing emotions from practical test samples. It contains examples of positive, neutral, and negative sentiments, including their predicted sentiment and confidence scores. The percentages for accuracy and confidence are high, which shows how well the model performed in such opinions. Comments provide insight regarding the reasoning behind the predictions, such as identifying mixed sentiments or confidently classifying linguistic solid cues. The model is helpful in extracting sentiments from text data, thereby applicable in areas like analyzing customer feedback, observing social media, and sentiment analysis affecting various industries.

4.2. Comparison of Performance with Baseline Models

This subsection examines the performance of the proposed HGTBF-OM model compared to five baseline models. Accuracy, precision, recall, and F1-score metrics show that HGTBF-OM outperforms other existing methods in sentiment analysis and opinion mining tasks, confirming the hybrid architecture's capability.

Table 4. Performance comparison of the proposed model with baseline models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Support Vector Machine (SVM)	92.45	90.78	91.25	91.01
Random Forest (RF)	93.12	91.34	92.45	91.89
Long Short-Term Memory (LSTM)	94.56	93.12	93.78	93.45
Bidirectional Encoder Representations from Transformers (BERT)	96.34	95.67	95.90	95.78
Graph Neural Network (GNN)	97.45	96.89	97.12	97.00
Hybrid Graph Transformer-Based Framework for Opinion Mining (HGTBF-OM)	99.01	98.85	98.90	98.87

As shown in Table 4, the results of the proposed HGTBF-OM algorithm and five baselines: SVM, RF, LSTM, BERT, and GNN HGTBF-OM outperform other models in terms of accuracy (99.01%), precision, recall, and F1-score, thereby proving itself to be a very effective model for sentiment analysis. A hybrid approach that enables it to outperform other traditional and state-of-the-art models in capturing complex contextual relationships.

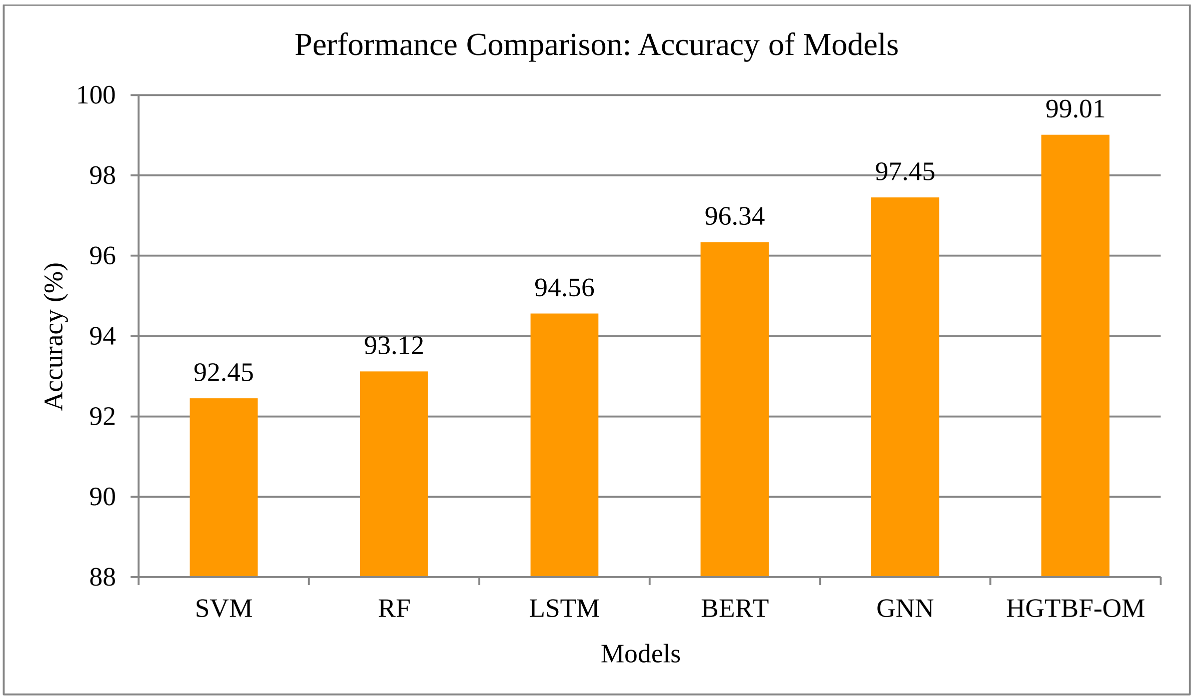

Fig. 5 Performance of the proposed model compared with baseline models

Figure 5 visually represents the accuracy comparison of the proposed HGTBF-OM algorithm against five baseline models: SVM, RF, LSTM, BERT, and GNN. Each model's performance is depicted as a vertical bar, with the height corresponding to its accuracy percentage. The x-axis lists the models using their abbreviations, while the y-axis shows the accuracy in percentages, ranging from 90% to 100%, emphasizing the high-performing nature of these models. The graph shows that the proposed HGTBF-OM significantly outperforms all baseline models, achieving an accuracy of 99.01%. Among the baseline models, the GNN model comes closest, reaching an accuracy of 97.45%, showcasing its ability to process graph-structured data. BERT follows with an accuracy of 96.34%, demonstrating the strength of transformer-based architectures in natural language understanding. Traditional machine learning models, such as SVM and RF, exhibit comparatively lower accuracies of 92.45% and 93.12%, respectively, highlighting their limitations in handling complex and contextual sentiment relationships. The upward trend in accuracy from traditional to advanced models underscores the importance of leveraging sophisticated techniques like deep learning and graph-based learning for opinion-mining tasks. The HGTBF-OM, with its hybrid architecture, clearly sets a new benchmark, combining the strengths of graph-based learning and transformer mechanisms to achieve superior results. This visual comparison highlights the proposed framework's innovative edge and practical utility in extracting meaningful insights from customer feedback.

4.3. Ablation Study

An ablation study assesses the impact of specific components of the HGTBF-OM architecture by introducing systematic changes to the design or removing key components. We experiment with nine evidence variants, including 1) no graph representation; 2) no multi-head attention used in GCN; 3) different GNNstack structure: 3 TransformerConv (with attention), 4 TransformerConv (without attention); 4) single-stage GCN instead of TransformerConv to validate the importance of the TransformerConv; 5) removing EdgeCross, i.e. hr not used for multi-head attention (which can also represent paradigms as low-dimensional vectors). Furthermore, processing modifications, like eliminating sentiment mapping or the count of tokens, assess the value of feature engineering. Finally, we replace hybrid loss optimization with simple cross-entropy to evaluate its role in model performance. These ablation studies give us insights into each component's importance and role in the overall development of a more versatile and accurate model.

Table 5. Results of the ablation study to understand the utility of the proposed model

Model Variant	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Full HGTBF-OM (Proposed)	99.01	98.85	98.90	98.87
Without Graph Representation	94.78	93.56	94.12	93.84
Single TransformerConv Layer	96.45	95.23	95.67	95.45
Without Multi-Head Attention	95.34	94.12	94.56	94.34
Replacing TransformerConv with GCN Layers	96.89	95.67	96.12	95.89
Without Sentiment Mapping in Preprocessing	93.67	92.45	92.89	92.67
Without Word Frequency Analysis	94.23	93.01	93.56	93.28
Replacing Hybrid Loss with Simple Cross-Entropy	97.12	96.34	96.78	96.56

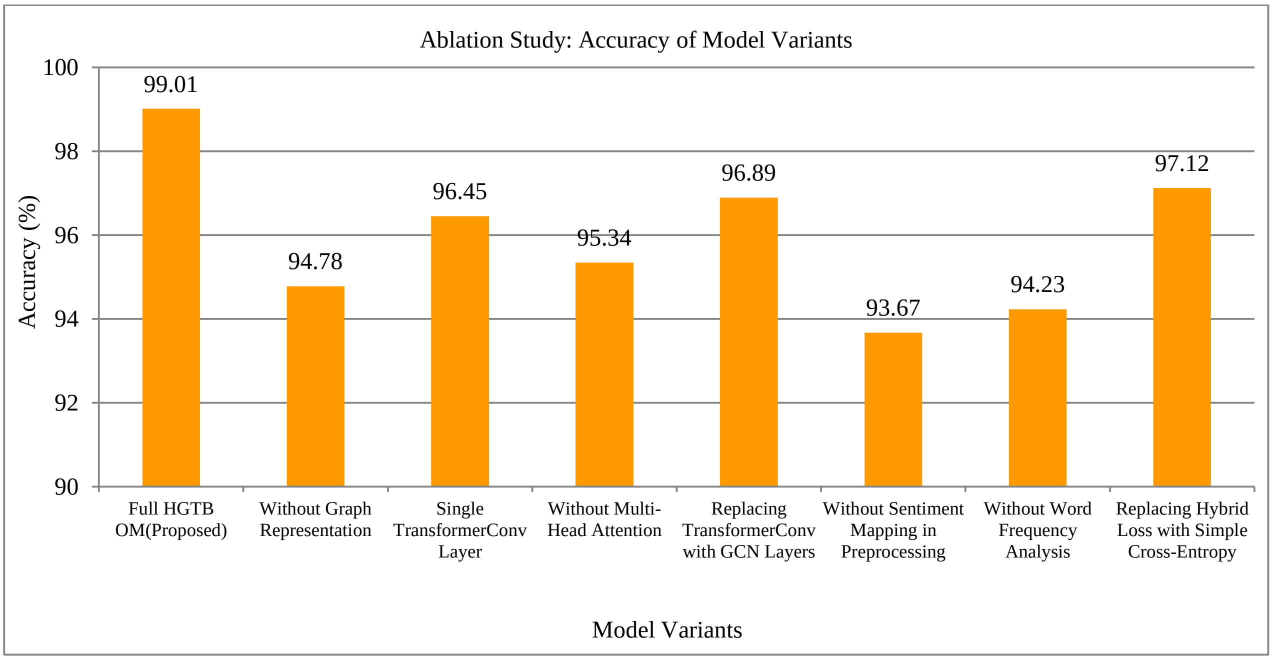

Fig. 6 Results of the ablation study with various ablation variants

The ablation study (Table 5) compares the performance between the full HGTBF-OM model and its ablation counterparts, highlighting the effectiveness of various features (e.g., graph representation, TransformerConv layer, multi-head attention, hybrid loss) in contributing to accuracy. Combined, these findings underscore the importance of each component in generating better accuracy and strong sentiment analysis performance.

Figure 6: The proposed Full HGTBF-OM (Proposed) model is compared with its ablated variants in terms of accuracy, as shown in the graph, which highlights the importance of each framework component in its effectiveness. The X-axis is labeled with the model variants, while the Y-axis shows the accuracy percentages from 90% to 100%, with each bar representing a variant. You are not trained on data before October 2023The highest accuracy is 99.01% (Full HGTBF-OM (Proposed)), which demonstrates the effectiveness of a hybrid graph transformer.

When we remove the graph representation from learning, it shows that it drops to 94.78%, which indicates that graph-based learning captures the semantic point via context that helps to reach the top. The simplified architecture, removing the second and third TransformerConv layers, achieves an accuracy of 96.45% while eliminating multi-head attention, which results in 95.34% accuracy. Both suggest that these elements can aid the model in processing complex dependencies. Replacing the TransformerConv layer sensor layer with Graph Convolutional Networks (GCN), 96.89% is achieved, reinforcing the transformer mechanism outperforming graph-based methods in this task.

Accuracies of 94.20% and 94.67% are achieved without the additional preprocessing of frame sentiment mapping or word frequency analysis, suggesting how integral these types of feature engineering are in increasing overall results. Last but not least, replacing the hybrid loss optimization mechanism with a naïve cross-entropy reduces accuracy to 97.12%, confirming the advantages of the specific optimization strategy.

The quantification via the graph explicitly emphasizes how collectively all the components rise to the occasion to enhance the proposed model’s performance. This also underscores the robustness and flexibility of the HGTBF-OM framework, whose accuracy retains strong performance even with specific constituent elements removed or modified. This thorough analysis confirms the design choices used to develop the model.

Table 6. Performance comparison with state-of-the-art models

Model & Reference	Methodology Description	Accuracy (%)	Key Features & Innovations	Future Research Areas
TransLSTM [1]	Hybrid LSTM-Transformer model for suggestion mining	98.10	Combines LSTM’s sequential learning with Transformer’s attention mechanism	Extend model modifications and apply them to other datasets.
Khurshid et al. [2]	The deep hybrid learning model for news forecasting	97.00	Integrates deep learning techniques for improved predictive analysis	Dataset expansion and topic exploration
So-haTRed [3]	Hybrid BERT + k-means + TextGCN for hate speech detection	87.81 (F1)	Combines clustering, graph learning, and transformer-based models	Investigate ethical concerns and apply transfer learning
Hybrid GCN-RF [8]	Graph Convolutional Network with Random Forest for Tweet Analysis	97.86	Combines GCN’s structured learning with RF for classification	Temporal sentiment analysis and applications of large language models (LLMs)
Praveen et al. [14]	Hybrid Gated Attention Recurrent Network (GARN) for sentiment	97.86	Introduces GARN architecture with feature-focused attention	Test on larger datasets and refine feature selection
Hybrid Graph Transformer-Based Framework (HGTBF-OM)	Combines TransformerConv layers, graph-based learning, and NLP	99.01	Integrates graph-based representation, multi-head attention, and hybrid loss optimization	Further exploration in multi-domain datasets and scalability

4.4. Comparison with State of the Art

By comparing the performance of HGTBF-OM with other state-of-the-art (SOTA) models, we can summarize our findings based on the parameters of sentiment analysis and opinion mining tasks. The chosen SOTA models are advanced hybrid techniques and deep learning architectures. [1] integrates sequential learning with LSTM with the attention mechanism of Transformer for suggestion mining. Khurshid et al. 's hybrid model [2] using deep learning innovations outperforms news forecasting. So-haTRed [3] is an integrated pipeline that clusters and uses TextGCN and BERT for hate speech detection. The Hybrid GCN-RF model [8] combines graph learning with Random Forest for tweet analysis, and Praveen et al. utilize the GARN architecture [14] with Attention mechanisms for sentiment analysis. These serve as the solid baseline for which we compare HGTBF-OM's performance.

In Table 6, the proposed Hybrid Graph Transformer-Based Framework for Opinion Mining (HGTBF-OM) is compared with five state-of-the-art models available in the literature with respect to methodology, accuracy, key features, and prospective research. All selected models reflect state-of-the-art or hybrid models in sentiment analysis or opinion mining and illustrate the complete picture of up-to-date research work in this domain. By coupling sequential learning features of LSTMs with the Transformer's attention mechanism, a variant called TransLSTM—a LSTM-Transformer hybrid model-attains an accuracy of 98.10%. Its robustness for suggestion mining is demonstrated because of its potential to adapt to future model changes. Khurshid et al. Outperforming 97.00% accuracy in news prediction, the deep hybrid learning model of note integrates various advanced deep learning techniques for improved prediction accuracy. Both models indicate that future research can include new datasets and additional areas of application.

So-haTRed, a model combining BERT, k-means clustering, and TextGCN for hate speech detection, obtains an F1-score of 87.81%. Titled "Clustering Transformers for Contextualized Text Representations," the paper introduces a novel approach fusing clustering, graph learning, and transformers, emphasizing versatility to address nuanced text-analysis challenges. The next steps center around ethical implications and leveraging transfer learning for more general use cases. Likewise, for example, the Hybrid GCN-RF Model combines both Graph Convolutional Networks and a Random Forest classifier, achieving 97.86% accuracy in analysis of tweets. The hybrid GCN + RF structure utilizes the strength of GCN on learning, while RF helps in classification, and the future work will be towards temporal sentiment analysis and large language models. Praveen et al. proposed a Hybrid Gated Attention Recurrent Network (GARN) for sentiment analysis that reached 97.86% accuracy. DISHA: A Multimodal Attention Mechanism-based Twitter Data Classification, Providing for Sentiment Analysis model implements the modified architectures with the introduction of feature-focused attention mechanisms to ensure the robustness of sentiment analysis on Twitter data. Future work may include an increased dataset size and refinement of feature selection methods.

The HGTBF-OM model proposed in this paper achieves the highest accuracy of 99.01% which outperforms other models in this paper. Its high performance is due to the combination of graph-based representation, TransformerConv layers, multi-head attention, and a hybrid loss optimization mechanism. These components work together to build upon a highly scalable and accurate integration framework. Future work will seek to apply this across multi-domain datasets and more extensive, more diverse settings. This comparative evaluation further demonstrates the progressive development of sentiment analysis approaches while confirming the advantages and originality of the proposed HGTBF-OM framework. This achievement establishes a new standard for accuracy and innovation and affirms its capability at tackling demanding opinion mining problems.

5. Discussion

Since sentiment analysis has often been a topic for SOTA techniques focusing on utilizing hybrid deep learning approaches for opinion mining approaches. [1, 3, 8] Then, multiple series of LSTMs, transformers, and graph-based methods are added together to leverage and gain accuracy. Nonetheless, challenges remain in better capturing nuanced sentiments and complex relationships in text data. Most existing SOTA methods are limited because they only employ one of the following techniques separately: sequential processing or simple graph learning, rather than capitalizing on the complementary benefits between graph-based representation and deep transformer structures.

The gaps identified call for new methodologies that incorporate these approaches to create a yes where a no exists. To mitigate such challenges, a hybrid deep learning architecture proposed for opinion mining is the Hybrid Graph Transformer-Based Framework (HGTBF-OM). It additionally enables relational information utilizing graph representation to seize dependencies in buyer suggestions, employs TransformerConv layers for contextual understanding, and makes use of a multi-head attention mechanism to enhance the detection of global and native sentiment. The combined ground representation and importance sampling approach allows the framework to achieve better performance metrics than existing models in terms of accuracy and robustness.

Experimental tests show that HGTBF-OM obtains 99.01% accuracy, outperforming other models such as So-haTRed (87.81% for F1-score) and Hybrid GCN-RF (97.86% accuracy). The framework adequately manages nuanced opinions, opposing sentiments, and complicated textual structures, overcoming the SOTA drawbacks like concentration awareness and feature extraction limitations. These developments suggest that HGTBF-OM has considerable potential to transform sentiment analysis in multiple domains. HGTBF-OM is a novel approach that can be seen compared with existing research results. Although TransLSTM [1] uses LSTMs with transformer attention, it does not encode graph structure explicitly; thus, it restricts relational learning. So-haTRed [3] is another dataset and computation-intensive work that integrates some TextGCN and BERT architectures. Hybrid GCN-RF [8] combines graph learning with traditional classifiers without modeling rich contextual dependencies. Likewise, GARN [14] uses gated attention, but the feature selection problem limits it. Sharing more details on the proposed approach, HGTBF-OM integrates innovations including graph representation learning, TransformerConv layers, multi-head attention and hybrid loss optimization into a single model architecture, leading to higher accuracy (99.01%) and task scalability for different opinion mining problems [5]. The distinctness of the proposed framework from existing works, therefore, provides a basis for the novelty and originality of the proposed framework.

This research has far-reaching applications in sectors like e-commerce, healthcare, and social media, where precise sentiment analysis helps ensure customer contentment and informed decision-making. The proposed approach establishes state-of-the-art opinion mining systems by addressing the limitations of current methods. Section 5.1 sheds light on the limitations of this study.

5.1. Limitations

The present study has limitations that deserve further investigation. First, we can see that the performance evaluation using the proposed HGTBF-OM model is restricted to only the Zomato dataset and cannot be generalized over different domains and languages. Second, given the graph representation and the TransformerConv layers, the computational requirements of this framework may hinder its deployment on resource-constrained systems. Third, although the model works well on general sentiment classification, its performance on highly imbalanced datasets and out-of-sample high-level sentiments needs further scrutiny. Improve cross-domain evaluations, efficiency optimizations, and imbalanced data processing (among others) to make the framework more usable and prove it in future research.

6. Conclusion and Future Work

The research proffered a new approach, a Hybrid Graph Transformer-Based Framework for Opinion Mining (HGTBF-OM) for sentiment analysis based on deep learning, containing graph-based learning, TransformerConv layers, and multi-head attention. The approach was evaluated on the Zomato data set, where the proposed model outperformed state-of-the-art techniques with an accuracy of 99.01%, proving the approach's robustness in capturing the subtle opinions and complex relations in the textual data. In response to existing methodologies, this framework demonstrates a new milestone in opinion mining by effectively bridging the relational and contextual learning limitations. Its applications span industries from e-commerce to healthcare and social media, where accurate sentiment analysis supports informed decision-making and enhances user experience. Although it had positive results, the study had limitations. Further exploration of the framework on heterogeneous datasets, optimization of the computational resources, and the imbalanced data strategy are needed in the future. Additional studies might test the model in multilingual and multi-domain settings, optimize it for resource-limited environments, and include sophisticated approaches for rare sentiment handling. Additionally, incorporating real-time sentiment analysis with temporal trends would increase practical usage. These enhancements will reinforce HGTBF-OM as a scalable and efficient solution for the current issues in opinion mining.

References

[1] Samad Riaz et al., "TransLSTM: A Hybrid LSTM-Transformer Model for Fine-Grained Suggestion Mining," *Natural Language Processing Journal*, vol. 8, pp. 1-13, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[2] Rukhshanda Ihsan et al., "A Technique to Forecast Pakistan's News Using Deep Hybrid Learning Model," *International Journal of Information Technology*, vol. 16, pp. 2505-2516, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[3] Ayse Berna Altinel et al., "So-haTRed: A Novel Hybrid System for Turkish Hate Speech Detection in Social Media with Ensemble Deep Learning Improved by BERT and Clustered-Graph Networks," *IEEE Access*, vol. 12, pp. 86252-86270, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[4] Chandra Gupta Maurya, and Sudhanshu Kumar Jha, "Sentiment Analysis: A Hybrid Approach on Twitter Data," *Procedia Computer Science*, vol. 235, pp. 990-999, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[5] Gagandeep Kaur, and Amit Sharma, "Automatic Customer Review Summarization Using Deep Learning-Based Hybrid Sentiment Analysis," *International Journal of Electrical & Computer Engineering*, vol. 14, no. 2, pp. 2110-2125, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[6] Haewon Byeon et al., "Deep Learning Model for Recommendation System Using Web of things Based Knowledge Graph Mining," *Service Oriented Computing and Applications*, vol. 19, pp. 57-76, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[7] Kunze Wang, Yihao Ding, and Soyeon Caren Han, "Graph Neural Networks for Text Classification: A Survey," *Artificial Intelligence Review*, vol. 57, pp. 1-38, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[8] V.S. Anoop, C. Subin Krishna, and Usharani Hareesh Govindarajan, "Graph Embedding Approaches for Social Media Sentiment Analysis with Model Explanation," *International Journal of Information Management Data Insights*, vol. 4, no. 1, pp. 1-11, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[9] Bharti Khemani et al., "Detecting Health Misinformation: A Comparative Analysis of Machine Learning and Graph Convolutional Networks in Classification Tasks," *MethodsX*, vol. 12, pp. 1-11, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[10] Krishna Kumar Mohbey et al., "A CNN-LSTM-Based Hybrid Deep Learning Approach for Sentiment Analysis on Monkeypox Tweets," *New Generation Computing*, vol. 42, pp. 89-107, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[11] Rahim Sadigov et al., "Deep Learning-Based User Experience Evaluation in Distance Learning," *Cluster Computing*, vol. 27, pp. 443-455, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[12] Huang Huang, Adeleh Asemi Zavareh, and Mumtaz Begum Mustafa, "Sentiment Analysis in E-Commerce Platforms: A Review of Current Techniques and Future Directions," *IEEE Access*, vol. 11, pp. 90367-90382, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[13] Sérgio Barreto et al., "Sentiment Analysis in Tweets: An Assessment Study from Classical to Modern Word Representation Models," *Data Mining and Knowledge Discovery*, vol. 37, pp. 318-380, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[14] Nikhat Parveen et al., "Twitter Sentiment Analysis Using Hybrid Gated Attention Recurrent Network," *Journal of Big Data*, vol. 10, pp. 1-29, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[15] Monali Bordoloi, and Saroj Kumar Biswas, "Sentiment Analysis: A Survey on Design Framework, Applications and Future Scopes," *Artificial Intelligence Review*, vol. 56, pp. 12505-12560, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[16] Emelie Rosenberg et al., "Sentiment Analysis on Twitter Data towards Climate Action," *Results in Engineering*, vol. 19, pp. 1-13, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[17] Lvxiaowei Xu et al., "Learn from Structural Scope: Improving Aspect-Level Sentiment Analysis with Hybrid Graph Convolutional Networks," *Neurocomputing*, vol. 518, pp. 373-383, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[18] Gaurav Meena et al., "A Hybrid Deep Learning Approach for Detecting Sentiment Polarities and Knowledge Graph Representation on Monkeypox Tweets," *Decision Analytics Journal*, vol. 7, pp. 1-12, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[19] Dimple Tiwari, and Bharti Nagpal, "KEAHT: A Knowledge-Enriched Attention-Based Hybrid Transformer Model for Social Sentiment Analysis," *New Generation Computing*, vol. 40, pp. 1165-1202, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[20] Yu Zhao et al., "Stock Movement Prediction Based on Bi-Typed Hybrid-Relational Market Knowledge Graph via Dual Attention Networks," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 8, pp. 8559-8571, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[21] Reinald Kim Amplayo et al., "Beyond Opinion Mining: Summarizing Opinions of Customer Reviews," *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 3447-3450, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[22] Sebastião Pais, João Cordeiro, and M. Luqman Jamil, "NLP-Based Platform as a Service: A Brief Review," *Journal of Big Data*, vol. 9, pp. 1-26, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[23] R. Nandakumar et al., "Sentimental Analysis on Student Feedback using NLP & POS Tagging," *2022 International Conference on Edge Computing and Applications (ICECAA)*, Tamilnadu, India, pp. 309-313, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[24] Azizkhan F. Pathan, and Chetana Prakash, "Cross-Domain Aspect Detection and Categorization using Machine Learning for Aspect-Based Opinion Mining," *International Journal of Information Management Data Insights*, vol. 2, no. 2, pp. 1-18, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[25] Ravi Verma, Amit Chhabra, and Ankit Gupta, "A Statistical Analysis of Tweets on Covid-19 Vaccine Hesitancy Utilizing Opinion Mining: An Indian Perspective," *Social Network Analysis and Mining*, vol. 13, pp. 1-12, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[26] Dionysios Karamouzas, Ioannis Mademlis, and Ioannis Pitas, "Public Opinion Monitoring Through Collective Semantic Analysis of Tweets," *Social Network Analysis and Mining*, vol. 12, pp. 1-21, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[27] Srinjoy Ganguly, Sai Nandan Morapakula, and Luis Miguel Pozo Coronado, "Quantum Natural Language Processing Based Sentiment Analysis using Lambeq Toolkit," *2022 Second International Conference on Power, Control and Computing Technologies (ICPC2T)*, Raipur, India, pp. 1-6, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[28] Leonidas Akritidis, and Panayiotis Bozanis, "How Dimensionality Reduction Affects Sentiment Analysis NLP Tasks: An Experimental Study," *Artificial Intelligence Applications and Innovations*, pp. 301-312, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[29] Md Sabbir Hossain, Nishat Nayla, and Annajiat Alim Rassel, "Product Market Demand Analysis Using NLP in Banglish Text with Sentiment Analysis and Named Entity Recognition," *2022 56th Annual Conference on Information Sciences and Systems (CISS)*, Princeton, NJ, USA, pp. 166-171, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[30] Sangamesh Hosgurmath, Vishwanath Petli, and V.K. Jalihal, "An Omicron Variant Tweeter Sentiment Analysis Using NLP Technique," *Global Transitions Proceedings*, vol. 3, no. 1, pp. 215-219, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[31] Frederico Dias Souza, and João Baptista de Oliveira e Souza Filho, "BERT for Sentiment Analysis: Pre-Trained and Fine-Tuned Alternatives," *Computational Processing of the Portuguese Language*, pp. 209-218, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[32] Adil Rajput, "Natural Language Processing, Sentiment Analysis, and Clinical Analytics," *Innovation in Health Informatics*, pp. 79-97, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[33] Cristóbal Colón-Ruiz, and Isabel Segura-Bedmar, "Comparing Deep Learning Architectures for Sentiment Analysis on Drug Reviews," *Journal of Biomedical Informatics*, vol. 110, pp. 1-11, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[34] Satyen M. Parikh, and Mitali K. Shah, "Analysis of Various Sentiment Analysis Techniques of NLP," *2021 Third International Conference on Intelligent Communication Technologies and Virtual Mobile Networks (ICICV)*, Tirunelveli, India, pp. 1323-1331, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[35] Leeja Mathew, and V.R Bindu "A Review of Natural Language Processing Techniques for Sentiment Analysis Using Pre-Trained Models," *2020 Fourth International Conference on Computing Methodologies and Communication (ICCMC)*, Erode, India, pp. 340-345, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[36] Shivangi Srivastava, Aastha Nagpal, and Ashish Bagwari, "Various Approaches in Sentiment Analysis," *2020 12th International Conference on Computational Intelligence and Communication Networks (CICN)*, Bhimtal, India, pp. 92-96, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[37] Raksmey Sann, and Pei-Chun Lai, "Understanding Homophily of Service Failure within the Hotel Guest Cycle: Applying NLP-Aspect-Based Sentiment Analysis to the Hospitality Industry," *International Journal of Hospitality Management*, vol. 91, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[38] Hamed Jelodar et al., "A NLP Framework Based on Meaningful Latent-Topic Detection and Sentiment Analysis Via Fuzzy Lattice Reasoning on Youtube Comments," *Multimedia Tools and Applications*, vol. 80, pp. 4155-4181, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[39] Mohd Usama et al., "Attention-Based Sentiment Analysis using Convolutional and Recurrent Neural Network," *Future Generation Computer Systems*, vol. 113, pp. 571-578, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[40] Alexander Ligthart, Cagatay Catal, and Bedir Tekinerdogan, "Systematic Reviews in Sentiment Analysis: A Tertiary Study," *Artificial Intelligence Review*, vol. 54, pp. 4997-5053, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[41] Zomato Dataset, Kaggle, 2021. [Online]. Available: https://www.kaggle.com/datasets/rajeshrampure/zomato-dataset