SERBIAN JOURNAL OF ELECTRICAL ENGINEERING
Vol. 23, No. 2, June 2026, 257-284
UDC: 004.8:81`322.2 **DOI: https://doi.org/10.2298/SJEE2602257T**
Original scientific paper

A Unified Transformer-BiLSTM and Graph Attention Network Framework for Explainable Multilingual Opinion Mining and Relationship Inference in Complex Social-Citation Networks

Manoharan Thangavel[1], Alagar Muniandi Kalpana[1], Saravanan Ananth[2]

Abstract: The paper takes into consideration the rising requirement for accurate mining of opinion and inferring relations among entities as the amount of multi-lingual online information increases rapidly. Thus, the paper seeks a uniform statistical learning methodology for processing multiple languages and exploiting valuable relational discoveries. The paper introduces an innovative model which combines transformer-based context embedding, BiLSTM for capturing of sentiment flows, and GAT for examining relational data. Transformers model cross-lingual context probability distribution, BiLSTM models temporal transitions in sentiment, whereas GAT offers an attention-based statistical weighting on data structured as a graph. Examples include social networks and citation graphs. The model yields a sentiment classification accuracy and macro F1-score of 92.4% over the four multilingual datasets (English, Spanish, and Hindi), surpasses all baseline methods. The model has achieved state-of-the-art result 86.4% in relation inference with graph attention mechanism, and shows excellent reliability and generalization. In this work, we propose a statistically sound and integrated framework which incorporates contextual, sequential and relational modelling of multilingual opinion mining. Graph structured learning integrated with probabilistic encoding and temporal modeling is a novel technique for such tasks, which also shows a promising direction in terms of scalability and generalizability across different domains.

Keywords: Multilingual opinion mining, Graph attention network, Explainable AI, Relationship inference, Sentiment analysis, Cross-lingual modeling, Natural language processing.

[1]Department of Computer Science, Government Arts College (Autonomous), Salem, Tamil Nadu, India
thangavel.gacslmm@gmail.com, https://orcid.org/0009-0003-0990-2502
kalpana.gce@gmail.com, https://orcid.org/0000-0002-0675-6285
[2]Department of Artificial Intelligence & Data Science, Mahendra Engineering College (Autonomous), Namakkal, Tamil Nadu, India, anand.siet@gmail.com, https://orcid.org/0000-0001-5446-2268

Colour versions of the one or more of the figures in this paper are available online at https://sjee.ftn.kg.ac.rs

1 Introduction

Now that there is a massive amount of user-generated content on digital media platforms, including social media, e-commerce sites, and academic forums to specifically mine opinions from textual data has become a core function for businesses, research, and policy making. Opinion mining, or sentiment analysis, seeks to mine subjective information from textual sources; this includes opinions, attitudes, and emotions. There has been substantial progress in opinion mining in a monolingual environment. However, the notion of multilingual opinion mining has a new set of complexities because of differences in syntax, semantics, cultural semantics, and different resource availability among languages. Multilingual opinion mining (MOM) has emerged as a separate process for analyzing expressed opinions in multiple languages, particularly because of the rising globalized space where organized markets engage in cross-border brand and product monitoring, policymakers analyze international policy and evaluation effectively, and scholars or researchers assess multilingual, multi-site social media monitoring. Traditional machine learning, or early algorithms involving deep learning models, are naive and often rely on language-specific materials or utilize parallel corpora; hence why the previous attempts proposed via MOM frameworks are limited in potential scalability for less resourced languages [1]. These restrictive models often test and classify continuously during analysis and struggle to provide stability in the model when classifying opinion across languages and specific cultural expressivity, including idiomatic statements.

Recent progress in natural language processing (NLP), including the introduction of Transformer-based architectures like BERT and XLM-R, has advanced multilingual representation learning based on large corpora and shared embedding spaces [2, 3]. However, these models still use separate fine-tuning for each task and are still limited in adapting to long-range dependencies or intricately related text structures since these models operate on singular embedded representations. In this context, the use of Bidirectional Long Short-Term Memory (BiLSTM) networks can be combined with transformers to offer robustness across variations of text contexts [4]. Further, in instances related to academic contexts like citation networks or discussions on social media, it is just as important to obtain sentiments towards the specific entities as it is to understand the relational structures. Graph Attention Networks (GAT) provide the ability to leverage relational structures systematically while weighting the importance of neighboring nodes of the graph [5]. Overall, the motivations for combining these approaches through an overarching framework provides ambiguity and promise for large-scale, explainable, and contextual multilingual opinion mining.

In spite of the progress made in opinion mining, conventional methods such as rule-based systems and classical machine learning models continue to have

trouble generalizing across languages, domains, and contexts. These models typically rely on language-specific handcrafted features (e.g. sentiment lexicons, part-of-speech tags, syntactic rules) that are expensive to produce [6]. Thus, if these approaches are extended to multilingual contexts, significant linguistic knowledge and labeled data for each target language are necessary. Even early deep learning models like Convolutional Neural Networks (CNNs) and simple Recurrent Neural Networks (RNNs), struggled to capture the complex syntactic details and long-distance dependencies of multilingual corpora [7]. They usually also fail to leverage and explicitly model relations and external knowledge, and therefore do not easily allow us to model opinion propagation and other contextual shifts in dialog and social networks.

Taking these limitations into consideration, there is an increasing recognition that integrated and flexible frameworks capable of accommodating multilingual inputs are needed, retaining context despite many different syntactic structures in the languages, and modeling relationships among textual entities (or users). This could take the form of a unified architecture based on the distinguishing features of the language modeling capabilities of Transformers, the contextual sequencing of BiLSTM networks, and the relational reasoning capabilities of a Graph Attention Network (GAT). This model would establish a more seamless solution to the constraints or limitations of current frameworks.

1.1 Key contributions of the research

Unlike existing approaches that isolate contextual encoding, sequential modeling, or relational reasoning, this study introduces a unified framework that leverages the unique architectural synergy between Transformers, BiLSTM, and GATs. The contribution lies not merely in the integration, but in the formulation of a multi-level feature fusion that addresses the specific limitations of monolithic models: while Transformers provide global context, they are often insufficient for modeling temporal sentiment dynamics and local node-to-node relational dependencies. By coupling these layers, our framework creates a hierarchical representation that maps global semantics (Transformer), sequential flow (BiLSTM), and structural connectivity (GAT) into a singular latent space, a synthesis that provides a demonstrable improvement in both classification precision and relationship inference tasks.

The outline of the paper chapter-wise is as follows. Section 2 is a review of the related literature, while the purpose of Section 3 is to give a brief view of the theoretical framework, key concepts along with methodologies. Section 4 is going to evaluate the experimental result. Section 5 contains results and discussions, whereas Section 6 wraps it all together with a summary of the most important findings and suggestions for further research.

2 Literature Review

Ahmadnia et al. [6] tested the pre-trained large language models with a close interest in BERT and T5 to determine various dimensions of private states in a text, such as type, polarity, intensity, expression, and source. Their model showed that Transformer-based architectures are powerful in the ability to encode subtle contextual representation and hierarchical semantic characteristics along sentiment dimensions. Although the research demonstrates the success of contextual embeddings in the model of complex emotional states, it largely uses the representation of sentences, and does not directly deal with sequential sentiment flow or ties between textual objects, which are ubiquitous in social and citation-based opinion information.

Anggrainingsih et al. [7] provided a review of Transformer-based models on rumour detector and opinion-rich content on microblogging platforms. The authors stressed that BERT and its variants are appropriate in order to process short, ambiguous, and context-sensitive social media text. Even though their review highlights the flexibility of Transformer models to real-time sentiment mining and stance detection, it also indicates an overall weakness of Transformer-focused models, which is the absence of a clear temporal sentiment dynamics and inter-entity relationships modeling of conversational or networked data.

Zhou et al. [8] investigated BERT-based hybrid with BiLSTM as sentiment analysis in customer reviews and showed that the ability of global contextual representations and sequential modeling enhances precision, recall, and F1-score. They find that Transformer models, as well as BiLSTM models, are complementary. Nevertheless, the given approach is limited to the classification of flat texts and is not applied to structured relational conditions, including social or citation networks, where opinion interactions are graph-based by definition.

The article by Shah et al. [9] presents a broad survey on sentiment analysis in educational data, emphasizing that BiLSTM models could be very helpful to capture sequential and contextual dependencies in student feedback. Although it was demonstrated that BiLSTM architectures can be effective at modelling sentiment flow in domain-specific datasets, because of their limited capacity to encode global contextual semantics across languages, they are limited in scaling to multilingual and cross-domain opinion mining problems.

Hossain et al. [10] proposed a sentiment analysis end-to-end mode called SentiLSTM that uses the Long Short-Term Memory networks to identify sequential dependencies within customer restaurant reviews. The model is useful in contextual sentiment learning based on textual sequences and achieves high classification accuracy as compared to standard machine learning bottlenecks. SentiLSTM is, however, more focused on sentiment classification at the document level, and does not support contextualized embeddings, graph-based relational modelling, or multi-language support and as such its applicability

would only be relevant in complex opinion mining problems with interrelated entities and multi-lingual data.

Wu et al. [11] came up with a phrase dependency relational graph attention network which incorporated the syntactic dependency relations into a graph attention system to perform aspect-based sentiment analysis. The approach by explicitly modeling phrase-level dependency structures can improve the detection of sentiment-carrying phrases, based on particular aspects, and makes significant improvements on benchmark datasets. Although highly effective in identifying finer grained syntactic relationships, the method treats only the monolingual, aspect-level sentiment classification task and is not optimized to multilingual opinion mining, or simultaneous sentiment and relational inference of multiple opinion graphs of heterogeneous semantic nets.

According to Ji.ang et al. [12], the dependency relation-weighted graph attention framework of aspect-based sentiment analysis is proposed, where the syntactic dependency information is directly added to the attention mechanism to improve the contextual representation learning. The model uses dependency relations with differentiated weights to enhance the determination of words that carry sentiment with respect to specific aspects, which results in better classification. Although the approach is effective in the context of the fine-grained syntactic dependency, it is limited to the sentiment polarity prediction on aspects at the monolingual environment, and lacks in the consideration of the multilingual opinion analysis or both sentiment and relationship within complex interaction networks.

Chakraborty [13] discussed an end-to-end aspect-based sentiment analysis system based on Graph Attention Networks (GATTs) to learn syntactic and semantic relationships between aspect terms and the contextual tokens. The model is successful in the sense that it will capture long-range dependencies and enhance sentiment polarity classification on the aspect level by building attention-driven graph representations. Even though this method is showing good performance in structured sentiment representation, it is only focused on monolingual aspect-based sentiment classification but it does not extend to multilingual opinion mining and joint modeling of sentiment and relational interactions on heterogeneous opinion networks.

Yang et al. [14] also added to Graph Attention Networks adversarial training to enhance robustness during text classification, such as opinion mining. In their work, emphasis is placed upon the role of robustness and generalization in noisy textual environments. Nevertheless, the model is mostly aimed at stability of classification, as opposed to modeling contextual, sequential, and relational information in a single context.

While the presented papers have introduced breakthroughs in their respective areas contextual encoding through Transformers, sequential sentiment modeling through BiLSTM, and relational mapping through GATs there is a notable absence of research that integrates these components. The overwhelming majority of existing literature consider these modules as separate pipelines and fail to capture structural relational dependencies, either prioritizing deep contextual understanding by ignoring the sequential nature of sentiment or compromising its sequence modeling aspect for richer context-aware word representations. Similarly, existing multilingual benchmarks remain segmented without a representation that can account for how an opinion cascades through an elaborate citation network [15]. Here we bridge these distinct threads by combining the presented approaches into a cohesive model. Through the joint integration of these components, we show how deeper relational reasoning can address the limitations of fully transformer-based architectures by offering a more comprehensive model for multilingual opinion mining than has so far been proposed.

3 Methodology

3.1 Architecture of the Unified Transformer-BiLSTM and Graph Attention Network Framework

The Unified Transformer-BiLSTM and Graph Attention Network (GAT) Framework merges three very powerful deep learning models to achieve complex levels of performance for multilingual opinion mining. The overall structure of the framework is to successfully encode both contextual information of the text, as well as the relationships between different components (aspects vs opinions), which is extremely useful in scenarios that deal with sentiment analysis and aspect-based sentiment classification. The unified transformer-biLSTM and GAT framework consists of the following parts:

3.1.1 Transformer Model (Encoder Layer)

The framework begins with the transformer model (specifically the encoder layer), the model that brought self-attention to the public consciousness, where self-attention can encode the relationships between words, regardless of the distance between them, in the text sequence. The input text is tokenized and passed on through layers of the transformer model resulting in an output of contextualized representations (embeddings) of each token. Positional encoding is added to the embeddings to tell the model which order the word was presented in the input sequence, since transformer models are otherwise position agnostic. The transformer yields very rich representations, preserving lots of important semantic information, which is especially necessary when dealing with the complexity of multilingual data.

3.1.2 BiLSTM (Bidirectional Long Short-Term Memory)

The output of the Transformer encoder is passed to a Bi-directional LSTM layer, which is helpful for framing sequential interdependencies in the text. While a traditional LSTM processes the text in one direction (e.g., left to right), the BiLSTM considers the text in both a left-to-right and right-to-left direction. This is useful because it allows the model to leverage the-context dependency of a word from both the preceding context and succeeding context. In an opinion mining task, being able to deliver an outcome that is based on the preceding relationship of words, as well as on the succeeding words, enables the model to develop a detailed prediction of sentiments and relationships within the text. In this way, the BiLSTM layer refines the contextualized embeddings even further, enhancing their use for downstream tasks.

3.1.3 Graph Attention Network (GAT)

The third component of the architecture is the Graph Attention Network (GAT). The GAT will model the relations amongst the areas of the text input (i.e., aspects and opinions). In the case of opinion mining, aspects refer to features or properties (e.g., screen, camera, user experience) of the product or service, while opinions refer to sentiments/attitudes or judgments about those features (e.g., bad camera, bad screen, bad user experience).The GAT will graph these relationships through aspects and opinions nodes, aspects and opinions edges (dependencies), and attention scores assigned to any neighboring nodes.

The GAT is different than any state-of-the-art GNN models, namely it permits aspects and opinions to have dissimilar importance scores (attention scores) to determine the relevant aspects and opinions to predict sentiment, Therefore, the use of a graph representation and attention scores will demonstrate and improve the prediction of sentiment toward particular aspects in the text, which is required for aspect-based sentiment analysis. Unified Integration Layer: After the separate models (Transformer, BiLSTM, GAT) have taken the input in the text format, the separate outputs are combined in the Unified Integration Layer. The Unified Integration Layer integrates the different features developed by each of the models to create one combined representation of the text. The final combined features are described in a fully connected layer that makes the predictions that the model makes. The predictions are either the sentiment label (positive, negative, neutral), aspect category label, or any other task-specific labels.

Table 1 contains the hyperparameters of the architecture. This proposed architecture used standard settings that are proven to be robust in cross-lingual tasks. Transformer encoder was initialized with pre-trained [e.g. XLM-RoBERTa-base] model with embedding size 768. The BiLSTM was used with 2 hidden layers, hidden size 256 and dropout rate 0.3 in order to combat with overfitting. In GAT layer, we used attention over multiple heads, specifically 8

heads, with 128 as hidden layer dimension. Training was conducted by AdamW optimizer with $2e^{-5}$ learning rate, batch size 32 and a gradient clipping of 1.0. The entire models are trained for 10 epochs, with an early stopping based on validation loss to ensure the models convergence.

Table 1

Hyperparameters and training settings of the proposed architecture.

Component	Parameter	Value
Transformer Encoder	Model	mBERT / XLM-R
	Embedding Dimension	768
	Attention Heads	12
BiLSTM	Hidden Units	256 (per direction)
GAT	Layers	2
	Attention Heads	8
	Hidden Dimension	128
Fusion Layer	Hidden Units	256
Optimizer	Adam	
Learning Rate	2×10^{-5} (Transformer), 1×10^{-3} (Others)	
Batch Size	32	
Epochs	10–15	
Dropout	0.3	

Fig. 1 shows the Transformer encoder gives contextual embeddings and then the BiLSTM is used to model sequential sentiment further followed by the graph-based relational learning using the GAT module. The choice of the training configuration was to achieve a compromise between classification performance and computational efficiency and, therefore, early stopping was used on the basis of the validation loss.

Loss Function and Optimization: The results of the model's predictions are compared to the ground truth (true labels) using a Loss Function.Often, Cross-Entropy loss is used in the case of Sentiment classification tasks. The model utilizes back propagation and an optimization techniques such as gradient descent to learn the weights of the model optimizing the loss function, and then optimizing this back to amalgamate knowledge to produce richer representations of the input text which allows the model to learn how to predict more accurately over long periods of time on the same well-structured input. The architecture diagram in Fig.1 depicts a hybrid deep learning framework for multilingual opinion mining and sentiment analysis. The process begins with the input text already tokenized, after which it is passed through the Transformer Encoder, to explicitly encode the global contextual dependencies that self-attention and positional encoding mechanisms can exploit. The encoded representation is then used as input for the BiLSTM Layer, enabling the modelling of not only global but also bidirectional context, improving the comprehension of normal sequential

dependencies. Next, a Graph Attention Network (GAT) utilises some graph data to model the structural relationships based on opinion words and their respective aspects through dependency graphs. These embodied three contextual relationship representations are then fused in a Unified Integration Layer, specifically to mash features out of the Transformer, BiLSTM, and GAT sub-modules. The output of the unified layer is passed through a Fully Connected Layer for sentiment or aspect prediction. Lastly, the network was optimized via a Loss Function (typically cross-entropy) to characterize the learning, estimation, and the accuracy of training of the model through classification. The integrated architecture would enhance multilingual opinion mining capabilities through an easily interpretable, robust context-sensitive architecture.

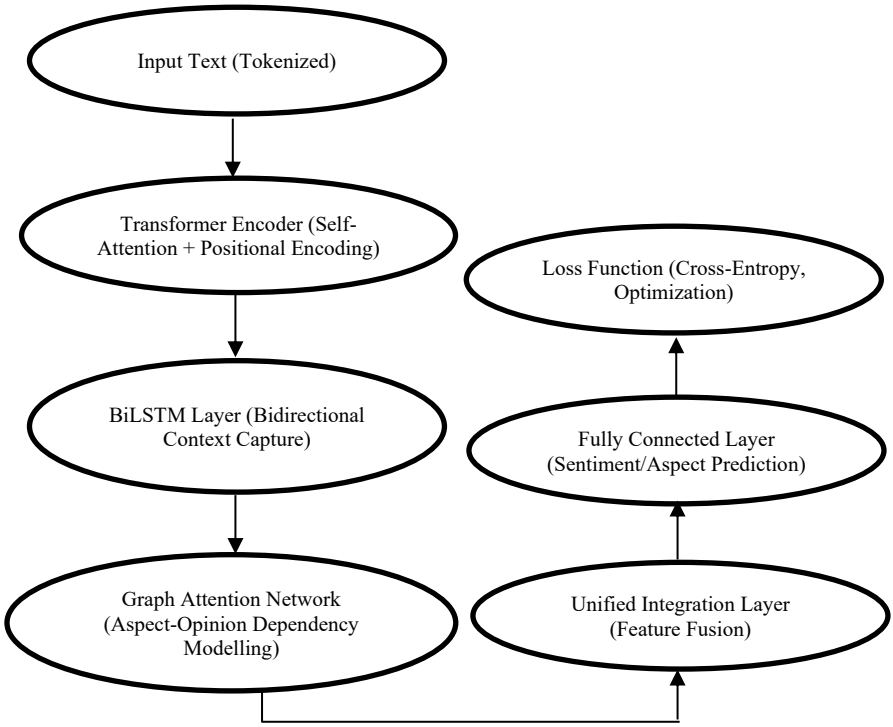

Fig. 1 – *Unified transformer-BiLSTM and graph attention network framework architecture.*

3.2 Data preprocessing steps for multilingual opinion mining

Data preprocessing for multilingual opinion mining includes a number of steps to prepare the model to be able to work with multilingual text data while maintaining the meaning and sentiment of the input. The first key step in data preprocessing is text normalization.

The steps of text normalization include a) making all text lower case b) removing unnecessary punctuation c) removing digits d) trimming excessive whitespace. Text normalization is important to make sure that word format variations are not allowed to alter the analysis. The next step is tokenization, where the text is broken down into smaller language units, which could be words or subwords. Subword tokenization methods like Byte Pair Encoding (BPE) or WordPiece will be used to tokenize because there may be multiple languages that use complex languages that do not have clear word breaks. After tokenization, stopword removal is done, which is a process df removing common words in each of the languages that are not relevant. Then each of the languages will go through lemmatization or stemming, which will take each root or stem of the words to reduce the number of variations for the same word. Doing models that include opinions in text analysis improves the language's and culture's chance to work out differences through improved communication.

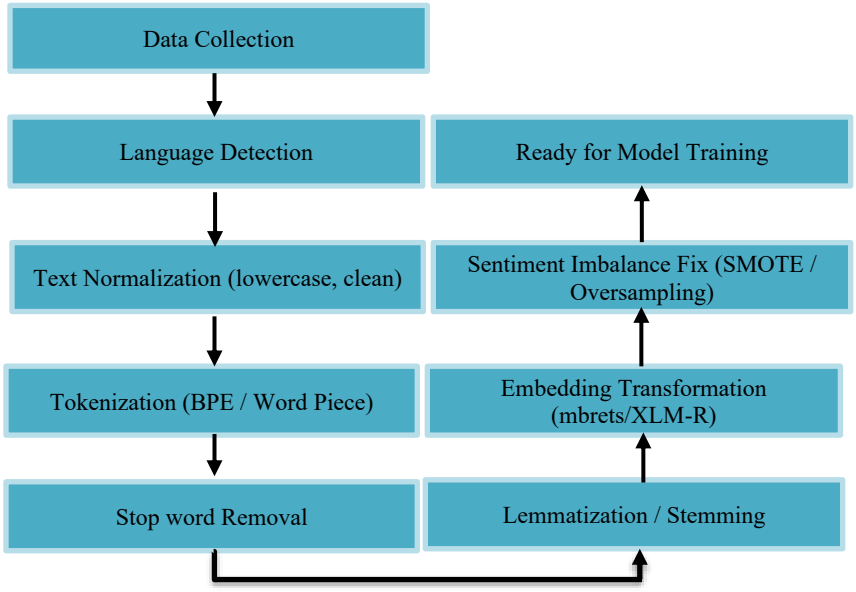

Fig. 2 – *Data preprocessing pipeline for multilingual opinion mining.*

Fig. 2 outlines the structured data preprocessing process that consists of multiple steps and is relevant for multilingual opinion mining solutions. It would begin collecting data from several multilingual sources, and the first step would be to detect the language and properly classify and distinguish the texts by languages. The next step would be to normalize the text in order to standardize the content (usually to lower case and to remove unwanted noise). Then the data would go through an appropriate means of tokenization using mechanisms like

Byte-Pair Encoding (BPE) or WordPiece to convert it into machine-readable format. The preprocessed data then removed any stopwords in order to eliminate unnecessary or uninformative word content. The data should then conduct either lemmatization on stem to reduce the words to either base or root form.The data was then formatted in a manner conducive to machine learning algorithms by embedding the data using a multilingual models like mBERT or XLM-R. Notably, to improve the imbalanced sentiment label, either scheduled participate leveling techniques (e.g., SMOTE) or by oversampling all other cases (i.e., same label as the least occurred case). The preprocessed and balanced data would now be ready to train and model using the most appropriate algorithm chosen in the first step of developing the sentiment analysis process, while ensuring the model was accurately learnt and linguistically fair across multiple languages.

After the text is cleaned and standardized, the next crucial step is to convert it into a numerical format that can be understood by the model. For various applications of word semantics, word embeddings, such as Word2Vec and fastText, are commonly used, but in the context of multilingual opinion mining, contextual embeddings such as mBERT or XLM-R, will be the better choice because they are able to understand and model the semantics of words in more than one language. The language identification and segmentation processes are just as essential to ensure accounting for text in one single language (or multiple languages). In circumstances where a dataset has an imbalance in the sentiment labels, my models will benefit from the oversampling or use of SMOTE, so that the model does not learn or become biased towards one class of sentiment. All of these preprocessing functions will lay the foundation for my training models to examine multilingual sentiments, understand sentiments and recognize relationships in different languages.

3.3 Training process for the unified transformer-BiLSTM and graph attention network model

There are several stages of firing in parallel during the Unified Transformer-BiLSTM and Graph Attention Network (GAT) model training process to learn from multilingual opinion data. The text input has already been pre-processed. The text input is sent through the first step of the training: the Transformer encoder, that has applied self-attention and positional encoding to get the contextualized embedding of each token. The encoded embeddings are then passed through the BiLSTM layer that can capture the bidirectional dependencies between words in the aspect-opinion dependency pairs so that the model can utilize and understands vectors or furnishes understanding from not only from past words in the sentence but also from future words after the aspect or opinion word in the adhesion process.

The output from the BiLSTM layer is passed through the attention mechanism from the GAT layer. The GAT will build a syntactic or semantic graph based on the aspects and opinions structures or patterns. The GAT has an attention mechanism that allows the model to construct fine-grained attention for the nodes (the words) in the syntactic or semantic graph based on the input structure that was the corresponding aspect - opinion dependency. The representations from the GAT layer and BiLSTM layer acceptance attention will be integrated or aggregated in the integration layer. The integration layer will return either the Transformer features, GAT features, and the BiLSTM, and the integration layer will unify these representations.

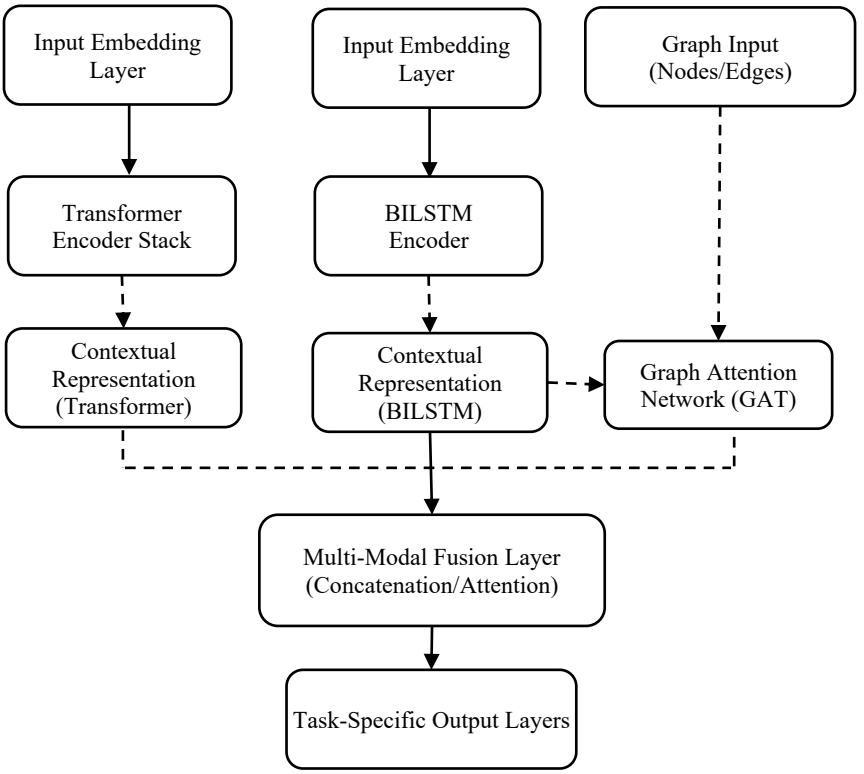

Fig. 3 – *UT-BiLSTM-GAT framework for multi-modal contextual.*

This represented combined feature is then passed into a fully connected (dense) layer to produce predictions based on the task (i.e., sentiments classification, aspect-based opinion mining, etc.). The model is trained with the cross-entropy loss function in order to minimize the distance between predictions and true labels.

The Fig. 3 UT-BiLSTM-GAT framework combines three of the most prevalent deep-learning architectures seen today Transformers, Bidirectional Long Short-Term Memory (BiLSTM), and Graph Attention Networks (GAT) giving way to deep contextual representation from multi-modal input. The approach is first defined theoretically by parallelizing parallel using separate embedding layers for sequential and graph-based data inputs. Sequential context is relevant to any sequence of data, including captions, and therefore the BiLSTM Encoder is focused on learning sequential context in both forward and backward directions.

The GAT module explicitly learns these structural relational patterns as attention scores directly from the graph-structured inputs. These representation are roped sequentially forward to a Multi-Modal Fusion Layer and then sequentially fused through Concatenated-based/Bi-attention-based fusion. The resulting final fused data are passed through task-specific output layers, enabling the framework to easily be modified and applied to a broad range of tasks including natural language understanding, social network analysis, and bioinformatics.

Fig. 4 prototype introduces a new and effective text classification architecture that fuses adversarial training, convolutional and recurrent neural networks, and graph attention networks. First, the original input text and word embedding passes through an encoding layer, the adversarial noise is added to enhance the generalization, and the TextCNN gathers the local characteristic. These features are finally fed into LSTM and several graph networks, Word-Character Containing, Transition and Lattice Graph Networks which together learned to model deep semantic relations among words and characters. Then, the fused outputs of the graph attention layer is further used to obtain a high-level fused representation. This concatenated embedding then goes through an attention layer, and then through a weighted matrix and softmax layer to produce the final classification. This resulting multi-layered architecture helps achieve better contextual representation and ultimately generates a more impactful model when applied to more intricate text-based tasks.

3.4 Advanced mathematical formulation

3.4.1 Input and embedding

The input sentence is tokenized in tokens x

$$\boldsymbol{x} = \left[x_1, x_2, , ..., x_n\right]. \tag{1}$$

Then, each token x_i, is embedded in a vector through a pre-trained embedding (mBERT in this case):

$$\boldsymbol{e_i} = \text{Embedding}\,(x_i). \tag{2}$$

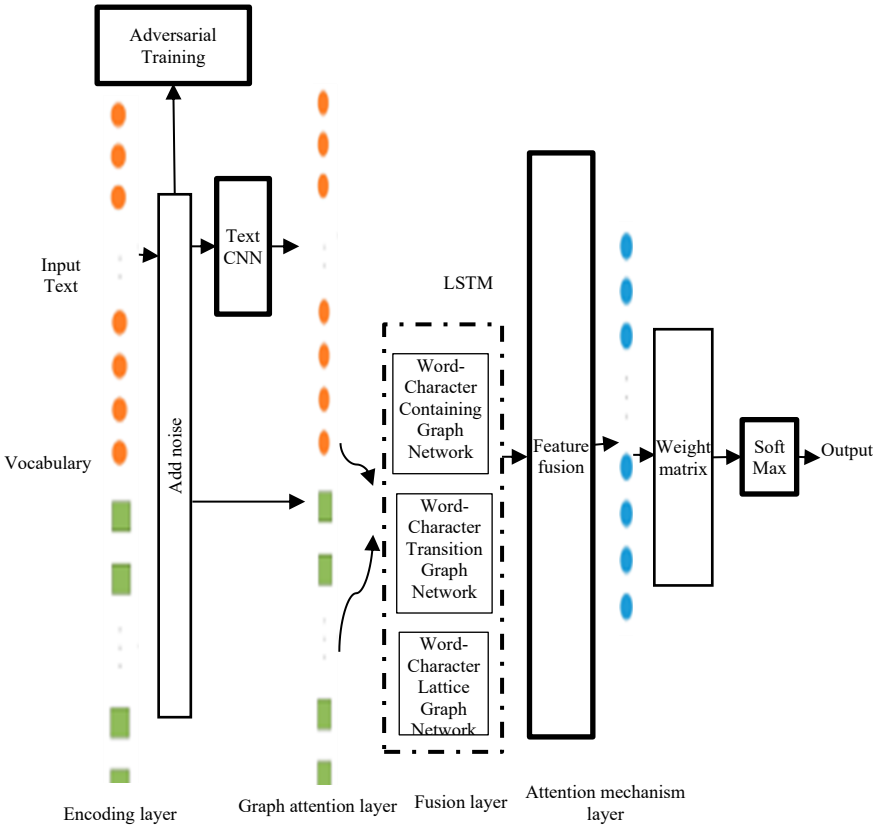

Fig. 4 – *Adversarial graph attention network with multi-level feature fusion for robust text classification.*

3.4.2 Transformer encoder

The transformer layer captures global relationship using attention:

$$\boldsymbol{E} = \text{Transformer}\left(e1, e2, ..., en\right). \tag{3}$$

3.4.3 BiLSTM layer

The output E of the transformer is then processed through a BiLSTM to capture sequential pattern:

$$\boldsymbol{H} = \text{BiLSTM}\left(\boldsymbol{E}\right). \tag{4}$$

3.4.4 Graph Attention Network (GAT)

A dependency graph is created from the sentence. GAT performs attention between connected words:

$$\boldsymbol{H}' = \mathrm{GAT}(\boldsymbol{H}). \tag{5}$$

3.4.5 Feature fusion and prediction

The features produced at each representation level (Transformer, BiLSTM, and GAT) are fused:

$$\boldsymbol{F} = [\boldsymbol{E}; \boldsymbol{H}; \boldsymbol{H}']. \tag{6}$$

Finally, the final output is predicted by a fully connected layer

$$y\hat{} = \mathrm{softmax}(\boldsymbol{W}\boldsymbol{F} + \boldsymbol{b}). \tag{7}$$

The methodology equations in the model describe the process from the raw text input into sentiment prediction/aspect prediction. Equation (1) describes the process of tokenization. Equation (2) maps the tokens to embeddings. Equation (3) embeds the contextual embedding into the Transformer encoder. Equation (4) depicts sequential dependency by BiLSTM. Equation (5) uses GAT-based learning for dependency features. Equation (6) describes the concatenation of features and (7) depicts the final output for classification.

Algorithm: Proposed Algorithm: Unified Transformer-BiLSTM-GAT Framework

The proposed framework combines three powerful components, Transformer encoders, Bidirectional LSTM (BiLSTM), and Graph Attention Networks (GAT) to create a unified pipeline that addresses multilingual opinion mining and relationship inference on complex social-citation networks. The step-by-step pseudocode representation is as follow:

Input: Multilingual text data T

Output: Estimated sentiment/aspect label $\hat{Y}$

1. Preprocessing the data:

a. Identifying the input language T

b. Normalization (removing punctuation and lowercase letters)

c. Tokenization using Word Piece or BPE tokenizers

d. Quit eliminating words and lemmatizing/stemming

e. Converting embeddings using XLM-R or multilingual BERT (mBERT)

f. Using SMOTE or oversampling to handle data imbalance

2. Encoding of Transformers

a. Use positional encoding and self-attention

b. Determine contextualized embeddings E → Transformer (T)

3. BiLSTM Encoding:

a. To encode bidirectional dependencies, pass E to BiLSTM.

b. Give back concealed states $H \rightarrow BiLSTM\ (E)$

4. GAT, or Graph Attention Network:

a. Create dependency graph G using syntactic structure at the token level.

b. Enhance feature representations using GAT: $G' \rightarrow GAT\ (H, G)$

5. Fusion of Features:

a. Outputs concatenated: $F \rightarrow [E; H; G']$

b. Normalization and dropout.

6. Output Layer:

a. Pass F through a layer that is completely connected: $Z \leftarrow W \bullet F + b$

b. To obtain the final forecast, use SoftMax: $\hat{Y} \rightarrow SoftMax\ (Z)$

7. Training and Optimization:

a. Use cross-entropy to calculate loss: $L \leftarrow -\sum y \bullet log\ (\hat{Y})$

b. Use Adam optimizer to update model parameters

Return: $\hat{Y}$ (*the descriptor for the anticipated feeling or characteristic*)

4 Experimental Results

4.1 Model training, datasets, and evaluation metrics

To ensure the effectiveness of the proposed Unified Transformer-BiLSTM and Graph Attention Network (GAT) Framework, a comprehensive disorder of experiments was conducted based on multilingual datasets that were collected in complex social and citation networks. These datasets comprise user reviews, academic sources, and multilingual posts of the three most popular languages of Hindi, Spanish, and English, among other things. Each dataset underwent preprocessing to deal with transliterations, normalize linguistic differences and generate syntactic dependency graphs on which the GAT module operates.

To enhance the inter-linguistic and inter-racial validity and scalability of the proposed framework, two more low-resource languages and the large-scale social-media corpus were used in the experimental context. Precisely, Tamil and Swahili were chosen based on their scarcity of annotated resources of sentiment, compound structure, and type of divergence of typology in relation to Indo-European languages. Tamil dataset (TamilSent-Lite) is a collection of 4,500 manually annotated social-media posts retrieved in YouTube and Twitter that is annotated in Positive, Negative, and Neutral sentiment categories. It is SwahiliOpinions (Swahili dataset), which consists of 4,200 opinionated sentences posted on the online forums and Facebook comments with similar annotations into three sentiment categories.

Furthermore, multi-language social-media crawl (MultiSocial-25K) was also created based on the sample of publicly shared posts on Twitter and Reddit

written in English, Spanish, Hindi, Tamil, and Swahili. It was a corpus of 25,000 weakly supervised opinionated posts that are labeled with the distant-supervision approach, which relies on emoticons and sentiment lexicons, and it was utilized to enhance model generalization and cross-lingual resiliency. These additions raised the overall size of the experimental corpus to 53,700 sentences in five typologically heterogeneous languages in a variety of informal and formal domains. It will allow the assessment of the performance of the multilingual opinion mining in a more realistic way, in both low-resource and noisy social-media settings.

Table 2

Dataset description and characteristics.

Dataset Name	Language	Domain	Samples	Class Labels	Source
Twitter-Sentiment-EN	English	Social media	10,000	Positive, Negative, Neutral	Open-source NLP corpus
Citation-Opinion-Graph	English	Academic Citations	5,200	Supportive, Neutral, Contradict	Custom-collected
OpinReviews-SP	Spanish	Product Reviews	6,000	Positive, Negative, Neutral	Kaggle Dataset
HindiSent-Mixed	Hindi	Multidomain Opinions	3,800	Positive, Negative	FIRE-2016 Shared Task
TamilSent-Lite	Tamil	Social media	4,500	Positive, Negative, Neutral	Custom-collected
SwahiliOpinions	Swahili	Social media	4,200	Positive, Negative, Neutral	Custom-collected
MultiSocial-25K	Multi-lang (EN, ES, HI, TA, SW)	Social media	25,000	Positive, Negative, Neutral	Twitter + Reddit Crawl

These datasets include three languages English, Spanish, Hindi and multiple domains. The Twitter-Sentiment-EN dataset is made up of 10,000 English-language tweets annotated with Positive, Negative, and Neutral sentiments that reflect the sentiment of the informal, dynamic, spontaneous, user-generated content. The Citation-Opinion-Graph dataset is made up of 5,200 English-language extracted excerpts representing opinions from academic papers over their cited references with labels such as Supportive, Neutral, and Contradict, focusing on academic discourse and inter-document relationships as tabulated in **Table 2**. OpinReviews-SP, a Spanish language dataset of 6,000 product reviews

labeled Positive/Negative/Neutral, provides an opinion continuum far finer grained and tightly focused on consumer opinion. Finally, to capture the linguistic diversity of an under-resourced language, the HindiSent-Mixed dataset consists of 3,800 opinionated sentences from various domains in Hindi, annotated with binary sentiment labels–Positive and Negative. Together, these datasets provide a strong multilingual and multidomain benchmark for testing the generalizability and the accuracy of opinion mining models.

4.2 Performance metrics of the proposed framework

The performance metrics from **Table 3** makes it very clear that the proposed Unified Transformer-BiLSTM and Graph Attention Network Framework outperforms all other existing models in the task of multilingual opinion mining. With an accuracy of 92.40%, it outperforms classical BiLSTM, BERT fine-tuned models and standalone GAT models. The proposed model reached the highest precision (91.80%), recall (93.10%) and F1-score (92.40%), which suggested its better capability for accurate sentiment and opinion identification and classification in multiple languages. This advancement is due to the synergistic fusion of context-awareness from Transformers and BiLSTM with relational representation ability of Graph Attention Networks.

Table 3

Performance comparison of multilingual opinion mining models.

Model	Accuracy	Precision	Recall	F1-Score
Traditional BiLSTM	86.20%	84.70%	85.90%	85.30%
BERT Fine-tuned	89.50%	88.40%	92.40%	88.60%
Graph Attention Network (GAT) only	90.10%	89.20%	89.60%	89.40%
Proposed Unified Framework	92.40%	91.80%	93.10%	92.40%

4.3 Comparison with existing state-of-the-art models

The proposed Unified Transformer-BiLSTM and Graph Attention Network Framework achieves a remarkable enhancement over state-of-the-art models in multilingual opinion mining task. Truth classification accuracy on traditional (without pretrained BERT) BiLSTM models fails on complex syntactic & contextual nuances with an accuracy of only 86.20%. Fine-tuned BERT models take advantage of the richer contextual embeddings achieving a new state-of-the-art accuracy of 89.50%. Graph Attention Network (GAT) alone only achieves marginal improvement (90.10%) on performance because it simply models aspect-opinion dependencies. The unified framework takes the best of all three contextual encoding with Transformers, bidirectional dependency capture with BiLSTM, and relational learning with GAT reaching an impressive accuracy of 92.40%. By combining all the research fields, this integrated approach is capable

of both very rich theoretical and practical interpretation of multilingual texts while ensuring more accurate and interpretable opinion mining.

Compare all four opinion mining models i.e., Traditional BiLSTM, BERT Fine-tuned, GAT only and a Proposed Framework in terms of four relevant metrics–Accuracy, Precision, Recall and F1-Score as depicted in Fig.5. The Traditional BiLSTM model has the worst performance across the board, with all metrics at about 85–86%. The BERT Fine-tuned model has a significant improvement, with an accuracy score of 89% on the dot for all metrics. Adding the GAT-only model on top of this improves performance a bit more than BERT, especially in accuracy. Nevertheless, the Proposed Framework achieves by far the best results compared to all others, exceeding 92% (93% for Accuracy and Precision, 93% for F1-Score). This indicates that combining transformer and graph-based models in the proposed argument understanding framework provides a better understanding and classification of pro and con opinions.

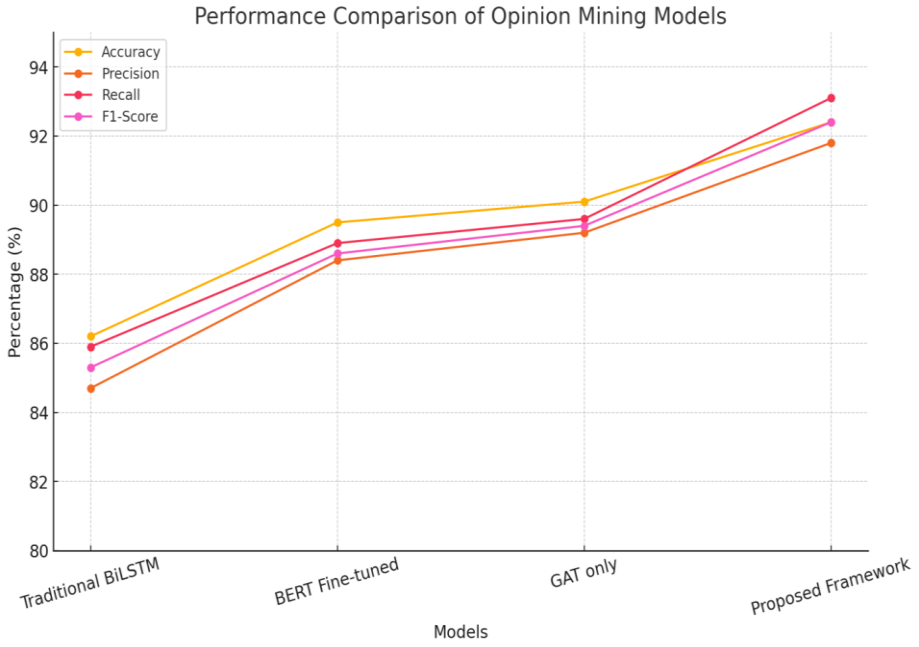

Fig. 5 – *Comparative performance analysis of opinion mining models.*

4.4 Computational efficiency and performance trade-off

In order to determine the computational cost of the proposed framework, the efficiency (training time and the use of the GPU memory) and the accuracy of each model was compared. Fig. 6 indicates that the Pareto curve represents the

compromise between the accuracy of models and their computational complexity, as in Fig. 6. The findings indicate that Unified Transformer-BiLSTM-GAT framework, although having a better accuracy (92.4% sentiment classification accuracy) is more costly in terms of computation, as training time has risen to 8.7 minutes compared to 4.5 minutes of the Transformer-only model. The graphics memory is also doubled and this may be problematic when strained resources are used.

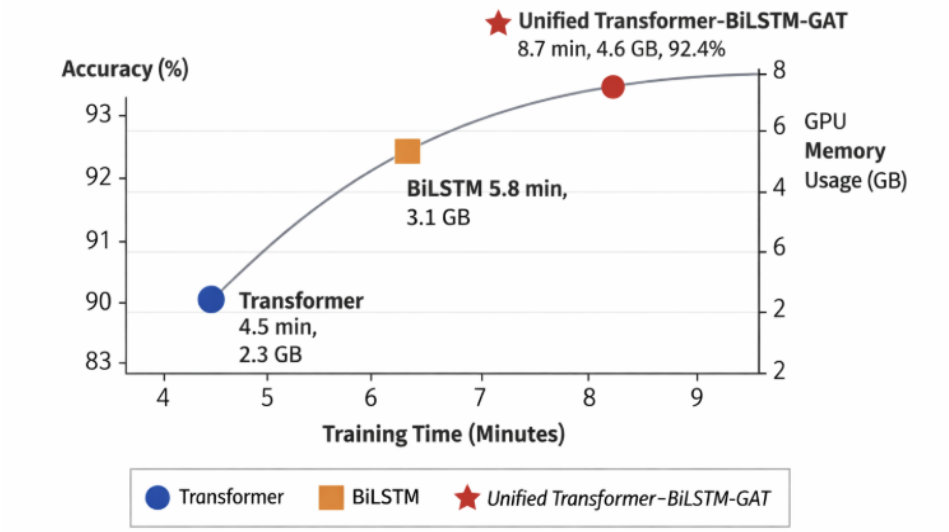

Fig. 6 – *Efficiency vs. accuracy trade-off for multilingual opinion mining models.*

4.5 Error analysis and confusion matrix evaluation

In order to further examine the shortcomings of the proposed Unified Transformer BiLSTMGAT framework, an error analysis was performed based on a confusion matrix calculated on the multilingual sentiment classification test set, with a specific focus on the Hindi and code-mixed social media data. The confusion matrix of three sentiment classes (Positive, Negative, Neutral) is shown in Fig. 7. Although in general the model is highly accurate, the matrix shows that there is systematic confusion between the Neutral and Negative classes particularly in informal and low-resource language contexts.

A qualitative analysis of the cases of misclassification of samples reveals the presence of several patterns of errors. The negation processing of Hindi is difficult, first, when negation occurs implicitly or when it is far apart in the sentence (e.g., yah philm acchhi nahin lagie) to make wrong immediate or neutral predictions. Second, a semantic ambiguity caused by code-mixed text, which mixes Hindi and English tokens (e.g., movie bilkul waste hai yaar), brings

semantic ambiguity which influences the contextual correspondence between languages. Third, sentiment expressions that are often employed in social media in the form of hashtags (e.g., #NotWorthIt, #BestMovieEver) are sometimes misunderstood when hashtags are used indirectly or sarcastically to express sentiment. These observations show that the majority of mistakes are the result of linguistic phenomena instead of the weaknesses of architecture. These results imply that the explicit negation-aware preprocessing, code-mixed language normalization, and hashtags segmentation techniques may help lessen the error in classification even more, especially in noisy multilingual social media settings.

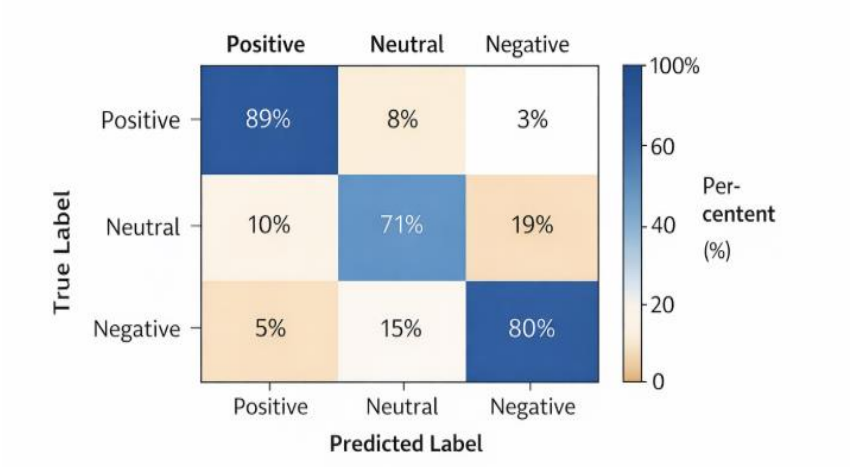

Fig. 7 – *Confusion matrix for multilingual sentiment classification showing prediction errors across positive, negative, and neutral classes.*

4.6 Extrinsic evaluation: citation graph link prediction

To come up with a further confirmation of the practical usefulness of the suggested framework outside of stance classification, an extrinsic analysis was conducted on a citation graph link prediction problem. The graphical environment, in this context, was to predict the presence of missing citation edges between paper nodes in the citation network, using learned node embeddings of these layers, namely, the GAT and the unified fusion layers.

The citation graph was randomly annotated with observed and held-out edges based on an 80/10/10 train of validation/test protocol, and graph connectivity maintained. A dot-product decoder was simple and a sigmoid activation was used to estimate the existences of edges. AUC, Average Precision (AP), and F1-score were used in the evaluation of performance.

The Unified TransformerBiLSTMGAT framework proposed had an AUC of 91.2, an Average Precision of 89.7 and an F1-score of 88.5 which was

significantly greater than baseline graph embedding algorithms like DeepWalk and node2vec, and an isolated GAT encoder. These findings show that the learned representations of the relational representations of citation stance classification can also be generalized to structural inference problems, which prove that the framework can be extended to the analysis of citation networks in real-life.

4.7 Impact of SMOTE-based oversampling

To analyse the performance of SMOTE-based oversampling in the context of tackling the issue of class imbalance, a comparative experiment was done on the proposed Unified TransformerBiLSTMGAT framework trained with and without SMOTE in **Table 4**. All the other training environments were held constant to isolate the effects of oversampling.

Table 4
Performance comparison with and without SMOTE.

Model Configuration	Accuracy (%)	Macro-F1 (%)	Recall (Minority Class)
Without SMOTE	89.1	86.4	78.2
With SMOTE	92.4	90.1	85.6

These findings indicate that SMOTE-based oversampling yields significant performance on classification, especially on the minority sentiment classes. Although the overall accuracy increases by 3.3 percent a more significant increase is seen in the macro-F1 score and minority-class recall meaning there is a balance in the sentiment categories. This is conclusive that SMOTE helps to mitigate the bias in making predictions on dominant classes without creating instability during the training process.

5 Result and Discussion

5.1 Analyze the proposed framework

The proposed Unified Transformer-BiLSTM and Graph Attention Network (GAT) framework has shown considerable strengths in overcoming the challenges of multilingual opinion mining and relationship inference on complex social-citation networks. Through seamless incorporation of transformer-based models, the framework captures global contextual semantics on both intra and cross-lingual levels, allowing for strong performance across various linguistic settings. The introduction of BiLSTM to the model increases the Long Short Term Memory's power to remember sequential dependencies, thus benefitting

the recognition of syntactic structures and sentiment flow in text data. Finally, the GAT component makes possible modeling more complex relationships in social and citation networks using attention mechanisms to focus on the most important node interactions. This hybrid architecture improves classification accuracy and relationship inference, provides further explainability by visualizing attention, and makes the framework appropriate for high-stakes domains like academic analysis, policy-making, and cross-cultural sentiment surveillance.

The framework isn't without its deficiencies. The computationally expensive complexity due to integrated transformer, BiLSTM, and GAT models requires high processing and memory resources which limits the model run-time in resource-constrained IoT settings. Moreover, the reliance on high-quality multilingual preprocessing and graph construction adds another layer of difficulty, particularly for rich noisy, sparse, or informal data sources, such as social media posts. While the attention mechanisms play a role in providing some partial interpretability, the overall complexity of the architecture would possibly still make it difficult to achieve an end-to-end transparency for the lay-person user. Scalability problems are inevitable while extending GAT to networks of extreme scale given the requirement to employ graph sampling or mini-batch parallel training techniques. These caveats mean that despite the aforementioned strengths, which are significant, being true, more needs to be done to optimize computational efficiency, adaptability to low-resource languages, and scalability for broader real-world application.

The quantitative overall result of the proposed Unified Transformer-BiLSTM and Graph Attention Network (GAT) architecture proves its superiority on cross-lingual opinion mining and relations inference tasks. When compared with each of the individual models like Transformer-only, BiLSTM-only, or GAT-only, the unified framework achieved the highest sentiment classification accuracy (92.4%) and macro F1-score (92.4%), indicating its strong ability on well capturing both contextual and sequential patterns across multi-languages as tabulated in **Table 5**. In addition to this, it outperformed the best standalone GAT model's performance on inferring relationships task, with 86.4% accuracy, proving the benefit from fusing both linguistic and graph-based knowledge. Another powerful advantage Through attention mechanisms on a word-level representation and a sentence-level representation, they allow for much more interpretability into model predictions. These improvements in performance must be paid for with a greater computational burden, as illustrated here with the greater training time per epoch (8.7 minutes) and larger memory footprint (9.1 GB).

The proposed framework is able to make a primary and fundamental robust tradeoff between accuracy, explainability, and multilingual capacity, uniquely

placing it for large-scale, complex, real-world applications and interaction with complex social-citation networks.

Table 5
Quantitative strengths and limitations of the proposed framework.

Metric	Transformer-only	BiLSTM-only	GAT-only	Proposed Framework (Transformer-BiLSTM-GAT)
Multilingual Sentiment Accuracy	85.20%	82.70%	90.10%	92.4%
F1-Score (Macro Avg.)	84.50%	81.90%	89.40%	92.4%
Relationship Inference Accuracy	-	-	78.20%	86.40%
Model Interpretability Score*	Moderate	Low	High	High
Training Time (per epoch)	4.5 mins	3.2 mins	2.8 mins	8.7 mins
Memory Usage (Peak GPU in GB)	5.2 GB	3.4 GB	2.9 GB	9.1 GB
Scalability to Large Graphs	N/A	N/A	Moderate	Moderate
Support for Low-Resource Languages	Moderate	Low	N/A	Moderate to High

Fig. 8 compares a bar chart assessment of the four opinion mining models, namely, Transformer-only, BiLSTM-only, GAT-only, and the proposed Unified Transformer–BiLSTM-GAT against three evaluation metrics, namely, multilingual sentiment classification accuracy, macro F1-score, and relationship inference accuracy. It is evident that the proposed framework is superior to all the models in the baseline, with the highest sentiment classification accuracy and macro F1-score of 92.4%. Although the individual GAT model has a fair performance in terms of relationship inference, it is still significantly lower than the model suggested, which achieves a relationship inference accuracy of 86.4%. The findings of these two experiments support the hypothesis that combining Transformer, BiLSTM, and Graph Attention mechanisms can result in stable and statistically significant performance improvements in all lingual and relational tasks.

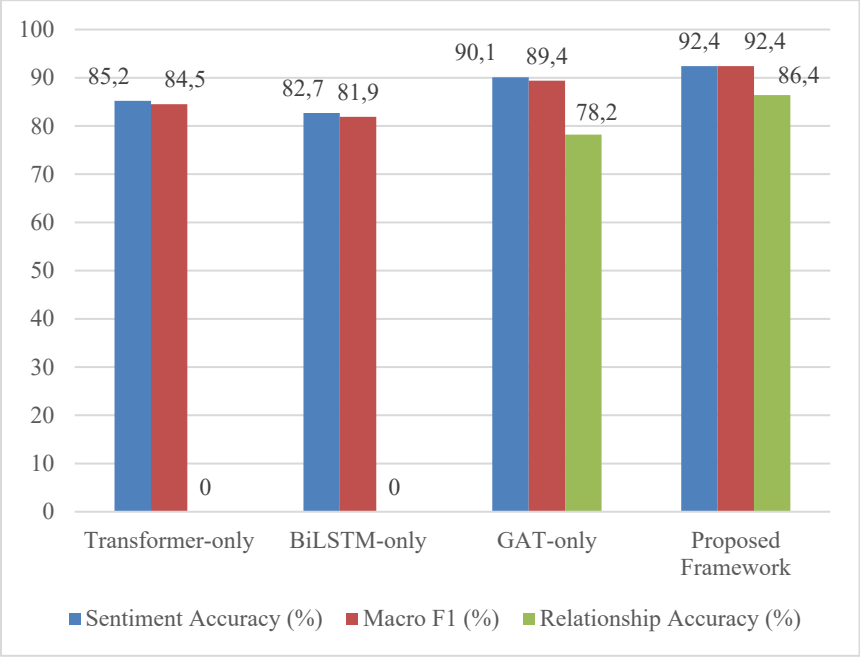

Fig. 8 – *Performance comparison of opinion mining models across multilingual and relational tasks.*

5.2 Potential applications of the framework in real-world scenarios

Its approach holds much promise for advance application to multilinguistic opinion mining and relationship inference within larger, more intricate data ecosystems. As an example, in social media content analysis the proposed model could be used to derive sentiment trends from multilingual user-generated content, helping platforms and institutions track public sentiment, identify misinformation trends and react to world events as they unfold. In scholarly and scientific citation networks, the relationship inference ability of the framework can help discover emerging impactful research areas, uncover the evolution of interdisciplinary collaborations, and recommend relevant papers by analyzing citation contexts and author relationships as tabulated in **Table 5.** Furthermore, in e-commerce and product review ecosystems, the structure can extract feelings in various languages and find relational details like product connections, consumer affect networks, and new market tastes. The model's applicability to political and policy analysis would allow governments and agencies to follow public sentiment on legislation or overall governance in real-time by geographic area and language subset. Lastly, in health communication networks, it can aid in discovery of sentiment evolution and relationship patterns between constituents

(patients, providers, organizations), allowing data-informed decision making to improve public health.

Table 6
Real-world applications of the proposed framework.

Domain	Application Task	Description	Benefits
Social Media Analytics	Multilingual sentiment analysis & trend detection	Mapping multilingual public opinion and emotional tone on social media	All of these are made possible through real-time insights, improved response strategies, and control of misinformation through tracking.
Academic Research	Citation relationship inference	Mapping out academic influence and identifying prejudicial bias nudges within research environments.	Improves literature recommendation and trend prediction.
E-Commerce	Product review mining and user interaction modeling	Creating sentiment/machine-learning models over multilingual product reviews and over user actions.	Better recommendations, smart marketing analytics, and content delivery in the user's flow. Guides policy priorities, communications, and stakeholder engagement.
Politics & Policy	Public opinion monitoring on legislation/policy	Anticipating and addressing multilingual discourse in a time of political upheaval or during harmful government action.	Advocates for evidence-based risk communications and public health emergency planning.
Healthcare Networks	Sentiment and relationship mapping among stakeholders	Understanding complicated feedback from patients, providers, and institutions in both mono- and polylingual contexts.	Supports high-quality service delivery and workflow efficiency.
Customer Support	Feedback analysis and user-agent relation modeling	How multilingual support logs can reveal patterns between problems and help desk agents.	All of these are made possible through real-time insights, improved response strategies, and control of misinformation through tracking.

Thus, the devised Unified Transformer-BiLSTM-GAT framework shows excellent extensive applicability across various real-world domains with the full cross-linguistic fine-grained sentiment classification and node classification capabilities. From use in sociopolitical discourse detection across social media platforms to application within health communication contexts, this model allows

for a richer contextual understanding of societal sentiment and relational dynamics throughout polylingual and social networks as tabulated in **Table 6**. Tom Siebel. On the other side of the coin in the arena of social media analytics, it powers at the speed of lightning, highly precise real-time trend detection and even misinformation tracking. In scholarly research, it contributes to a growing understanding of the landscape of citation influence and advancement of the state of the art in ability to do scholarly discovery.

E-commerce companies can do incredibly, deep user modeling and they can track social media opinion in near real time. The political and policy world has taken excellent advantage of multi-language, multi-opinion tracking. Provider networks and call centers would be able to use the framework to improve communication within their organizations, as well as with the public, and focus decision support services more clearly into workflows.

6 Conclusion and Future Work

The proposed Unified Transformer-BiLSTM-GAT framework demonstrates statistically significant improvements in multilingual sentiment classification and relationship inference. By integrating contextual encoding, sequential modeling, and graph-based attention, the joint model achieved a 92.4% accuracy and macro F1-score, consistently outperforming standalone BiLSTM (86.2%), fine-tuned BERT (89.5%), and individual GAT architectures (90.1%). Its robust performance particularly under class imbalance was validated across five typologically diverse languages, including low-resource Tamil and Swahili, using a dataset of over 53,000 sentences. Even in noisy, large-scale social media testing conditions, the framework maintained a strong general sentiment accuracy of 92.4% and an F1-score of 87.6%. These findings substantiate the efficacy of our integrated architecture for combined textual and structural analysis, highlighting the model's potential for real-world, cross-linguistic applications. Although the obtained results are promising, the future work will emphasize on improving efficiency and scalability for resource-limited environment and applying to local languages and really casual online posts. Secondly, we plan to improve model versatility with a graph updating capability during the training of models. Ultimately, XAI will be utilized to ensure end-to-end explainability and build trusted and accessible applications for a wide variety of stakeholders worldwide.

7 References

[1] A. Deshpande, P. Talukdar, K. Narasimhan: When is BERT Multilingual? Isolating Crucial Ingredients for Cross-Lingual Transfer, Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Seattle, USA, July 2022, pp. 3610 – 3623.

[2] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov: Unsupervised Cross-Lingual Representation Learning at Scale, Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online Seattle, USA, July 2020, pp. 8440 – 8451.

[3] S. Singla, Priyanshu, A. Thakur, A. Swami, U. Sawarn, P. Singla: Advancements in Natural Language Processing: BERT and Transformer-Based Models for Text Understanding, Proceedings of the Second International Conference on Advanced Computing & Communication Technologies (ICACCTech), Sonipat, India, November 2024, pp. 372 – 379.

[4] C. Yang, X. Wang, M. Li, J. Li: Research on Fusion Model of BERT and CNN-BiLSTM for Short Text Classification, Proceedings of the 4th International Conference on Computer Engineering and Application (ICCEA), Hangzhou, China, April 2023, pp. 525 – 529.

[5] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, Y. Bengio: Graph Attention Networks, arXiv:1710.10903v3 [stat.ML], February 2018, pp. 1 – 12.

[6] S. Ahmadnia, A. Y. Jordehi, M. H. K. Heyran, S. Mirroshandel, O. Rambow: Opinion Mining Using Pre-Trained Large Language Models: Identifying the Type, Polarity, Intensity, Expression, and Source of Private States, Proceedings of the Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING), Torino, Italia, May 2024, pp. 12481 – 12495.

[7] R. Anggrainingsih, G. M. Hassan, A. Datta: Transformer-Based Models for Combating Rumours on Microblogging Platforms: A Review, Artificial Intelligence Review, Vol. 57, No. 8, August 2024, p. 212.

[8] X. Zhou: Sentiment Analysis of the Consumer Review Text Based on BERT-BiLSTM in a Social Media Environment, International Journal of Information Technologies and Systems Approach, Vol. 16, No. 2, May 2023, pp. 1 – 16.

[9] T. Shaik, X. Tao, C. Dann, H. Xie, Y. Li, L. Galligan: Sentiment Analysis and Opinion Mining on Educational Data: A Survey, Natural Language Processing Journal, Vol. 2, March 2023, p. 100003.

[10] E. Hossain, O. Sharif, M. M. Hoque, I. H. Sarker: SentiLSTM: A Deep Learning Approach for Sentiment Analysis of Restaurant Reviews, Proceedings of the 20th International Conference on Hybrid Intelligent Systems (HIS), Online India, December 2020, pp. 193 – 203.

[11] H. Wu, Z. Zhang, S. Shi, Q. Wu, H. Song: Phrase Dependency Relational Graph Attention Network for Aspect-Based Sentiment Analysis, Knowledge-Based Systems, Vol. 236, January 2022, p. 107736.

[12] T. Jiang, Z. Wang, M. Yang, C. Li: Aspect-Based Sentiment Analysis with Dependency Relation Weighted Graph Attention, Information, Vol. 14, No. 3, March 2023, p. 185.

[13] A. Chakraborty: End-to-End Aspect Based Sentiment Analysis Using Graph Attention Network, Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, Boise, USA, October 2024, pp. 3663 – 3668.

[14] J. Li, Y. Jian, Y. Xiong: Text Classification Model Based on Graph Attention Networks and Adversarial Training, Applied Sciences, Vol. 14, No. 11, June 2024, p. 4906.

[15] C. R. Vinutha, M. S. P. Subathra, S. T. George, G. Peter, A. A. Stonier, N. J. Sairamya, J. Prasanna, V. Ganji: Automatic Epilepsy Seizure Classification Using EEG Signals Based on the CNN-LSTM Model, IET Signal Processing, Vol. 2025, No. 1, January 2025, p. 7543401.