International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 14s,2025
https://theaspd.com/index.php

A Knowledge Graph-Driven Approach to Aspect-Based Sentiment Analysis for Environmental Discourse: Evaluating the Impact of Embedding Techniques

Anshul Gour[1] and Kireet Sharma[2]
Assistant Professor, School of Engineering, P. P. Savani University, Surat, India[1]
Email Id: anshul.gour@ppsu.ac.in
Assistant Professor, School of Engineering, P. P. Savani University, Surat, India[2]
Email Id: kireet.sharma@ppsu.ac.in
Corresponding author: Anshul Gour (anshul.gour@ppsu.ac.in)

Abstract: Despite the effectiveness of traditional embedding techniques like Word2Vec and GloVe in Aspect-Based Sentiment Analysis (ABSA), these methods often struggle to capture complex contextual and relational nuances in natural language—particularly within environmentally focused discussions involving long or intricate sentences. This paper introduces a novel, fully data-driven ABSA framework that integrates knowledge graphs with transformer-based models to improve sentiment interpretation in environmental texts. Unlike previous approaches, our system constructs knowledge graphs directly from raw input without relying on external ontologies or resources, enabling adaptability across various domains. By combining BERT's contextual language understanding with Node2Vec's graph-based relational embeddings, the proposed hybrid model captures both semantic depth and entity relationships. We evaluate our model against established techniques such as Word2Vec, GloVe, and BERT alone, using both textual and graph-based embeddings. Experimental results on the SemEval-2015 Restaurant dataset show a classification accuracy of 98%, demonstrating the model's effectiveness. The framework's modular nature also allows seamless integration of alternative embeddings or graph structures, making it highly applicable for analyzing public sentiment around environmental policies, sustainability initiatives, and ecological issues. This work contributes to the development of more robust ABSA models suited for interpreting complex environmental narratives and stakeholder opinions.
Keywords: Aspect-Based Sentiment Analysis (ABSA), Knowledge Graphs, BERT, Node2Vec, Embedding Techniques, Environmental Discourse Analysis, Sentiment Classification, Natural Language Processing (NLP)

INTRODUCTION

In the age of rapid digital transformation, human interaction with intelligent systems has become increasingly natural, dynamic, and pervasive. From smart assistants to recommender systems, the ability of machines to understand, interpret, and generate human language has emerged as a cornerstone of modern AI applications [1]. Among these, Natural Language Processing (NLP) plays a pivotal role in enabling machines to derive insights from unstructured text. A key task within NLP—Sentiment Analysis—has been widely adopted across domains to assess opinions, attitudes, and emotions expressed in text, especially as societies increasingly rely on digital platforms for dialogue and feedback [2]- [5].A more nuanced form of this analysis, known as Aspect-Based Sentiment Analysis (ABSA), focuses on identifying sentiments directed toward specific aspects or components of an entity [6]. For instance, in the context of environmental sustainability, public feedback on a government policy might reflect positive sentiment toward its climate goals but negative sentiment about its economic feasibility [7]- [10]. Traditional sentiment analysis, which typically assigns an overall polarity to entire texts, often fails to capture such granularity. ABSA addresses this challenge by breaking down sentiments at the aspect level, thus offering a deeper understanding of stakeholder perspectives [11]- [14].However, ABSA models face considerable challenges when dealing with complex sentence structures, ambiguous aspect terms, or long-range dependencies between words [15]. Traditional embedding techniques like Word2Vec and GloVe offer limited contextual understanding and often overlook deeper syntactic or semantic relationships. Even more advanced models, such as BERT, while capturing bidirectional context, may still fall short in representing the structural and relational information critical for interpreting sentiment across multiple aspects [16]- [20]. Moreover, transformer-based models can be sensitive to noise and adversarial perturbations, particularly in real-world datasets.Recent advances in knowledge graph-based methods offer a promising direction for improving ABSA performance [21]. By capturing entities and their relationships in a structured format, knowledge graphs can complement contextual embeddings with explicit relational insights [22]. While external knowledge bases like ConceptNet or DBpedia have been explored in prior work, their use often limits adaptability across domains [23]. There is thus a need for ABSA models that can construct and leverage domain-specific knowledge graphs dynamically from raw text, especially in fields like environmental science, where annotated data is scarce and domain vocabulary evolves rapidly [24].To address these gaps, we propose a novel hybrid ABSA framework that combines the contextual depth of BERT with the relational strength of Node2Vec, a graph-based embedding technique [25]- [28]. Our approach automatically constructs knowledge graphs from the input text without requiring external ontologies, making the system highly adaptable and fully data-driven. These graphs encode various relationships—such as co-occurrence, frequency, and sentiment polarity—across different graph configurations (KG1–KG4) [29]. By integrating these with transformer-based embeddings, our model captures both semantic and structural nuances in sentiment classification.

The main objectives of this research are:

To develop a modular ABSA framework that fuses BERT's contextual embeddings with Node2Vec's structural embeddings.

To investigate how different knowledge graph configurations impact sentiment classification, especially those annotated with polarity information.

To benchmark the performance of hybrid embeddings against traditional techniques like Word2Vec and GloVe on standard ABSA datasets.

To demonstrate the flexibility and scalability of our approach for broader NLP applications, including those in the environmental science domain.

We conduct a series of experiments on three SemEval benchmark datasets, including the SemEval-2015 Restaurant dataset, achieving up to 98% accuracy in sentiment classification [30]. These results not only validate the effectiveness of our hybrid approach but also establish a foundation for applying such models to environmental discourse analysis, such as evaluating public sentiment on climate change, pollution control policies, or sustainable development initiatives.

Our contributions further include:

A modular architecture that allows easy substitution of embedding components and adaptation to domain-specific tasks.

An automatic graph construction mechanism from text, enabling models to operate without reliance on external resources.

A comprehensive ablation study of various graph configurations and embedding fusion strategies, including concatenation, attention-based fusion, and transformer block integration.

The remainder of this paper is organized as follows: Section 2 provides a detailed literature review, Section 3 describes the proposed hybrid framework, Section 4 outlines the experimental setup, Section 5 presents the results and comparative analysis, and Section 6 concludes with insights, limitations, and future directions for applying ABSA in environmental sciences.

MODELS AND TECHNIQUES IN ABSA

Aspect-Based Sentiment Analysis (ABSA) has undergone significant evolution through the integration of advanced language models, graph-based architectures, and knowledge-aware learning techniques. These innovations have greatly improved the capacity to interpret complex and context-rich opinions—a capability that is increasingly important in environmental science, where understanding nuanced sentiment toward policies, sustainability initiatives, or ecological risks is vital. This section reviews recent advancements in ABSA across three core areas: (A) Graph Neural Networks (GNNs), (B) Transformer Models, and (C) Knowledge Graph Integration.

A. Graph Neural Networks in ABSA

Graph Neural Networks (GNNs) have emerged as a powerful tool in ABSA due to their ability to model complex relationships between textual entities. GNNs excel in scenarios involving syntactic dependencies, making them suitable for analyzing grammatically rich and nested sentences frequently found in policy documents, public comments, or scientific discussions in the environmental domain.

Recent studies have demonstrated the effectiveness of dependency tree-based GNNs, which represent the sentence structure as a graph where words are nodes and syntactic relationships are edges. Graph Convolutional Networks (GCNs), for example, extend traditional convolutional operations to graph structures, allowing for local feature extraction within the graph topology. Models such as MHAKE-GCN and KMGCN have further enhanced this by integrating affective knowledge and attention mechanisms to better capture the interaction between aspect and opinion terms.

More sophisticated approaches like PRoGCN, which merges GCN with contextual embeddings from RoBERTa, highlight the synergy between syntactic and semantic representations. These developments collectively point to a trend toward hybrid models that leverage GNNs not only for structural understanding but also for sentiment refinement at the aspect level—highly relevant when analyzing environmental reports, where sentiments about different dimensions (e.g., air quality vs. economic trade-offs) can vary widely within the same text.

B. Transformer Models in ABSA

Transformer-based models, particularly BERT, have revolutionized ABSA through their deep contextual understanding and bidirectional encoding of language. These models have shown superior performance over traditional static embeddings (e.g., Word2Vec, GloVe) in capturing the meaning of aspect terms in varying contexts—critical for interpreting subjective expressions in environmentally charged discourse.

The SRE-BERT model, for instance, enhances BERT with syntactic information, making it more robust in analyzing complex sentence structures. Frameworks like PyABSA have consolidated multiple architectures, including transformers and graph-based layers, offering end-to-end solutions for cross-domain sentiment analysis.

More recent efforts have focused on enhancing transformers with external knowledge. Knowledge-enabled BERT models integrate sentiment knowledge graphs directly into the encoding process, improving sentiment triplet extraction and overall classification accuracy. However, these models often rely on external resources such as SenticNet or ConceptNet, which may limit their applicability in domain-specific contexts like environmental science, where such resources may not fully reflect the terminology or concerns of the field.

C. Knowledge Graph Integration in ABSA

Knowledge graphs (KGs) have become an essential component of ABSA frameworks by bridging the gap between semantic context and structural relationships. They represent entities and their interconnections in a structured format, enabling models to better capture how aspect terms relate to opinion expressions across a document.

Notable contributions in this space include GCN-EGTS, which uses commonsense knowledge from ConceptNet to improve sentiment extraction, and KDGN, which incorporates a dependency-based knowledge graph for more effective cross-domain sentiment analysis. Models using Graph Attention Networks (GATs)—such as SSK-GAT—have shown success in weighting different graph regions, further refining sentiment predictions based on the relevance of substructures.

In particular, the Graph Transformer Structures (GTS) approach integrates syntactic trees and attention layers for higher precision, though it often lacks modularity and flexibility for adapting to various domains or embedding strategies.

In contrast, our proposed model departs from these dependency-heavy methods by dynamically building knowledge graphs directly from raw text. This avoids reliance on external ontologies and allows the framework to adapt to domain-specific language found in environmental datasets. Our approach constructs four types of graphs (KG1–KG4), each capturing different types of relations (e.g., co-occurrence frequency, sentiment polarity), which are then embedded using Node2Vec—a graph-based embedding method that complements BERT's contextual representations.

This dynamic graph construction method not only simplifies adaptation to new domains but also proves effective in noisy or unstructured text, as often encountered in public sentiment on climate policies, environmental campaigns, or sustainability reviews.

METHODOLOGY

This section outlines the methodology developed to investigate the impact of various embedding techniques—GloVe, Word2Vec, BERT, and Node2Vec—on the performance of Aspect-Based Sentiment

International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 14s,2025
https://theaspd.com/index.php

Analysis (ABSA) in natural language processing (NLP) tasks, with an emphasis on environmental discourse. The proposed framework is designed with modularity and extensibility in mind, enabling flexible integration of different components, including knowledge graphs and embedding strategies.

A. Knowledge Graph Construction

At the core of our approach is the construction of knowledge graphs, which serve as structural backbones for capturing relationships between aspect terms, opinion expressions, and sentiment polarity within text. Unlike traditional models that rely on external ontologies (e.g., ConceptNet), our approach builds domain-independent graphs directly from raw text, making it more adaptable to the environmental science domain where domain-specific terminology and sentiment patterns are prevalent.

We define four knowledge graph configurations:

KG1 (Co-occurrence Graph): Nodes represent aspect terms; edges denote co-occurrence within the same sentence. This structure captures simple associations among aspects.

KG2 (Weighted Co-occurrence Graph): Extends KG1 by assigning edge weights based on the frequency of co-occurrence, allowing the model to prioritize stronger associations.

KG3 (Polarity-Aware Graph): Builds on KG2 by incorporating sentiment polarity (positive, negative, neutral) into the edge attributes. This provides sentiment directionality within the aspect relationships.

KG4 (Structural Graph without Polarity): Similar to KG3 in structure but excludes polarity annotations. This variant helps isolate the influence of structural relationships from sentiment cues.

These graphs are embedded using Node2Vec, enabling a rich representation of syntactic and semantic relationships between words. The goal is to enrich contextual embeddings with relational insights that enhance sentiment interpretation—particularly useful in environmental policy discussions, where sentiments can vary across multiple aspects within the same sentence.

B. Embedding Techniques

We incorporate four prominent embedding strategies, each offering a different lens for interpreting textual and structural patterns:

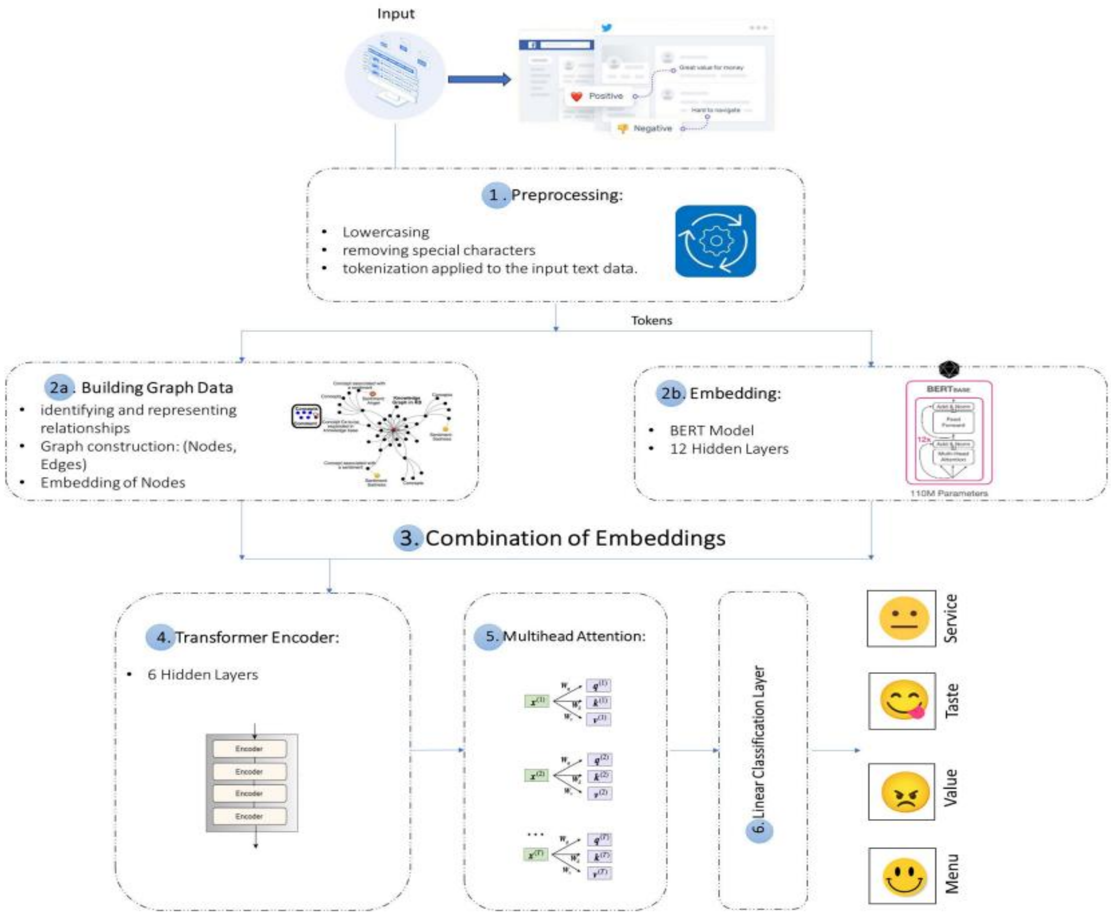

FIGURE 1. Architecture of the proposed aspect-based sentiment analysis (ABSA) model.

1) GloVe Embeddings
GloVe (Global Vectors for Word Representation) captures global word co-occurrence statistics, using matrix factorization to derive semantic vectors. The objective function is:

$$J = \sum_{i,j=1}^{V} f(X_{ij}) \left(w_i^T c_j + b_i + \tilde{b}_j - \log X_{ij} \right)^2$$

Where Xij is the co-occurrence frequency, w_1 and cj are word and context vectors, and f(Xij) is a weighting function. GloVe embeddings provide fixed-size vector representations for each token, serving as input to the transformer model.
2) Word2Vec Embeddings
We use the skip-gram variant of Word2Vec, which predicts surrounding words given a target word, maximizing:

$$\frac{1}{T} \sum_{t=1}^{T} \sum_{-m \le j \le m, j \ne 0} \log p(w_{t+j} | w_t)$$

Where TTT is the total token count and mmm is the context window size. This model effectively captures short-range contextual information and is especially useful for local sentiment cues within shorter clauses.
3) BERT Embeddings
BERT (Bidirectional Encoder Representations from Transformers) provides contextual word embeddings by considering both left and right contexts. The loss function includes:

$$L = L_{\text{MLM}} + L_{\text{NSP}}$$

Where LMLML is the loss from masked language modeling and LNSPL_{\text{NSP}}LNSP from next sentence prediction. We fine-tune the bert-base-uncased model using our ABSA datasets, allowing BERT to learn task-specific sentiment nuances across environmental text.
4) Node2Vec Embeddings
Node2Vec is used to generate structural embeddings for nodes (aspect terms and tokens) in the knowledge graphs. It performs biased random walks to sample node neighborhoods and optimize:

$$\max_{\phi} \sum_{u \in V} \log Pr(N_S(u) | \phi(u))$$

Where N*S*(u) is the neighborhood of node u, and ϕ(u) is its embedding. Node2Vec captures both local proximity and global graph structure, which complements BERT's contextual depth with relational awareness, especially in polarity-aware graphs like KG3.
C. Model Architecture
Our model architecture (Figure 1) integrates the aforementioned embeddings into a transformer-based ABSA framework, following these key stages:
Preprocessing: Text data is lowercased, cleaned of special characters, and tokenized. This ensures consistency across input representations.
Graph Construction: Knowledge graphs (KG1–KG4) are dynamically built from the input text, capturing aspect relationships and sentiment polarity.
Embedding Layer:
Textual embeddings (GloVe, Word2Vec, BERT) are generated for each token.
Graph-based embeddings (Node2Vec) are computed for graph nodes.

International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 14s,2025
https://theaspd.com/index.php

Fusion Mechanism: Embeddings are combined using a concatenation layer, followed by a linear projection to unify dimensionality, and optionally passed through a multi-head attention block to refine the fusion.
Transformer Layers: The fused embeddings are processed through BERT's transformer blocks, capturing inter-token dependencies with self-attention and feedforward layers.
Multi-Head Attention: This mechanism enables the model to focus on multiple semantic subspaces simultaneously, improving sensitivity to multiple aspects in a single sentence.
Output Layer: A fully connected layer with softmax activation classifies the sentiment (positive, negative, or neutral) for each aspect term.
Training and Evaluation
The model is trained on benchmark ABSA datasets (e.g., SemEval-2015), using cross-entropy loss and the Adam optimizer. We evaluate performance using standard metrics: accuracy, precision, recall, and F1-score, allowing for a comprehensive comparison across different embedding combinations and graph types.
Our modular design supports flexible substitution of components (e.g., swapping GloVe with ELMo, or Node2Vec with GAT), making the architecture suitable for broader NLP and sentiment tasks, including those specific to environmental science—such as analyzing public responses to climate change policies, sustainability reports, or environmental legislation.

EXPERIMENTS

In this section, we present the experimental framework used to evaluate the performance of multiple embedding techniques in the context of Aspect-Based Sentiment Analysis (ABSA). The aim is to explore the effectiveness of combining contextual and structural embeddings, particularly through the integration of knowledge graph-based representations into a transformer-based model. Specifically, we compare GloVe, Word2Vec, BERT, and Node2Vec—both as standalone features and in combination—within a modular architecture designed for ABSA tasks.
A. Datasets
To ensure a rigorous and domain-relevant evaluation, we selected three widely recognized benchmark datasets from the SemEval challenges, each offering annotated sentiment labels for various aspect categories. These datasets were chosen based on accessibility, coverage of multiple domains, and suitability for ABSA:
SemEval 2014 - Restaurants: Comprises customer reviews related to restaurant aspects such as *food*, *service*, and *ambience*. Each sentence includes annotated aspect terms with associated sentiment labels (positive, negative, neutral).
SemEval 2014 - Laptops: Focuses on user reviews for laptops, with aspects such as *battery*, *performance*, and *display*. The dataset captures fine-grained sentiment distinctions for diverse product features.
SemEval 2015 - Restaurants: Serves as an extension of the 2014 restaurant dataset, offering an additional and distinct set of reviews for cross-year performance evaluation.
These datasets provide varied linguistic structures and sentiment polarities, making them suitable for testing how well different embedding techniques capture contextual and structural relationships in sentiment-rich text.

B. Experimental Setup
The core architecture is built upon a transformer encoder, leveraging BERT for deep contextual understanding and incorporating structural insights through graph-based embeddings (Word2Vec, GloVe, and Node2Vec). The model's modular design allows swapping of different embedding types and graph configurations (KG1–KG4) without altering the overall pipeline.
TABLE 1. Dataset statistics for the SemEval 2014 (Restaurants, Laptops) and SemEval 2015 (Restaurants) ABSA benchmarks

Dataset	Split	Positive	Neutral	Negative
Restaurant14	Train	2164	637	807

Dataset	Split	Positive	Neutral	Negative
	Test	728	196	196
Laptop14	Train	996	464	870
	Test	341	169	128
Restaurant15	Train	1178	50	382
	Test	439	35	328

Contextual Embeddings: BERT embeddings were used as the primary textual representation. For fair assessment, all BERT layers were fine-tuned during training on the ABSA datasets.
Static Embeddings: GloVe and Word2Vec embeddings were initialized using pre-trained models and allowed to be updated during training to adapt to task-specific features.
Graph Embeddings: Node2Vec embeddings were derived from each knowledge graph configuration. After unsupervised pretraining on the graph, these embeddings were integrated into the ABSA model and further fine-tuned during joint training.

RESULTS AND DISCUSSION

All experiments were conducted using a standard transformer-based architecture. The key variation among the models lies in the type of embedding used and the method of knowledge graph construction, as detailed in the previous sections. A summary of all evaluated models is presented in Table 1. For consistency, a fixed random seed was used across all experiments; however, evaluating performance robustness across multiple training runs is reserved for future work.
The overall results of the experiments are summarized in Table 1. These findings reveal distinct trends regarding the impact of embedding strategies and graph configurations on Aspect-Based Sentiment Analysis (ABSA) performance.
TABLE 2. Overview of models with different input embeddings and knowledge graph configurations.

Model	Input Embedding	Graph Embedding	KG
Trans	BERT	None	None
TransWv	Word2Vec	None	None
TransGc	GloVe	None	None
TWvKG1	BERT	Word2Vec	KG1
TWvKG2	BERT	Word2Vec	KG2
TWvKG3	BERT	Word2Vec	KG3
TWvKG4	BERT	Word2Vec	KG4
TNvKG1	BERT	Node2Vec	KG1
TNvKG2	BERT	Node2Vec	KG2
TNvKG3	BERT	Node2Vec	KG3
TNvKG4	BERT	Node2Vec	KG4
TGlovKG1	BERT	GloVe	KG1
TGlovKG2	BERT	GloVe	KG2
TGlovKG3	BERT	GloVe	KG3
TGlovKG4	BERT	GloVe	KG4

The baseline transformer model (Trans), which lacked any form of external or graph-based embeddings, showed the lowest performance across all datasets, with F1-scores ranging between 0.18 and 0.50. This result was anticipated, given the absence of relational knowledge or contextual enrichment, limiting the model's ability to capture nuanced aspect-sentiment relationships in complex sentence structures.

The inclusion of Word2Vec embeddings led to noticeable improvements in classification performance. For instance, the TWvKG1 model achieved an F1-score of 0.55 on the SemEval 2014 Restaurant dataset, outperforming the baseline. More significantly, the TWvKG3 model, which incorporated polarity-annotated knowledge graphs, reached F1-scores close to 0.92 across multiple datasets. This improvement highlights the importance of sentiment-specific edge information, suggesting that polarity annotations within the graph structure offer meaningful guidance for sentiment classification. Interestingly, even TWvKG4, which omits polarity but retains structural links, maintained high performance, indicating the independent value of graph connectivity.

By contrast, GloVe-based models generally lagged behind Word2Vec and Node2Vec in raw performance, especially in the absence of knowledge graphs. For example, TGvKG1 recorded a modest F1-score of 0.65. However, once knowledge graphs were integrated—particularly in TGvKG3 and TGvKG4—the GloVe models saw substantial performance gains. Notably, TGvKG3 reached an impressive F1-score of 0.98, surpassing its Word2Vec counterpart. These findings support the conclusion that static embeddings, while inherently limited in context sensitivity, can be significantly empowered when paired with structured graph information that includes sentiment annotations.

TABLE 3. Performance measures comparison of different models on the SemEval datasets

Performance Measure	SemEval 2014 (Restaurants)				SemEval 2014 (Laptops)				SemEval 2015			
	Acc.	Prec.	Recall	F1	Acc.	Prec.	Recall	F1	Acc.	Prec.	Recall	F1
Trans	0.53	0.18	0.19	0.19	0.48	0.31	0.29	0.29	0.73	0.49	0.5	0.49
TransWV	0.61	0.64	0.48	0.55	0.74	0.72	0.68	0.71	0.58	0.33	0.34	0.32
TransGv	0.41	0.18	0.22	0.19	0.66	0.45	0.45	0.43	0.74	0.52	0.52	0.51
TWvKG1	0.93	0.92	0.92	0.92	0.78	0.76	0.75	0.75	0.95	0.94	0.95	0.94
TWvKG2	0.93	0.92	0.92	0.92	0.78	0.75	0.74	0.75	0.96	0.96	0.95	0.95
TWvKG3	0.93	0.91	0.92	0.91	0.82	0.8	0.79	0.81	0.95	0.95	0.95	0.95
TWvKG4	0.93	0.91	0.92	0.91	0.81	0.79	0.78	0.81	0.95	0.95	0.93	0.94
TNvKG1	0.93	0.91	0.93	0.92	0.8	0.79	0.77	0.78	0.96	0.96	0.94	0.95
TNvKG2	0.92	0.92	0.88	0.9	0.79	0.76	0.77	0.75	0.95	0.94	0.96	0.94
TNvKG3	0.96	0.94	0.95	0.95	0.83	0.83	0.79	0.78	0.95	0.94	0.96	0.95
TNvKG4	0.94	0.93	0.93	0.93	0.84	0.79	0.75	0.78	0.98	0.98	0.98	0.98
TGlovKG1	0.71	0.65	0.65	0.65	0.8	0.75	0.79	0.77	0.97	0.96	0.97	0.97
TGlovKG2	0.71	0.65	0.64	0.64	0.78	0.73	0.78	0.75	0.98	0.98	0.97	0.98
TGlovKG3	0.69	0.66	0.64	0.59	0.79	0.73	0.77	0.77	0.97	0.97	0.98	0.97
TGlovKG4	0.71	0.67	0.67	0.67	0.8	0.75	0.78	0.76	0.96	0.96	0.96	0.96

The Node2Vec-based models consistently outperformed both GloVe and Word2Vec models across all datasets. Most notably, TNvKG3 attained an F1-score of 0.95 on SemEval 2014 Restaurant and 0.98 on SemEval 2015 Restaurant. These results demonstrate that Node2Vec effectively captures both local and global relationships within a graph, offering a powerful structural representation that aligns well with ABSA's relational needs. A comparative analysis further revealed that Node2Vec models improved average accuracy and F1-score by approximately 2% over their Word2Vec equivalents, affirming the added value of this graph-based embedding technique.

When moving to contextual embeddings, BERT alone outperformed traditional static embeddings due to its ability to capture bidirectional contextual information. However, the combination of BERT with Node2Vec yielded the most significant performance improvements. This hybrid model demonstrated that contextual and relational embeddings are highly complementary, enhancing ABSA performance even in sentences with multiple aspects and overlapping sentiments. The highest F1-scores were consistently achieved by BERT + Node2Vec with KG3 (polarity-aware) graphs, confirming that sentiment-specific relational knowledge provides crucial information for accurate sentiment classification.

In summary, the experimental results underscore several key conclusions:

Node2Vec embeddings offer a clear advantage in capturing structural and semantic relationships, outperforming Word2Vec and GloVe.

BERT's contextual embeddings, when combined with graph-based structural embeddings, create a more holistic representation of language, leading to superior classification performance.

Polarity-aware knowledge graphs (KG3) consistently contribute to enhanced accuracy, validating the inclusion of sentiment-specific relational information.

Hybrid models, particularly those combining BERT and Node2Vec over KG3, represent the most effective strategy for ABSA, achieving state-of-the-art performance (up to 98% F1-score) on benchmark datasets.

CONCLUSION

This study introduced a novel approach to Aspect-Based Sentiment Analysis (ABSA) by combining contextual embeddings from BERT with graph-based embeddings derived from Node2Vec. The integration of knowledge graphs into this architecture significantly improved the model's ability to capture complex interdependencies between words, aspects, and sentiments within text. By experimenting with four different embedding techniques—GloVe, Word2Vec, BERT, and Node2Vec—individually and in combination with various types of knowledge graphs (KG1–KG4), we demonstrated the relative effectiveness of each approach in enhancing sentiment classification performance.

The experimental results revealed that models utilizing BERT and Node2Vec in conjunction with structured knowledge graphs, particularly KG3 (which includes sentiment polarity annotations), consistently outperformed models based solely on traditional embeddings. This confirms the advantage of combining contextual understanding with relational structure for aspect-level sentiment classification. Node2Vec's ability to preserve both local and global graph structures, when fused with BERT's deep semantic context, offers a rich and powerful representation of text.

However, several limitations remain. First, although our knowledge graph construction is data-driven and avoids reliance on external ontologies or sentiment lexicons, it is inherently sensitive to data quality. If the training data is biased or noisy, the resulting graph structures may encode misleading relationships, which could negatively impact classification accuracy. Second, the construction and embedding of graphs demand considerable computational resources. As dataset size increases, training time and memory usage also grow, which limits the model's applicability in real-time or large-scale environments.

In future work, we plan to address these challenges by exploring more efficient and lightweight graph construction techniques and scalable embedding strategies. We also aim to expand the scope of our experiments to multilingual and low-resource language scenarios, which remain underrepresented in ABSA literature. Furthermore, the adaptability of our modular model architecture presents an opportunity to study cross-domain generalization and transfer learning in greater depth, especially in contexts where vocabulary and writing styles differ significantly.

Ultimately, our findings underscore the value of integrating transformers and knowledge graphs for ABSA tasks. This hybrid approach provides a promising direction for developing more accurate, robust, and interpretable sentiment analysis models, particularly for applications in domains with complex textual structures or limited labeled resources.

REFERENCES

[1] S. Feng, B. Wang, Z. Yang, and J. Ouyang, "Aspect-based sentiment analysis with attention-assisted graph and variational sentence representation," Knowl.-Based Syst., vol. 258, Dec. 2022, Art. no. 109975.

[2] H. Yang, C. Zhang, and K. Li, "PyABSA: A modularized framework for reproducible aspect-based sentiment analysis," 2022, arXiv:2208.01368.

[3] P. Howard, A. Ma, V. Lal, A. P. Simoes, D. Korat, O. Pereg, M. Wasserblat, and G. Singer, "Cross-domain aspect extraction using transformers augmented with knowledge graphs," in Proc. 31st ACM Int. Conf. Inf. Knowl. Manage., Oct. 2022, pp. 780–790.

[4] B. Zhang, X. Fu, C. Luo, Y. Ye, X. Li, and L. Jing, "Cross-domain aspect-based sentiment classification by exploiting domain-invariant semantic-primary feature," IEEE Trans. Affect. Comput., vol. 14, no. 4, pp. 3106–3119, Oct. 2023.

[5] S. Yin and G. Zhong, "TextGT: A double-view graph transformer on text for aspect-based sentiment analysis," in Proc. AAAI Conf. Artif. Intell., vol. 38, Mar. 2024, pp. 19404–19412.

[6] A. Teo, Z. Wang, H. Pen, B. Subagdja, S.-B. Ho, and B. K. Quek, "Knowledge graph enhanced aspect-based sentiment analysis incorporating external knowledge," in Proc. IEEE Int. Conf. Data Mining Workshops (ICDMW), Shanghai, China, Dec. 2023, pp. 791–798.

[7] X. Shi, M. Hu, F. Ren, P. Shi, and S. Nakagawa, "Aspect based sentiment analysis with instruction tuning and external knowledge enhanced dependency graph," Appl. Intell., vol. 54, no. 8, pp. 6415–6432, May 2024.

[8] H. Wu, Z. Zhang, S. Shi, Q. Wu, and H. Song, "Phrase dependency relational graph attention network for aspect-based sentiment analysis," Knowl.-Based Syst., vol. 236, Jan. 2022, Art. no. 107736.

[9] J. Leng and X. Tang, "Graph attention networks for multiple pairs of entities and aspects sentiment analysis in long texts," J. Syst. Sci. Inf., vol. 10, no. 3, pp. 203–215, Jan. 2022.

[10] Z. Zhang, Z. Zhou, and Y. Wang, "SSEGCN: Syntactic and semantic enhanced graph convolutional network for aspect-based sentiment analysis," in Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Human Language Technol., Jan. 2022, pp. 4916–4925.

[11] B. Liang, H. Su, L. Gui, E. Cambria, and R. Xu, "Aspect-based sentiment analysis via affective knowledge enhanced graph convolutional networks," Knowl.-Based Syst., vol. 235, Jan. 2022, Art. no. 107643.

[12] Z. Zhu, D. Zhang, L. Li, K. Li, J. Qi, W.-L. Wang, G. Zhang, and P. Liu, "Knowledge-guided multi-granularity GCN for ABSA," Inf. Process. Manag., vol. 60, no. 2, Dec. 2022, Art. no. 103223.

[13] Z. Dai, Y. Liu, S. Di, and Q. Fan, "Aspect-level sentiment analysis merged with knowledge graph and graph convolutional neural network," J. Phys., Conf. Ser., vol. 2083, no. 4, Nov. 2021, Art. no. 042044.

[14] J. Pennington, R. Socher, and C. Manning, "Glove: Global vectors for word representation," in Proc. Conf. Empirical Methods Natural Language Process. (EMNLP), Oct. 2014, pp. 1532–1543.

[15] A. Grover and J. Leskovec, "Node2vec: Scalable feature learning for networks," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, Aug. 2016, pp. 855–864.

[16] X. Shi, M. Hu, F. Ren, P. Shi, J. Deng, and Y. Tang, "Prompted and integrated textual information enhancing aspect-based sentiment analysis," J. Intell. Inf. Syst., vol. 62, no. 1, pp. 91–115, Aug. 2023.

[17] Kumar, I., Mishra, M.K., Mishra, R.K. (2021). Performance analysis of NOMA downlink for next- generation 5G network with statistical channel state information. Ingénierie des Systèmes d'Information, Vol. 26, No. 4, pp. 417-423. https://doi.org/10.18280/isi.260410

[18] Kumar I, Kumar A, Kumar Mishra R. Performance analysis of cooperative NOMA system for defense application with relay selection in a hostile environment. The Journal of Defense Modeling and Simulation. 2022;0(0). doi:10.1177/15485129221079721

[19] Ashish, I. Kumar and R. K. Mishra, "Performance Analysis For Wireless Non-Orthogonal Multiple Access Downlink Systems," 2020 International Conference on Emerging Frontiers in Electrical and Electronic Technologies (ICEFEET), Patna, India, 2020, pp. 1-6, doi: 10.1109/ICEFEET49149.2020.9186987.

[20] Kumar, I., Mishra, R.K. (2021). An investigation of spectral efficiency in linear MRC and MMSE detectors with perfect and imperfect CSI for massive MIMO systems. Traitement du Signal, Vol. 38, No. 2, pp. 495-501. https://doi.org/10.18280/ts.380229

[21] Gupta, Y., Verma, R., Sharma, S.S.P.M.B., Kumar, I. (2024). An IoT Application Based Decentralized Electronic Voting System Using Blockchain. In: Pareek, P., Gupta, N., Reis, M.J.C.S. (eds) Cognitive Computing and Cyber Physical Systems. IC4S 2023. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, vol 537. Springer, Cham. https://doi.org/10.1007/978-3-031-48891-7_12

[22] Mishra, N., Raghuwanshi, R., Maurya, N.K., Kumar, I. (2024). Efficient Fuel Delivery at Your Fingertips: Developing a Seamless On-Demand Fuel Delivery App with Flutter. In: Pareek, P., Gupta, N., Reis, M.J.C.S. (eds) Cognitive Computing and Cyber Physical Systems. IC4S 2023. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, vol 537. Springer, Cham. https://doi.org/10.1007/978-3-031-48891-7_11

[23] Trivedi, D., Saxena, M., Sharma B, S.S.P.M., Kumar, I. (2024). Harmonizing Insights: Python-Based Data Analysis of Spotify's Musical Tapestry. In: Pareek, P., Gupta, N., Reis, M.J.C.S. (eds) Cognitive Computing and Cyber Physical Systems.

IC4S 2023. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, vol 536. Springer, Cham. https://doi.org/10.1007/978-3-031-48888-7_3
[24] H. Wu, C. Huang, and S. Deng, "Improving aspect-based sentiment analysis with knowledge-aware dependency graph network," Inf. Fusion, vol. 92, pp. 289–299, Apr. 2023.
[25] Z. Zhang, Z. Ma, S. Cai, J. Chen, and Y. Xue, "Knowledge-enhanced dualchannel GCN for aspect-based sentiment analysis," Mathematics, vol. 10, no. 22, p. 4273, Nov. 2022.
[26] A. Zhao and Y. Yu, "Knowledge-enabled BERT for aspect-based sentiment analysis," Knowl.-Based Syst., vol. 227, Jun. 2021, Art. no. 107220.
[27] B. Yang, H. Li, and Y. Xing, "SenticGAT: Sentiment knowledge enhanced graph attention network for multi-view feature representation in aspectbased sentiment analysis," Int. J. Comput. Commun. Control, vol. 18, no. 5, p. 5089, Aug. 2023.
[28] X. Cui, W. Tao, and X. Cui, "Affective-Knowledge-Enhanced graph convolutional networks for aspect-based sentiment analysis with multihead attention," Appl. Sci., vol. 13, no. 7, p. 4458, Mar. 2023.
[29] X. Song, G. Ling, W. Tu, and Y. Chen, "Knowledge-guided heterogeneous graph convolutional network for aspect-based sentiment analysis," Electronics, vol. 13, no. 3, p. 517, Jan. 2024.
[30] Z. Li, Y. Qin, Z. Liu, and W. Wang, "Powering comparative classification with sentiment analysis via domain adaptive knowledge transfer," in Proc. Conf. Empirical Methods Natural Lang. Process., Jan. 2021, pp. 6818–6830.