

Knowledge-Enhanced BERT for Aspect-Based Sentiment Analysis of Tourism Destinations Using Social Media Data

A. Irumporai

Department of Computer Science and Engineering, Amrita School of Computing, Amrita Vishwa Vidyapeetham, Chennai, India
a_irumporai@ch.students.amrita.edu (corresponding author)

G. Anitha

Department of Computer Science and Engineering, Amrita School of Computing, Amrita Vishwa Vidyapeetham, Chennai, India
g_anitha@ch.amrita.edu

Received: 3 July 2025 | Revised: 13 September 2025 and 1 October 2025 | Accepted: 14 October 2025

Licensed under a CC-BY 4.0 license | Copyright (c) by the authors | DOI: <https://doi.org/10.48084/etasr.13139>

ABSTRACT

Tourists are increasingly relying on social media platforms to share their impressions and ideas about various places. To improve customer satisfaction and service quality, tourism professionals must understand these sentiments. However, conventional methods of sentiment analysis often struggle to capture the subtle context and specific language of tourism content. Another factor that affects the accuracy of sentiment classification is the lack of integration of domain-specific knowledge. Using social media data, this study presents a Knowledge-Enhanced BERT (KE-BERT) model to analyze feelings about various things related to tourism destinations, with the aim of tackling the above challenges. The KE-BERT model enhances its understanding of aspect terms and sentiment within the BERT framework by incorporating knowledge from the tourism industry. A comprehensive dataset of social media reviews and comments on popular tourist destinations around the world was used. The training and evaluation of the KE-BERT model utilized several performance metrics, including accuracy, F1-score, recall, and precision. The experimental results show that KE-BERT outperforms existing models, including traditional BERT- and LSTM-based models, achieving an accuracy of 92.7% and an F1 score of 91.5%.

Keywords-aspect-based sentiment analysis; Knowledge-Enhanced BERT (KE-BERT); tourism destination; social media data; sentiment polarity

I. INTRODUCTION

The travel industry greatly helps the world economy by generating major income and employment opportunities. Today, tourists post their experiences and comments on various locations in social media sites, including TripAdvisor, Yelp, and Google Reviews [1]. For visitors considering traveling and people involved in the tourism sector, User-Generated Content (UGC) on these sites is an excellent source of information, as it can facilitate the extraction of sentiment polarity (positive, negative, or neutral) related to specific elements such as accommodation, food, or transportation [2]. Sentiment analysis, including Aspect-Based Sentiment Analysis (ABSA) [3], has recently drawn much interest. Conventional methods of sentiment analysis sometimes have difficulties in correctly classifying sentiment due to the complex and context-dependent character of the language related to tourism [4].

ABSA for travel data presents several data-related difficulties. First, data acquired on social media sites embed informal language, slang, and abbreviations, complicating language processing [5]. Vague aspect terms in tourism content often depend on context and require substantial awareness to properly classify sentiment [6]. Third, current models mainly use statistical methods or simple neural networks, which struggle to understand specific knowledge and details in user reviews. Travel trends and the dynamic character of the language add yet another level of challenge [7]. Standard word embeddings and shallow neural network models can not adequately capture the subtle nuances and sophisticated sentiment structures of travel reviews [8]. Moreover, the situation is further exacerbated by the absence of external knowledge sources, such as specific ontologies for tourism and semantic data, to provide precise sentiment matching with elements [9].

This study presents a Knowledge Enhanced BERT (KE-BERT) model, which performs aspect-based sentiment analysis for travel destinations. Specific goals were set as follows:

1. To increase contextual awareness of aspect terms and sentiment polarity, it was advised that the BERT framework incorporate knowledge from the tourism industry and be specifically tailored for that.
2. Evaluate the performance of the model compared to baseline models using metrics such as accuracy, F1-score, precision, and recall.

The novelty of the proposed KE-BERT model is that it includes external tourism-related data in the BERT architecture. This feature improves sentiment correspondence and provides more contextual awareness. KE-BERT enhances understanding of meaning by using tourism-related knowledge graphs and ontologies, unlike other models that rely only on pre-trained embeddings. Among the various contributions, the following stand out as particularly significant:

- The development of a BERT model that is suitable for the tourism industry and has knowledge specific to the sector.
- Contextual learning enhanced by knowledge augmentation to improve sentiment classification.

II. RELATED WORKS

Since ABSA can extract fine-grained sentiment insights from user-generated content, it has recently attracted a lot of attention. ABSA models relied mainly on lexicon-based techniques, where predefined word lists and polarity scores helped classify sentiment. However, the limitations of these methods stem from their inability to address context-dependent sentiment fluctuations and ambiguous aspect terms [10]. Learning sentiment patterns from labeled datasets has allowed machine learning models, such as Support Vector Machines (SVM) and Naive Bayes (NB) classifiers, to perform better than methods that rely on word lists. These models struggled with complex language structures and semantic nuances on casual social media platforms, which negatively affected their performance [11]. Sentiment analysis has effectively used Long Short-Term Memory (LSTM) networks to understand the order of words and the context of sentences [12]. However, LSTM-based models have limited capacity to scale to large datasets because they usually suffer from vanishing gradients and are computationally expensive.

In sentiment analysis, transformer-based models, especially BERT (Bidirectional Encoder Representations from Transformers), have evolved into the most creative approach. However, standard BERT lacks domain-specific knowledge, which restricts its capacity to precisely classify aspect terms in travel data. To improve performance in particular fields, many studies have proposed domain-adapted BERT models [13], such as BLURB [14] for medical data and FinBERT for financial data [15]. The K-BERT algorithm showed better results in classifying sentiment by adding structured knowledge [16]. Despite the potential of the travel industry to enhance sentiment analysis through domain-specific knowledge integration, it still underutilizes similar techniques.

Recent studies have revealed shortcomings in the way current ABSA models manage informal language and ambiguous aspect terms in tourism reviews. Most of the current models are based on general-purpose knowledge sources, which reduces their relevance to tourism content. The proposed KE-BERT model aims to bridge this gap. Incorporating tourism-specific knowledge into the BERT framework can improve aspect-based sentiment alignment and general classification accuracy.

III. PROPOSED METHOD

The proposed method combines domain-specific knowledge with the BERT architecture to enhance ABSA for travel destination reviews. The KE-BERT model leverages knowledge from tourism-specific ontologies and knowledge graphs to enhance contextual awareness of aspect terms and sentiment polarity. The model aims to manage the complicated and casual language often found in social media by matching outside knowledge with the token representations created by the BERT framework. A knowledge augmentation layer adds domain-specific semantic information to the transformer architecture. The BERT model is tuned on a tourism dataset, including external knowledge, using attention-based fusion. A SoftMax layer is used to classify the sentiment. This layer projects the polarity of sentiment for each individually discovered feature. The Adam optimizer is used for fine-tuning hyperparameters, and a cross-entropy loss function guides model training.




Fig. 1. Proposed process.

A. Data Collection and Preprocessing

1) Data Collection

On many well-known websites, including TripAdvisor, Yelp, Google Reviews, and Twitter, visitors routinely share their experiences about locations, lodging, food, transportation, and services. This study used the Social Media Sentiment Analysis Dataset [17], including 13 features and 700 observations related to location, date, and free-text reviews, including phrases that express sentiment and aspect. Table I presents a sample of the dataset.

TABLE I. SAMPLE OF COLLECTED DATA

Review text	Date	Country	Aspect terms	Sentiment label
"The hotel was clean, and the staff was very friendly."	2023-01-10	USA	hotel, staff	Positive
"Food was expensive and not worth the price."	2023-02-15	Canada	food, price	Negative
"Location was great but the room was small."	2022-03-05	UK	location, room	Neutral
"Excellent service but the wait time was too long."	2021-03-12	Australia	service, wait time	Mixed

2) Preprocessing

The data was cleaned to remove noise and standardize it for input into the KE-BERT model. The preprocessing phase involved the following.

a) Data Cleaning

- **HTML Tag Removal:** Regular expressions were used to remove HTML tags and metadata.
- **Stopword Removal:** Using a predefined list of stopwords, common stopwords, such as "the," "is," and "and," were removed, thus enhancing processing efficiency and reducing noise level.
- **Lowercasing:** To ensure consistency, every single bit of text was turned to lowercase.
- **Punctuation Removal:** Except for sentiment relevance, i.e., exclamation marks, special characters and punctuation marks were eliminated.

b) Tokenization

The BERT tokenizer helps tokenize the cleaned text to break it into individual words or subwords. WordPiece tokenization allows BERT to break out rare or challenging-to-pronounce words into smaller subword units. For example, for an input sentence "The hotel was clean, and the staff was very Friendly!", the tokens were: [the, hotel, was, clean, and, the, staff, was, very, friendly]

c) Data Labelling

Sentiment scores were calculated using lexicon-based techniques and contextual embeddings derived from BERT. Four categories were used to classify sentiment polarity: positive, negative, neutral, and mixed. For example:

"The hotel was clean, and the staff was very friendly." → Positive

d) Data Formatting for BERT

Attention masks, input tokens, and token type IDs were the last preprocessing steps before formatting data to fit BERT. For example, BERT input comes in the following structure:

[CLS] hotel clean, staff friendly [SEP]

where:

- [CLS] is the start token for classification

- [SEP] is a separator token between sequences

Attention masks identify which tokens are relevant for processing.

B. Aspect Extraction

This phase identifies notable aspects of travel destinations using social media platform reviews.

1) Part-Of-Speech (POS) Tagging and Dependency Parsing

POS tagging is the method of identifying potential aspect terms, nouns, and noun phrases. Usually, when trying to mix elements and emotions, one considers nouns (e.g., "hotel" and "food"), adjectives (e.g., "clean" and "expensive"), and verbs (e.g., "liked" and "enjoyed"). For example, in the sentence "The hotel was clean, and the staff was very friendly," POS tagging identifies

- Noun phrases: hotel, staff
- Adjectives: clean, friendly

and dependency parsing identifies the relationships:

- "clean" → modifies → "hotel"
- "friendly" → modifies → "staff"

2) Aspect Clustering

Semantic similarity facilitates the clustering of extracted features, enhancing the efficiency of the model and reducing noise. Two similar components can mix under a broader category, in this case, hospitality: "staff" and "service." The extent of semantic similarity between elements can be determined using cosine similarity between word embeddings. The clustering process is performed using:

$$S(A_i, A_j) = \frac{\vec{A}_i \times \vec{A}_j}{\|\vec{A}_i \times \vec{A}_j\|} \quad (1)$$

where $S(A_i, A_j)$ is the cosine similarity between aspect terms A_i and A_j , and \vec{A}_i and \vec{A}_j are the embedding vectors of aspect terms A_i and A_j . A threshold similarity value of 0.8 was used to determine whether two terms belong in the same cluster or not (see Table II). Once this threshold is crossed, aspect terms cluster with other terms.

TABLE II. ASPECT CLUSTERING

Original aspect	Clustered aspect	Similarity score
Hotel	Accommodation	0.85
Room	Accommodation	0.87
Staff	Hospitality	0.82
Service	Hospitality	0.88
Food	Dining	0.91

C. Knowledge Integration

Knowledge related to tourism from ontologies and knowledge graphs helps the aspect embeddings created by BERT, improving its understanding of how different aspects relate to sentiments.

1) *Extraction of Domain-Specific Knowledge*

Data specific to tourism is gathered from Open Tourism Knowledge Graph (OTKG), Wikidata, and DBpedia, among other sources. "Hotel → provides → Service" and "Restaurant → serves → Food" map ontology-based relationships to aspect terms. This approach lets the model establish contextual links between elements and emotions that go beyond the local text setting.

2) *Attention-Based Fusion*

The learned knowledge is elegantly included in the BERT model using an attention-based fusion mechanism. Through a weighted attention mechanism, the external knowledge embedding (K_i) produced by BERT is coupled with the aspect embedding (E_i) generated by BERT.

$$H_i = W_1 E_i + W_2 K_i \tag{2}$$

where H_i is the final integrated representation, E_i is the aspect embedding from BERT, K_i is the knowledge embedding, and W_1, W_2 are trainable weights adjusted during model training.

The degree of relevance of the knowledge for the context of the aspect determines the dynamically shifting scores of attention.

3) *Knowledge-Enhanced Sentiment Classification*

After passing through the embedding, the BERT transformer layers improve the knowledge of aspect sentiment alignment even more using multi-head attention. Improved embeddings let a SoftMax classifier forecast the polarity of the sentiment, that is, whether it is positive, negative, neutral, or mixed.

TABLE III. KNOWLEDGE-INTEGRATED OUTPUT

Review text	Extracted aspect	External knowledge	Predicted sentiment
"The hotel was clean and the staff was very friendly."	hotel, staff	Hotel → provides → Service	Positive
"Food was expensive and not worth the price."	food, price	Restaurant → serves → Food	Negative

D. *BERT Fine-Tuning*

The pre-trained BERT model was fine-tuned for the specific task of sentiment prediction for extracted aspects. BERT's transformer-based architecture helps it to identify intricate interactions found in textual data by including multi-head self-attention and positional encoding. Fine-tuning helps to change the last few layers of the BERT model so that it can be specialized for sentiment classification.

1) *Input Construction*

The review text and the aspect terms provide the BERT model's input. The BERT tokenizing mechanism directs the input structure as follows:

[CLS]Review Text[SEP]Aspect Term[SEP]

where [CLS] is a special token marking the beginning of the input, and [SEP] is a token separating the review text and the aspect term. For example, the sentence:

"The hotel was clean, and the staff was very friendly."

for the aspect term "hotel" is formatted as:

[CLS]The hotel was clean and the staff was very friendly.[SEP]hotel[SEP].

2) *Fine-Tuning with Knowledge Embedding*

After receiving the encoded input, BERT runs it through its transformer layers. BERT produces tokens with contextualized embedding. This embedding identifies the semantic link between the review text and the aspect term. At this point, shown by the HCLS, the last hidden state for the [CLS] token is regarded as the overall input representation. Combining the BERT contextualized embedding with the integrated knowledge embedding (K_i):

$$\hat{H} = [H_{CLS} || K_i] \tag{3}$$

where H_{CLS} is the final hidden state from BERT for the [CLS] token, K_i is the knowledge embedding from external sources. The concatenated representation ($||$ operator) is then passed through a dense layer with SoftMax activation to generate a prediction about the sentiment class (positive, negative, neutral, or mixed). The output layer is given by:

$$P(y) = \text{softmax}(WH_{final} + b) \tag{4}$$

where $P(y)$ is the predicted sentiment probability for each class, W is the trainable weight matrix, and b is the bias term.

3) *Loss Function and Backpropagation*

In training, the cross-entropy loss function reduces the variance between the expected and actual sentiment labels:

$$L = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C y_{ij} \log(P(y_{ij})) \tag{5}$$

where N is the number of training samples, C is the number of sentiment classes, y_{ij} is the ground truth label for i and class j , and $P(y_{ij})$ is the predicted probability for i and class j .

The Adam optimizer helps fine-tune the learning rate to prevent overfitting. This approach starts the learning rate low and progressively increases it in response to the computed gradients.

E. *Sentiment Classification*

After fine-tuning, the model will give every aspect a term taken from the review text sentiment polarity. The SoftMax layer on the neural network generates a probability distribution spanning the sentiment classes (positive, negative, neutral, and mixed). Table IV selects the most likely class to depict the expected attitude.

TABLE IV. SENTIMENT CLASSIFICATION

Review text	Extracted aspect	Knowledge integrated context	Predicted sentiment	Confidence score
"The hotel was clean and the staff was very friendly."	hotel	Hotel → provides → Service	Positive	0.92
"The food was overpriced and not tasty."	food	Restaurant → serves → Food	Negative	0.87
"The location was convenient, but the room was noisy."	location	Hotel → located at → City	Neutral	0.79

IV. RESULTS AND DISCUSSION

Training and evaluation were carried out in a high-performance computing environment equipped with an Intel Core i9-12900K processor, 64 GB of RAM, and an NVIDIA RTX 3090 GPU with 24 GB of VRAM. Python was used to run the simulation along with the TensorFlow and PARCH frameworks. To fine-tune the BERT model, the learning rate was set at 2e-5, the batch size at 16, and the maximum sequence length at 128 tokens. The Adam optimizer with weight decay was applied for optimization, and a three-epoch patience period helped prevent overfitting by using early stopping. Performance was compared with two existing BERT methods: ABSA-BERT and BERT-Sentiment-KG.

The proposed approach achieved better accuracy in all epochs, reaching 91.2%, outperforming the ABSA-BERT (85.0%) and BERT-Sentiment-KG (85.1%) in the tenth epoch. With 89.8% at epoch 10, the proposed approach achieved higher precision than ABSA-BERT (83.0%) and BERT-Sentiment-KG (85.1%). This demonstrates the improved ability of the model to spot cases of suitable sentiment free of false positives.

Comparing the proposed model at epoch 10 with ABSABERT (84.1%) and BERT-Sentiment-KG (86.0%) showed that it achieved a higher recall of 90.2%. This implies that the model can sufficiently detect cases of relevant sentiment, reducing false negatives. With an F1-score of 90.0% at epoch 10, the proposed model exceeded ABSA-BERT (83.6%) and BERT-Sentiment-KG (85.5%) by a noteworthy margin.




Fig. 2. Accuracy per epoch.

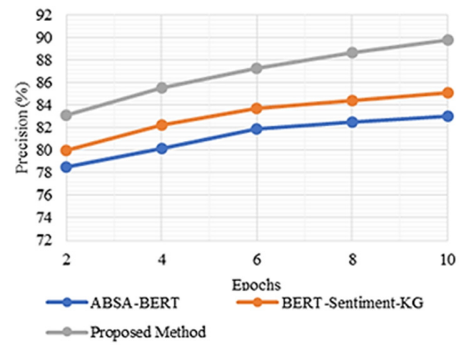


Fig. 3. Precision per epoch.

TABLE V. RECALL

Epoch	ABSA-BERT (%)	BERT-SentimentKG (%)	Proposed KE-BERT (%)
2	79.8	81.2	84.2
4	81.5	83.0	86.5
6	82.9	84.5	88.1
8	83.7	85.3	89.4
10	84.1	86.0	90.2

TABLE VI. F1-SCORE

Epoch	ABSA-BERT (%)	BERT-SentimentKG (%)	Proposed KE-BERT (%)
2	78.9	80.6	83.6
4	80.8	82.6	86.0
6	82.4	84.1	87.7
8	83.1	84.8	89.0
10	83.6	85.5	90.0

V. CONCLUSION

The proposed KE-BERT model for ABSA of tourism destinations showed notable improvements over current methods using data from social media platforms. KE-BERT maximizes semantic knowledge by including tourism-specific knowledge graphs and ontologies, unlike other models that depend just on pre-trained embeddings. Combining external information with BERT's understanding of context improved the model's accuracy, precision, recall, and F1-score. The model was able to better capture the subtleties of aspect-based sentiment using domain-specific knowledge, as KE-BERT preserves the sentence structure while injecting the domain knowledge. The proposed model outperformed ABSA-BERT and BERT-Sentiment-KG in all performance metrics, reaching an accuracy of 91.2%, a precision of 89.8%, a recall of 90.2%, and an F1-score of 90.0%. The fine-tuning method, along with aspect extraction and knowledge integration, helped the model be precise, allowing it to accurately determine sentiment polarities for each aspect.

REFERENCES

[1] F. Mirzaalian and E. Halpenny, "Exploring destination loyalty: Application of social media analytics in a nature-based tourism setting," *Journal of Destination Marketing & Management*, vol. 20, June 2021, Art. no. 100598, <https://doi.org/10.1016/j.jdmm.2021.100598>.

[2] B. Shin, S. Ryu, Y. Kim, and D. Kim, "Analysis on Review Data of Restaurants in Google Maps through Text Mining: Focusing on

- Sentiment Analysis,"*Journal of Multimedia Information System*, vol. 9, no. 1, pp. 61–68, 2022, <https://doi.org/10.33851/JMIS.2022.9.1.61>.
- [3] M. Hoang, O. A. Bihorac, and J. Rouces, "Aspect-Based Sentiment Analysis using BERT," in *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, Turku, Finland, Sept. 2019, pp. 187–196.
- [4] H. Li, B. X. B. Yu, G. Li, and H. Gao, "Restaurant survival prediction using customer-generated content: An aspect-based sentiment analysis of online reviews,"*Tourism Management*, vol. 96, June 2023, Art. no. 104707, <https://doi.org/10.1016/j.tourman.2022.104707>.
- [5] J. Hartmann, M. Heitmann, C. Siebert, and C. Schamp, "More than a Feeling: Accuracy and Application of Sentiment Analysis,"*International Journal of Research in Marketing*, vol. 40, no. 1, pp. 75–87, Mar. 2023, <https://doi.org/10.1016/j.ijresmar.2022.05.005>.
- [6] O. A. George and C. M. Q. Ramos, "Sentiment analysis applied to tourism: exploring tourist-generated content in the case of a wellness tourism destination,"*International Journal of Spa and Wellness*, vol. 7, no. 2, pp. 139–161, May 2024, <https://doi.org/10.1080/24721735.2024.2352979>.
- [7] C. Sun, L. Huang, and X. Qiu, "Utilizing BERT for Aspect-Based Sentiment Analysis via Constructing Auxiliary Sentence." arXiv, Mar. 22, 2019, <https://doi.org/10.48550/arXiv.1903.09588>.
- [8] B. Wan, P. Wu, C. K. Yeo, and G. Li, "Emotion-cognitive reasoning integrated BERT for sentiment analysis of online public opinions on emergencies,"*Information Processing & Management*, vol. 61, no. 2, Mar. 2024, Art. no. 103609, <https://doi.org/10.1016/j.ipm.2023.103609>.
- [9] M. M. Danyal, S. S. Khan, M. Khan, S. Ullah, F. Mehmood, and I. Ali, "Proposing sentiment analysis model based on BERT and XLNet for movie reviews,"*Multimedia Tools and Applications*, vol. 83, no. 24, pp. 64315–64339, July 2024, <https://doi.org/10.1007/s11042-024-18156-5>.
- [10] A. Booranakittipinyo, R. Y. M. Li, and N. Phakdeephrot, "Travelers' perception of smart airport facilities: An X (Twitter) sentiment analysis,"*Journal of Air Transport Management*, vol. 118, July 2024, Art. no. 102600, <https://doi.org/10.1016/j.jairtraman.2024.102600>.
- [11] I. Karabila, N. Darraz, A. EL-Ansari, N. Alami, and M. ELMallahi, "BERT-enhanced sentiment analysis for personalized e-commerce recommendations,"*Multimedia Tools and Applications*, vol. 83, no. 19, pp. 56463–56488, June 2024, <https://doi.org/10.1007/s11042-023-17689-5>.
- [12] N. A. Alabdulkarim, M. A. Haq, and J. Gyani, "Exploring Sentiment Analysis on Social Media Texts,"*Engineering, Technology & Applied Science Research*, vol. 14, no. 3, pp. 14442–14450, June 2024, <https://doi.org/10.48084/etasr.7238>.
- [13] A. Bello, S.C. Ng, and M.F. Leung, "A BERT Framework to Sentiment Analysis of Tweets,"*Sensors*, vol. 23, no. 1, Jan. 2023, Art. no. 506, <https://doi.org/10.3390/s23010506>.
- [14] S. Yu, J. Su, and D. Luo, "Improving BERT-Based Text Classification With Auxiliary Sentence and Domain Knowledge,"*IEEE Access*, vol. 7, pp. 176600–176612, 2019, <https://doi.org/10.1109/ACCESS.2019.2953990>.
- [15] D. Araci, "FinBERT: Financial Sentiment Analysis with Pre-trained Language Models." arXiv, Aug. 27, 2019, <https://doi.org/10.48550/arXiv.1908.10063>.
- [16] Y. Gu *et al.*, "Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing," *ACM Transactions on Computing for Healthcare (HEALTH)*, vol. 3, no. 1, July 2021, Art. no. 2, <https://doi.org/10.1145/3458754>.
- [17] "Social Media Sentiments Analysis Dataset ." Kaggle, Available: <https://www.kaggle.com/datasets/kashishparmar02/social-media-sentiments-analysis-dataset>.