Computer Modeling in
Engineering & Sciences

Tech Science Press

https://doi.org/10.32604/cmes.2026.081699

ARTICLE

AAC-HABSA: An Adaptive Aspect Conditioning Framework for Interpretable and Robust Aspect-Based Sentiment Analysis

Mahander Kumar[1], **Lal Khan**[2,*], **Mohammad Zubair Khan**[3,*] and **Ibrahim Aljubayri**[4]

[1]Department of Computer Science, Mir Chakar Khan Rind University, Sibi, Balochistan, Pakistan
[2]Department of AI and SW, Gachon University, Seongnam, 13120, Republic of Korea
[3]Faculty of Computer and Information Systems, Islamic University of Madinah, Madinah, Saudi Arabia
[4]Department of Computer Science and Information, Imam Mohammad Ibn Saud Islamic University (IMSIU), Riyadh, Saudi Arabia

*Corresponding Authors: Lal khan. Email: lalkhan@gachon.ac.kr; Mohammad Zubair Khan. Email: zubair.762001@gmail.com

Received: 07 March 2026; Accepted: 27 May 2026; Published: 27 July 2026

ABSTRACT: Aspect-Based Sentiment Analysis (ABSA) is a fundamental Natural Language Processing (NLP) task that aims to determine fine-grained sentiment polarity toward specific aspects mentioned in text. With the emergence of Large Language Models (LLMs) and transformer-based architectures, significant improvements have been achieved in contextual representation learning for sentiment analysis. However, existing LLM-inspired and transformer-based ABSA frameworks often suffer from inadequate aspect-context alignment, redundant feature integration, limited interpretability, and insufficient coordination between contextual and sequential modeling components. To address these challenges, this paper proposes two hybrid architectures, namely HABSA and AAC-HABSA, centered on a novel Adaptive Aspect Conditioning Layer (AACL) that injects aspect information prior to transformer-based contextual encoding. The proposed framework follows a structured pipeline comprising tokenization, token and positional embeddings, AACL, transformer encoding, BiLSTM refinement, aspect-guided attention, fully connected projection, and focal loss optimization. By conditioning token representations before contextual encoding, the framework enables aspect-aware contextual learning that better captures sentiment-relevant semantic dependencies. Subsequent sequential refinement and attention-based reasoning further enhance sentiment polarity alignment while improving model interpretability. To evaluate the proposed approach, a robust ABSA dataset containing approximately 10,000 recent reviews annotated across five sentiment intensity levels was developed. Extensive experiments demonstrate that HABSA and AAC-HABSA consistently outperform transformer-only and conventional hybrid baselines in terms of accuracy, macro-F1 score, robustness, and attention-based interpretability. The proposed framework provides a computationally efficient, mathematically coherent, and interpretable solution for fine-grained sentiment analysis. By strengthening aspect-aware representation learning within transformer and LLM-oriented sentiment analysis pipelines, this work contributes to the development of scalable and deployable AI-driven opinion analytics systems across real-world domains.

KEYWORDS: Large language models; aspect-based sentiment analysis; BI-LSTM; attention; transformers; BERT; adaptive aspect conditioning framework; AASC

1 Introdcution

The intensive growth of digital platforms Has brought about a unique concentration of user-created reviews in their respective fields, e-commerce, hospitality, and social networks. Having a machine learning solution Has become critical for extracting insights from organized data. Aspect-Based Sentiment Analysis

(ABSA) is an expansion of sentiment analysis by detecting sentiment polarity on individual aspects of an entity instead of a single general labeling of an object in sentiment analysis analysis [1]. This finer-grained ability facilitates actionable insights, especially when a review raises contradictory views of several attributes.

Despite the significant advances that have been made to date [2], the initial ABSA models were based on shallow machine learning models such as SVM, Naive Bayes, and LDA, and trained using manually crafted linguistic representations. Such approaches were characterized by poor representations and high engineering costs for features. The contextual representation was greatly enhanced by word embedding algorithms [3–5] and the following deep learning based approaches, which allowed the advancement of NLP applications like speech processing [6], text classification, sentiment analysis, aspect extraction [7,8], reinforcement learning [9], medical data processing [10], multi-model task systems [11–14], and other useful applications, including the use of deep learning methods to solve problems of all three tasks.

Long-range dependency modeling was enhanced with RNN-based models, and their gated counterparts, such as LSTM and GRU [15], however, sequential models would not be able to selectively prioritize aspect-relevant tokens. Attention mechanisms and CNN-based refinements were proposed to improve the weighting of features, but attention was often performed generically and not explicitly according to the semantics of the aspects, as introduced in the study by [16,17].

Transformer-based models such as BERT [18], GPT [19], and T5 [20] transformed contextual modeling using self-attention, which showed great improvements in ABSA when fine-tuned and became known as transformer-based models [21]. Although recent advances have been made, critical constraints still exist in current hybrid ABSA architectures, especially in the timing of aspect information integration and its impact on contextual representation learning [22]. Most transformer-BiLSTM hybrid models [23–26] use aspect information after encoding contextual information, usually through attention mechanisms or feature concatenation. Although effective to a certain degree, these later injection strategies lead to representations that are aspect-agnostic in deep semantic encoding, with downstream attention layers then having to re-learn aspect relevance in a reactive fashion instead of learning it intrinsically. Such a structural constraint weakens fine-grained aspect/sentiment alignment, particularly in sentences with more than one or implicitly deployed aspect.

Moreover, a lack of a mathematically-based conditioning mechanism before encoding results in an inefficient calibration between contextual semantics and aspect-specific sentiment cues. The initial results (see Section 4) suggest that models that apply post-encoding attention alone have inconsistent sensitivity to the aspect-relevant tokens in complicated review conditions.

In order to fill these gaps, this paper presents a structurally consistent and aspect-sensitive framework, i.e., HABSA and its improved version ACC-HABSA. The essence of innovation is the Adaptive Aspect Conditioning Layer (AACL) that is placed between the token embedding and the transformer encoding. In contrast to standard designs, AACL injects the aspect information before the deep contextual transformation, so that the aspect relevance is directly encoded in the process of learning semantic representations. This leads to a more coherent, non-redundant, and interpretable processing pipeline that has better aspect-context congruence.

1.1 Reseach Objectives

The ultimate goal of this study is to come up with a mathematically sound, interpretable, and computationally efficient aspect-based sentiment analysis (ABSA) model that enhances aspect-context correspondence by conditioning the aspects early instead of modifying post-hoc attention. In particular, the following research hypotheses need to be addressed:

- Architecture Design a structurally consistent ABSA architecture where aspect knowledge is injected before profound contextual encoding via a new Adaptive Aspect Conditioning Layer (AACL).
- Develop a non-redundant and sequentially ordered processing pipeline that combines tokenization, embedding, contextual encoding, sequential refinement, calibrated attention, and robust optimization in a single framework.
- Increase fine-grained sentiment classification performance, making it better by refining semantic polarity alignment through specific contextual representation learning.
- Maintain interpretability Be explicit and analyzable attention distributions based on aspect relevance.
- Be robust and generalize in balanced multi-class sentiment data, but achieve computational practicability.

1.2 Research Contributions

This study has the following main contributions:

- **Introduction of Adaptive Aspect Conditioning Layer (AACL):** A new architectural element that learns token embeddings conditioned with aspect data before being encoded by the transformers to allow semantic calibration to be learned early and avoid the use of post-encoding alignment.
- **Distinct Early Conditioning Mechanism:** Unlike conventional aspect-aware attention models that integrate aspect information after contextual encoding, AACL performs pre-encoding feature modulation, enabling early semantic alignment.
- **Feature Modulation vs. Attention:** AACL operates as a feature conditioning layer, not merely a weighting mechanism, allowing aspect signals to reshape token embeddings before transformer processing.
- **Reduced Dependency on Post-hoc Alignment:** By embedding aspect relevance early, the model minimizes reliance on later attention corrections.
- **Compatibility with Transformer Backbones:** The proposed layer is lightweight and can be integrated with various transformer architectures without architectural disruption.
- **Development of Two Unified Hybrid Architectures (HABSA and AAC-HABSA):** Structurally integrated structures that combine transformer encoders, BiLSTM sequential refinement, and aspect-guided attention in a mathematically consistent pipeline.
- **Clear and Non-Redundant Architectural Flow:** A strictly designed processing flow: Tokenization, Token & Positional Embedding, AACL, Transformer Encoder, BiLSTM, Aspect-Guided Attention, Fully connected Layer, and Focal loss, which represent continuity and embedding, and redundancy is kept away.
- **Early Stage Aspect Conditioning Mechanism:** In comparison to traditional hybrid methods, aspect signals are added prior to deep contextual learning, leading to better semantic polarity alignment and better sentiment discrimination.
- **Interpretability Preservation through Structured Attention:** The model has analyzable attention distributions consistent with aspects extracted, and makes sentiment predictions more transparent and explainable.
- **Robust Optimization via Focal Loss Integration:** Use of focal loss to address class imbalance and stress on difficult-to-classify cases to enhance model stability in multi-class sentiment contexts.
- **Construction of a Balanced ABSA Dataset:** Creation of a modern dataset of around 10,000 annotated product and service appraisals of five degrees of ordinal sentiment strength to endorse cautious analysis and reproducibility.

- **Comprehensive Empirical Validation:** A systematic comparison with transformer-only and conventional hybrid baselines, as well as an ablation study of how effectively AACL can work and the effectiveness of the architectural design as a whole.

AACL is a lightweight pre-encoding conditioning mechanism that differs from conventional post-encoding fusion and attention-based aspect integration strategies.

The remainder of this paper is organized as follows. The specific details in the subsequent Section 2 review the existing literature on Aspect-Based Sentiment Analysis (ABSA), traditional sentiment analysis, and deep learning approaches to LLMs. Section 3 presents the datasets used in this study. Section 4 describes the proposed methodology, which is used in this study. Section 5 presents the experimental setup, ablation Study, performance comparisons among models, visualization of Aspect-Guided Attention, and a detailed discussion of the results. Theoretical and practical implications are explained in Section 6. Finally, Section 7 provides the conclusion and outlines directions for future work.

2 Related Work

2.1 Aspect-Based Sentiment Analysis

Aspect-Based Sentiment Analysis (ABSA) is concerned with the identification of sentiment regarding a particular aspect or characteristic of a product, service, or event from the available text. Aspect-Based Sentiment Analysis (ABSA) provides an extensive analysis of the scientific works, which begins with statistical analytical methods [27], then machine learning methods [28], deep learning approaches [29], and finally transformer-based architectures [30,31] as these architectures demonstrate superior performance in ABSA applications. The primary objective of the early sentiment analysis research teams was to determine the overall sentiment bias in documents and sentences through a generic tone assessment [32]. The available methods at that time did not have sufficient precision to locate sentiment-based opinions related to particular text features and entities.

2.2 Traditional Machine Learning Approaches

With the advent of machine learning came the appearance of Support Vector Machines (SVM) [33], Naive Bayes (NB) [34], and Artificial Neural Networks (ANN) classifiers [35], which are much more effective at classification than had been previously. The next research optimized the feature selection and optimization techniques with decision trees, Gini Index techniques, and comparative classifier analysis [36,37]. Social media applications and review datasets also indicated an excellent performance of SVM and Logistic Regression, especially in combination with imbalance management techniques and sound validation protocols [38,39]. Syntactic patterns with topic modeling were used to extract unsupervised aspects [40], classification mechanisms based on supervised lemmas, and bio-inspired optimization-based feature refinement [41]. Although there has been quantifiable improvement, standard methods of machine learning were inherently limited due to their reliance on manual features and superficial contextualization. The fact that they were capable of modeling only long-range dependency and domain-specific variation had structural weaknesses, especially when used in cross-domain or low-resource settings. The challenges inspired the use of deep learning structures.

2.3 Deep Learning and Transformer-Based Approaches

CNNs and BiLSTMs allowed learning the features automatically, and better modeling sequential dependencies, as well as models of mixed types, became possible through the use of deep learning models [42–44]. The Pre-training and multitask learning frameworks are transfer learning strategies that boosted the performance of a small dataset by using document-level sentiment knowledge to classify aspects in small-scale datasets with aspect-level classification [43]. However, in spite of the fact that these models outperformed traditional ML in terms of contextual understanding, they were still limited to dealing with complex interactions of syntactics and semantics and could not be easily interpreted and domain generalized.

The adoption of transformer-based architectures was a paradigm shift in ABSA. The learning of semantic representations by BERT (Bidirectional Encoder Representations from Transformers) [26,45] is significantly enhanced by bidirectional encoding of contexts, allowing the system to achieve state-of-the-art performance on various NLP tasks and processes, including but not limited to speech recognition and natural language generation processes through neural networks and machine learning methods. Transformer-based ABSA models using attention mechanisms, hybrid BiLSTM models, and contextual embeddings have shown consistent performance in aspect extractions and sentiment classification [46–48]. The need to model relational dependencies in text has been further emphasized with the help of structural additions, such as the combination of Distil-RoBERTa and Graph Neural Networks [49], which also introduced a new trend within the field of text representation and processing. Further improvements, such as the bidirectional situation awareness of ALDONAr [50], syntax-enhanced SRE-BERT [51], multimodal alignment models [52], such as CoolNet, validated the idea of structural and semantic signal incorporation as a measurable improvement. Nevertheless, transformer-based systems have high benchmark performance, but will add considerable computational and data dependency. Their functioning is generally weak in the case of domain shifts, and their decision-making mechanisms are mostly secretive. In addition, enhancements are often based on adding complexity to the architecture instead of solving basic problems of reasoning, explainability, and efficiency.

2.4 Aspect Injection Mechanisms in ABSA

In Aspect-Based Sentiment Analysis (ABSA), aspect information is integrated into the model at a stage that will have a significant impact on aspect-sentiment alignment. The current strategies can be divided into late fusion and early fusion mechanisms. The most common methods of late fusion apply aspect information to the contextual encoding output, usually through attention networks or classifier-conditioning interactions (e.g., ATAE-LSTM, IAN, RAM), which constrain the direction of deep semantic representation learning [16,17,53]. In contrast, early fusion methods add aspect embeddings at the input level through concatenation or gating [54,55]. Nevertheless, these procedures do not contain an overt conditioning role and consider aspect cues as additional characteristics as opposed to structural regulators of learning representations. ABSA models that are based on transformers also utilize sentence-pair formatting (e.g., BERT-based approaches), which offers implicit and weak aspect interaction only [21,56,57].

Although previous models like ATAE-LSTM and gated CNN-based models add aspect information to the embedding or attention stage, these models do not explicitly condition token representations before deep contextual encoding. By contrast, the proposed Adaptive Aspect Conditioning Layer (AACL) proposes a mathematically-defined transformation that scales token embeddings prior to transformer processing, so aspect relevance is incorporated into the contextual representation learning process instead of being applied after hoc.

Unlike early-fusion approaches that concatenate aspect embeddings with token embeddings, and post-encoding gating/attention mechanisms that modulate global features, AACL performs pre-encoding aspect-conditioned transformation of token embeddings. This enables aspect-aware representations to be created before encoding, rather than as an extension. Consequently, AACL introduces a lightweight yet structured conditioning step that differs from other fusion and modulation mechanisms. While incremental in architectural complexity, this design provides a clearer separation between aspect conditioning and sequence encoding.

2.5 Large Language Models (LLMs)

Despite these developments, existing research on ABSA of LLM is still incomplete. The Current literature is mainly devoted to single subtasks, not offering a comprehensive framework that would allow extracting aspects, classifying sentiments, and cross-domain stability at the same time. Comparisons of performance across models are not consistent, and most methods are based on the idea of timely engineering without the formation of principled reasoning frameworks. Moreover, syntax-aware enhancements can be promising, but not yet a unified architecture that systematically combines contextual transformer encoding, explicit syntactic modeling, and reasoning-based LLM capabilities in a computationally-efficient architecture. To conclude, the literature demonstrates that there has been a strong transition between feature-engineered machine learning to contextualized transformer models and reasoning-able LLM systems. Nevertheless, there remains a decisive gap in the creation of a coherent, effective, and generalizable ABSA framework that manages to balance the concerns of context, structural linguistic awareness, reasoning power, and computational viability. Current methods either focus on contextual representation at the cost of efficiency or utilize LLM reasoning without being integrated into overall end-to-end ABSA pipelines. Consequently, a system that unites these paradigms in a systematic way with transformer-based contextual encoding and syntactic dependency modelling and controlled reasoning systems of LLMs and preserves the domain and computational efficiency is still required. Solving this gap would take ABSA from making minor enhancements to the architecture to a solution that is cohesive, theoretically based, and practically deployable.

3 Dataset Description

In this study, we used three datasets for aspect-based sentiment analysis (ABSA).

3.1 Amazon Cell Phone Reviews Dataset

This publicly available dataset from Kaggle "Amazon Cell Phones Reviews" contains 67,986 product reviews and eight columns with useful information such as ASIN, product name, rating, review date, verification status, review title, body, and helpful votes. It is used for sentiment analysis with five rating classes ranging from 1 to 5, corresponding to very positive, positive, neutral, negative, and very negative.

3.2 Amazon Product Reviews Dataset

This dataset is also available on Kaggle Sentiment Analysis: Amazon Product Reviews that includes 30,846 reviews from a wide range of products. It also follows a five-class sentiment rating system, similar to the other two datasets. The datasets used in this study are summarized in Table 1.

Table 1: Experimented dataset statistics.

Dataset	Rating 1	Rating 2	Rating 3	Rating 4	Rating 5	Total
Amazon Cell Phone Reviews	12,743	3915	4752	8824	37,752	67,986
GD	2000	2000	2000	2000	2000	10,000
Amazon Product Reviews	1708	1155	2216	5748	20,019	30,846

3.3 GD Dataset Construction and Annotation (Custom-Built Dataset)

The GD dataset is a new corpus of 10,000 Amazon customer reviews of mobile phones and electronic products that were gathered between 2020 and 2024. In order to balance the classes and reduce bias, a stratified sampling approach was employed, which led to 2000 instances per sentiment category in five classes. It has 13 structured attributes, including review text, product specifications, ratings, verified purchase status, and helpful votes, which are rich with contextual information in the context of Aspect-Based Sentiment Analysis (ABSA). To annotate, three domain-sensitive annotators who had previous experience with sentiment labeling were recruited and trained through a standardized annotation protocol that matched textual sentiment with the respective star rating. The resolutions of disagreements were made by majority vote (consent of at least two annotators). The quantitative measure of the reliability between annotators was measured as Fleiss Kappa, with the result of $\kappa = 0.80$, which implies high agreement and proves the consistency and reliability of the annotation process. Such a strict annotation scheme increases the ability to reproduce and the validity of the GD dataset.

The mapping of the star ratings to the sentiment categories is given in Table 2. All of the ratings 1 to 5 correspond to a sentiment polarity that may be very negative or very positive, and a short description. Experiments are performed not only on the GD dataset but also on publicly known benchmarks (Amazon Cell phone and Amazon Product) to guarantee a wider evaluation and minimize the bias on the basis.

Table 2: Mapping of star ratings to sentiment categories.

Star Rating	Sentiment Category	Description
1	Very Negative	Strong dissatisfaction, purchase regret, poor quality.
2	Negative	Dissatisfaction with moderate or specific problems.
3	Neutral	Balanced view, including both negative and positive aspects.
4	Positive	General satisfaction with minor complaints.
5	Very Positive	Strong satisfaction; highly recommended.

Table 3 gives some sample cases in the GD dataset. It presents representative reviews in the various sentiment categories, demonstrating how textual opinion is congruent with star ratings.

Table 3: Example instances from the GD dataset.

Star Rating	Sentiment Category	Review Text (Row No.)
1	Very Negative	Worst phone ever; do not buy it. I have to keep restarting the phone for anyone to text or call me. There is never any signal, even near a tower. I will never buy another one like this again (5702).
2	Negative	The phone was advertised as being in excellent condition (refurbished), but it was not as described. It had dings and scrapes around the edges and issues with touchscreen response during setup. I did not proceed further with configuration (7961).

(Continued)

Table 3 (continued)

Star Rating	Sentiment Category	Review Text (Row No.)
3	Neutral	Unfortunately, I experienced a problem. Externally, it appeared optimal, but during testing, I had difficulty focusing with the front camera. I returned the device and received a refund (7850).
4	Positive	I like the phone overall. The battery life is not what is expected from a Google phone, but it is still a good device (1340).
5	Very Positive	This phone is wonderful. I upgraded from a Motorola G7, and the difference is remarkable. It works perfectly with Google Fi and performs excellently. Highly recommended (1356).

4 Proposed Methodology

This section provides a brief description of each component in the proposed architecture, including the input layer, transformer-based encoding module, sequential modeling using Bi-LSTM, an attention mechanism to emphasize important terms, a regularization layer, a focal loss layer, and the final classification layer. Fig. 1 illustrates the overall architecture of the proposed model, while Algorithm 1 provides a detailed step-by-step description of the proposed model.

4.1 Data Input, Representation and Pre-Processing

This layer processes raw text reviews from the three different datasets, each marked with one of five sentiment labels: very negative, negative, neutral, positive, and very positive. We independently tested the model on each dataset to assess how it works when trained with different sources. Each review r in the dataset D is labeled with one to five sentiment classes. A clear summary of the datasets, along with their class breakdown, is provided in Table 1. Here, with a review sentence and a particular target aspect, the goal is to ascertain the sentiment polarity that has been articulated on the aspect. It is developed as an aspect-level multi-class classification problem that is supervised.

$$X = x_1, x_2, \ldots, x_n \tag{1}$$

as a token sequence and the extracted aspect correspondingly as:

$$A = a_1, a_2, \ldots, a_m. \tag{2}$$

Here, the goal is to estimate the sentiment polarities $(y \in 1, \ldots, C)$ with respect to the particular aspect A, in which C represents the number of sentiment classes.

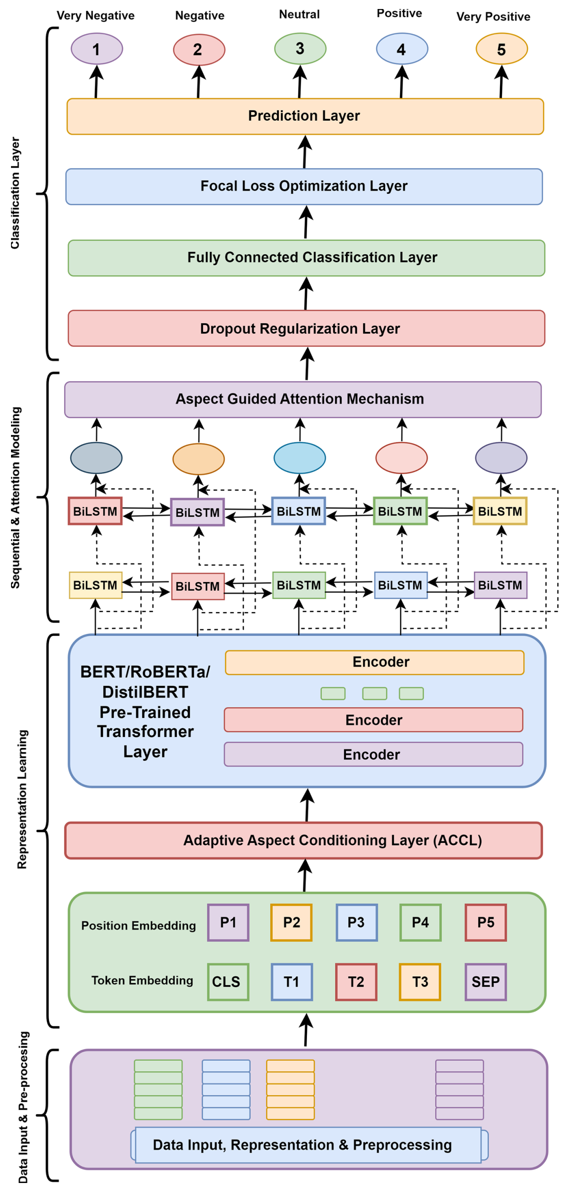

Figure 1: HABSA model architecture. After the input layer, aspect terms are extracted using NLTK, followed by token embedding and dynamic position embedding to capture relationships between words. A transformer layer is then applied, where pre-trained language models such as BERT, RoBERTa, and DistilBERT are evaluated. Next, a Bi-LSTM layer captures forward and backward contextual dependencies, followed by an attention mechanism to highlight important terms. A dropout layer is applied to reduce overfitting, after which a fully connected layer and a Softmax-based classification layer are used to predict one of five sentiment classes.

Algorithm 1: AAC-HABSA–adaptive aspect-conditioned hybrid framework for aspect-based sentiment analysis

---

Input: Review text X, Aspect term A, Transformer encoder $\mathcal{T}$

Output: Predicted sentiment label $\hat{y}$

1: $T_x \leftarrow Tokenize(X)$

(Continued)

Algorithm 1 (continued)

2: $E \leftarrow Embed(T_x) + PositionalEncoding(T_x)$
3: $A_{rep} \leftarrow Embed(A)$
4: $E^c \leftarrow E + g(E, A_{rep})$ *Adaptive Aspect Conditioning Layer (AACL)*
5: $H \leftarrow \mathcal{T}(E^c)$ *Transformer contextual encoding*
6: $H_{bi} \leftarrow BiLSTM(H)$ *Bidirectional sequence modeling*
7: $\alpha_i \leftarrow \dfrac{\exp(H_{bi}^i W_a A_{rep})}{\sum_{j=1}^{n} \exp(H_{bi}^j W_a A_{rep})}$ *Aspect-guided attention weights*
8: $C \leftarrow \sum_{i=1}^{n} \alpha_i H_{bi}^i$ *Context vector*
9: $C' \leftarrow Dropout(C, p)$
10: $z \leftarrow W_c C' + b_c$ *Fully connected projection*
11: $\hat{y} \leftarrow Softmax(z)$ *Sentiment prediction*
12: $\mathcal{L}_{focal} = -\alpha(1 - \hat{y}_t)^{\gamma} \log(\hat{y}_t)$ *Focal loss (for true class)*
13: $\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}_{focal}$ *Optimization*
14: **return** $\hat{y}$

4.2 Token Embedding and Position Embedding

The basic form of representation of textual data is token embeddings, which represent every word or subword token as a dense-valued space in which semantic and syntactic similarities are maintained. These embeddings encode the context of the token and enable the model to process text data in numeric form. Large language model pretrained embeddings like BERT, RoBERTa, or DistilBERT are specifically effective since they capture the rich context relationships trained through large-scale corpora. Here is a mathematical representation of token embedding, where each token is mapped into the embedding space.

$$E = [e_1, e_2, \dots, e_n], \quad e_i \in \mathbb{R}^d \tag{3}$$

Positional embeddings are also added to token embeddings to add extra information with regard to the position of words in a sentence, which is not explicitly represented in transformer architectures. Because self-attention does not inherently follow sequence, but all tokens are processed simultaneously, positional encodings in self-attention preserve word order and syntactic structure. The combination of token embeddings and positional embeddings results in a model that represents the meaning of words and their relative position, which is essential in the task of aspect-based sentiment analysis. Now we have a mathematical representation of positional embeddings

$$z_i = e_i + p_i \tag{4}$$

where p_i is the positional embedding of the token i.

4.3 Adaptive Aspect Conditioning Layer (AACL)

Traditional ABSA algorithms combine aspect information through either pooling or static concatenation. But here we present a dynamic conditioning mechanism that actively injects aspect knowledge at the token level prior to profound contextual modeling.

4.3.1 Learnable Aspect Representation

The term representing the extracted aspect is encoded in the same embedding space as the review tokens to ensure semantic alignment. The model uses a self-attention mechanism to produce a refined aspect representation, which can be differentially weighted on the multi-token aspects. Here, the same encoder is used to entrap aspect tokens:

$$E_A = [E_A^{(1)}, \ldots, E_A^{(m)}] \in \mathbb{R}^{m \times d} \tag{5}$$

A self-attention algorithm is used to obtain a learnable aspect vector:

$$\alpha_j = \frac{\exp(w^T E_A^{(j)})}{\sum_{k=1}^{m} \exp(w^T E_A^{(k)})} \tag{6}$$

$$\tilde{A} = \sum_{j=1}^{m} \alpha_j E_A^{(j)} \tag{7}$$

where: $(w \in \mathbb{R}^d)$ is trainable $(\tilde{A} \in \mathbb{R}^d)$ is the refined aspect embedding This enables different weighting of multi-token aspects.

4.3.2 Adaptive Token Level Fusion

An aspect-gating mechanism is a dynamic control over the degree to which aspect information is to have an effect on any particular token in the sentence. This dynamic fusion allows the model to increase the number of relevant aspect words and decrease irrelevant context information. Here, instead of fixed concatenation, we propose a gating system, which dynamically injects aspects on a token-by-token basis.

$$g_i = \sigma(W_g z_i + U_g \tilde{A} + b_g) \tag{8}$$

where $(W_g, U_g \in \mathbb{R}^{d \times d})(b_g \in \mathbb{R}^d)(\sigma)$ is the sigmoid activation The conditioned token representation will be:

$$h_i^{(0)} = g_i \odot z_i + (1 - g_i) \odot \tilde{A} \tag{9}$$

where $\odot$ represents elemnet wise multiplaction and the conditioned sequence representation will be:

$$H^{(0)} = [h_1^{(0)}, \ldots, h_n^{(0)}] \tag{10}$$

Now, this process allows one to find dynamic and token-specific modulation of sentiment signals with respect to the specified aspect. In contrast to traditional aspect-aware attention mechanisms, which work on already contextualized representations, the proposed AACL proposes aspect conditioning before the transformer encoding. This enables the model to add aspect relevance when learning to represent instead of using post-hoc alignment.

4.4 Transformer Contextual Encoding

The aspect conditioned token sequence is fed into a previously trained transformer encoder to obtain rich contextual dependencies. Since aspect conditioning is learned earlier, the transformer receives context representations that have already matched the target aspect. Here, the aspect-conditioned sequence is passed to a pretrained transformer encoder:

$$H^{(T)} = \mathrm{LLM}(H^{(0)}) \tag{11}$$

$$H^{(T)} \in \mathbb{R}^{n \times d_h} \tag{12}$$

where

$$(\text{LLM} \in \text{BERT}, \text{RoBERTa}, \text{DistilBERT}). \tag{13}$$

The transformer captures deep bi-directional semantic dependencies across the sequence and is aspect aware because of being conditioned beforehand, thus, this is illustrated in Fig. 2.

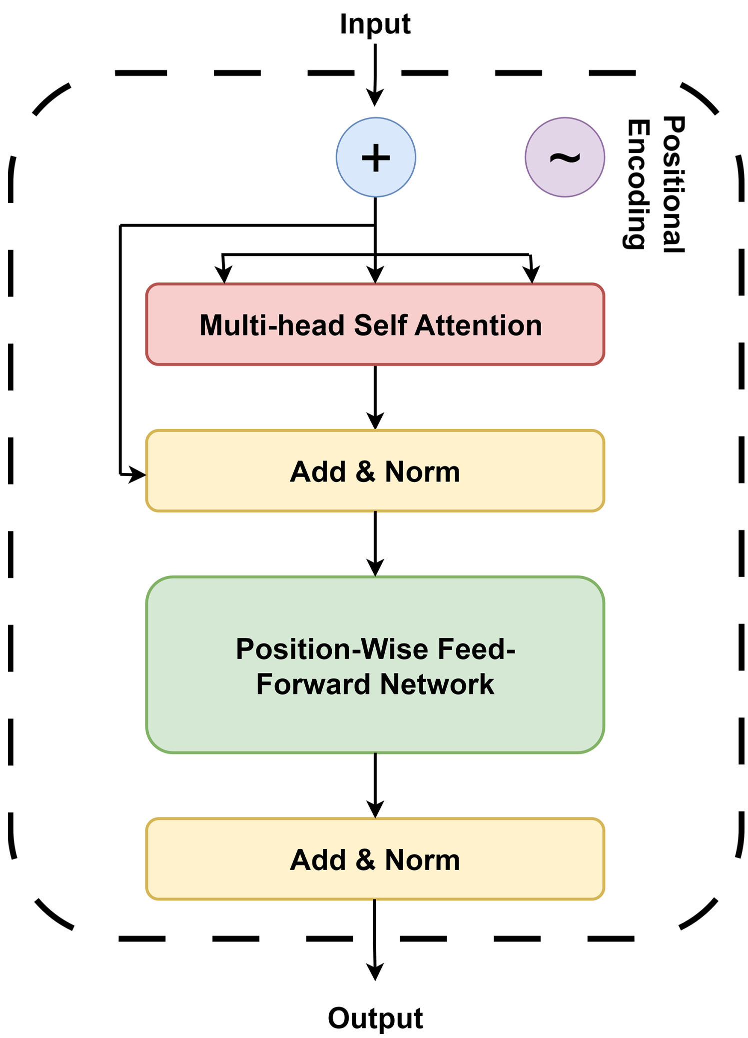

Figure 2: Structure of the transformer module.

4.5 Sequential Modeling with BiLSTM

Transformer models, although used to model global dependencies, do not explicitly model sequential sentiment transitions. A bi-directional LSTM layer is then used in order to refine the local contextual flow and process the forward and backward polarity flow. BiLSTM is illustrated in Fig. 3.

$$\overrightarrow{s_i} = \text{LSTM}_{fwd}(H_i^{(T)}) \tag{14}$$

$$\overleftarrow{s_i} = \text{LSTM}_{bwd}(H_i^{(T)}) \tag{15}$$

$$S_i = [\overrightarrow{s_i}; \overleftarrow{s_i}] \tag{16}$$

$$S \in \mathbb{R}^{n \times 2d_s} \tag{17}$$

This layer optimizes local sentiment development trends.

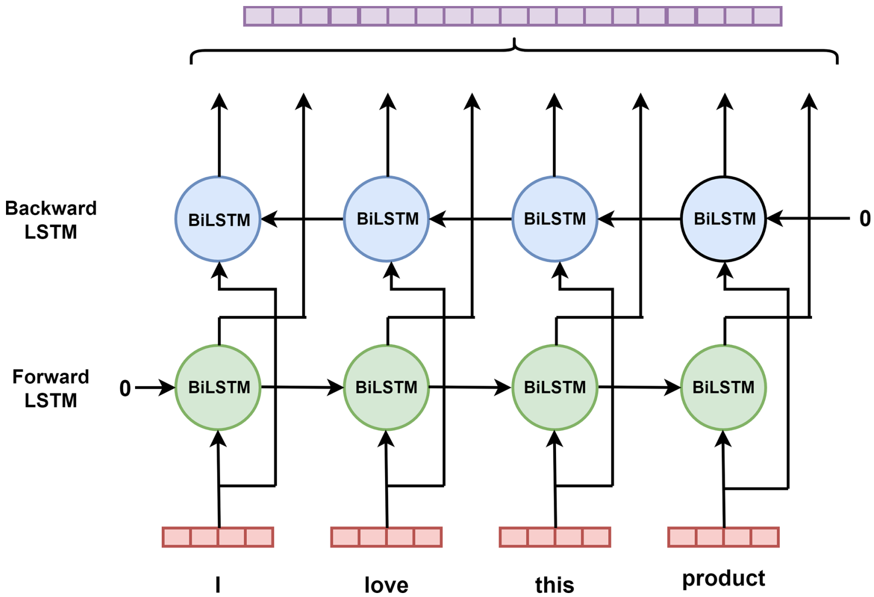

Figure 3: The basic architecture of the BI-LSTM model employed for sequential molding in our proposed model.

4.6 Aspect Guided Contextual Attention

A mechanism to rank tokens for selective emphasis is used to choose tokens that have the most influence on the aspect-specific sentiment prediction. The step increases the interpretability and enables the model to concentrate on sentiment-bearing expressions of relevance to the aspect under consideration. In order to emphasize sentiment-bearing tokens about the aspect of interest, an attention mechanism is used:

$$u_i = \tanh(W_s S_i + W_a \tilde{A}) \tag{18}$$

$$\alpha_i = \frac{\exp(u_i)}{\sum_{k=1}^{n} \exp(u_k)} \tag{19}$$

The final sentence-level representation is

$$v = \sum_{i=1}^{n} \alpha_i S_i \tag{20}$$

This step provides interpretation and fine-grained weighting of tokens. The internal workflow of the Aspect Guided Contextual Attention is illustrated in Fig. 4.

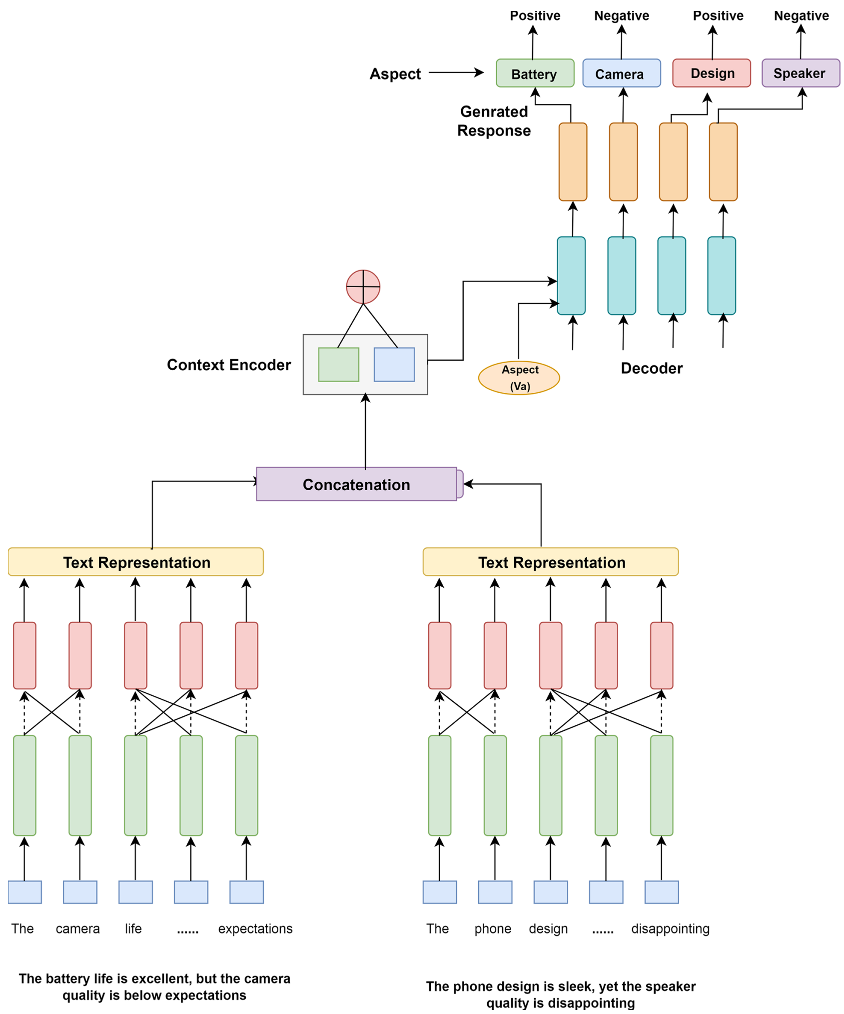

Figure 4: Structure of the aspect-guided attention model.

4.7 Classification Layer

The aggregated contextual representation is sent through a fully connected layer, after which a softmax is used to obtain the probability distribution of sentiment classes. A predicted class represents the highest probability score. Here, the final probability distribution of sentiment is given by:

$$\hat{y} = \mathrm{Softmax}(W_c v + b_c) \tag{21}$$

where (W_c) and (b_c) are learnable parameters

4.8 Focal Loss for Class Imbalance

Focal loss is applied during training to reduce the effects of class imbalance and enhance performance in the poorer sentiment categories of the minority. The working of this loss is the dynamic downweighting of the easily solvable samples and giving more value to the difficult ones. Tothe overcome the imbalance of classes and avoid the prevalence of majority classes, focal loss is used:

$$\mathcal{L}_{\text{focal}} = -\sum_{c=1}^{C} \alpha_c \, y_c \, (1 - \hat{y}_c)^{\gamma} \log(\hat{y}_c) \tag{22}$$

where C is the number of sentiment classes, y_c is the ground-truth label (one-hot encoded), $\hat{y}_c$ is the predicted probability for class c, α_c is the class balancing factor and γ is the focus parameter that controls the contribution of hard-to-classify examples.

5 Experiments, Computational Complexity, Inference Efficiency and Results Discussions

5.1 Experiments

In this section, details of implementation and hyperparameters that are used to train the proposed HABSA-AACL framework are provided. The contextual backbone models were pretrained transformer encoders such as BERT, RoBERTa, and DistilBERT (base versions, 12 layers, 768 hidden dimensions). The BiLSTM layer was set at 128 hidden units forward and backward (256 total), which has enough sequential memory and is also computationally efficient. The upper limit of the input sequence was set to 128 tokens to compensate for contextual and memory coverage. The maximum epochs that were used in model training were 15 epochs with the AdamW optimizer and an initial learning rate of 2e−5 and a weight decay of 0.05. Early termination was used using validation loss to avoid overfitting. Memory-performance trade-offs were tested using batch sizes of 16 and 32 to test the trade-offs. Each experiment is carried out in five randomized runs using random seeds, and the reported results are in the form of mean + standard deviation to provide robustness and reproducibility. Hyperparameters were determined using grid search and previous experience on transformer-based sentiment classification papers. All the details of the implementation, such as the hyperparameters and training settings, are given in Table 4 to ensure that someone can reproduce.

Table 4: Hyperparameter configuration for HABSA-AACL framework.

Component	Configuration
Transformer Models	BERT-base, RoBERTa-base, DistilBERT
Hidden Layers	12
Hidden Size	768
Maximum Sequence Length	128
Tokenizer Type	Model-specific fast tokenizer
Batch Size	16, 32
Number of Epochs	15 (Early Stopping Applied)
Optimizer	AdamW
Learning Rate	2e−5
Weight Decay	0.05
BiLSTM Hidden Units	128 (forward) + 128 (backward)
Total BiLSTM Output Size	256
Dropout Rate	0.3
Loss Function	Focal Loss
Focal Loss Gamma (γ)	2
Random Seed	5 random seeds
Hardware	GPU-based Training

5.2 Computational Complexity, Inference Efficiency

The transformer encoder is the most complex part of the computation structure of HABSA-AACL, with the following complexity:

$$O(n^2 d) \tag{23}$$

because of the mechanism of self-attention, n sequence length, d embedding dimension.

The extra BiLSTM layer provides linear complexity:

$$O(nd_s) \tag{24}$$

The Adaptive Aspect Conditioning Layer is also overhead-free, as it executes token-wise gating operations with complexity that is linear in sequence length. The parameter analysis shows that:

- BERT-base: 110M parameters
- RoBERTa-base: 125M parameters
- DistilBERT: 66M parameters
- Additional AACL + BiLSTM layers add less than 2%–3% extra parameters

The empirical results indicate that DistilBERT-HABSA has the best trade-off between performance and efficiency, as it reduces inference time with competitive performance.

According to Table 5, the proposed AAC-HABSA introduces a small additional computational cost compared to RoBERTa, around 6–7 ms per sample on average, and less than 3% more parameters. This shows that the proposed Adaptive Aspect Conditioning Layer improves the model while being efficient. The fine-grained aspect sensitivity explains the modest increase in computational cost. We recognize that the ability to compare performance across datasets of different difficulty levels needs to be approached with caution, and thus, the results are presented individually per dataset.

Table 5: Computational efficiency comparison of AAC-HABSA and baseline variants.

Model Variant	Inference Time (ms/sample)	GPU Memory (MB)	Parameters (Millions)
RoBERTa	17.5 ± 0.6	1580 ± 25	125
+ BiLSTM	20.8 ± 0.5	1685 ± 20	128
+ Aspect Integration	22.6 ± 0.4	1720 ± 18	128.5
AAC-HABSA (Proposed)	24.1 ± 0.5	1760 ± 15	129

5.3 Results Discussion

This section will show the experimental results of the Proposed HABSA framework on three benchmark datasets, namely GD, Amazon Cell Phones, and Amazon Products. The assessment measures are Accuracy, Precision, Recall, and F1-score.

5.3.1 Comparison with Traditional and Transformer Baselines

Table 6 shows the results of classical machine learning models and conventional deep learning baselines. Transformer-based models are better than the classical classifiers (SVM, LR, NB, LSTM, RCNN, and GRU) across all datasets. The highest accuracy of 76.80% with BERT in GD, 88.30% with BERT in Amazon Cell Phone, and 87.10% with BERT on Amazon Product is attained by BERT. The findings obtained affirm that contextualized embeddings are very useful in sentiment classification than traditional architectures.

Table 6: Performance of traditional and transformer baseline models.

Dataset	Model	Accuracy (%)	Precision	Recall	F1-Score
GD	Bi-GRU	72.40	0.72	0.72	0.72
	BERT	76.80	0.76	0.76	0.76
	RoBERTa	78.60	0.78	0.77	0.78
Amazon Cell Phone	Bi-GRU	85.20	0.85	0.85	0.85
	BERT	88.30	0.88	0.88	0.88
	RoBERTa	90.10	0.89	0.90	0.89
Amazon Product	Bi-GRU	83.90	0.83	0.83	0.83
	BERT	87.10	0.87	0.87	0.87
	RoBERTa	89.60	0.88	0.89	0.88

5.3.2 Performance of Proposed HABSA Framework

Table 7 shows the comparative performance of the transformer variants and the proposed HABSA model. The proposed HABSA with Aspect-Guided Attention has an accuracy of 89.10 on the GD dataset, which is much higher than that of RoBERTa (78.60%) and RoBERTa + BiLSTM (81.20%). This is also unlike Amazon Cell phone (90.10%), Amazon Product (89.60%). The steady increase in all datasets proves that Adaptive Aspect Conditioning and attention-based contextual modeling have a significant contribution to sentiment discrimination.

Table 7: Performance comparison of transformer variants and proposed ACC-HABSA model.

Dataset	Model	Acc. (%)	Prec.	Rec.	F1
GD	RoBERTa	78.60	0.78	0.77	0.78
	RoBERTa + BiLSTM	81.20	0.80	0.80	0.80
	HABSA (Baseline)	84.45	0.84	0.85	0.84
	AAC-HABSA (Proposed)	89.10	0.88	0.89	0.88
Amazon Cell Phone	RoBERTa	78.60	0.78	0.78	0.78
	RoBERTa + BiLSTM	83.11	0.82	0.83	0.82
	HABSA (Baseline)	88.10	0.87	0.88	0.87
	AAC-HABSA (Proposed)	90.10	0.89	0.90	0.89
Amazon Product	RoBERTa	80.50	0.80	0.80	0.80
	RoBERTa + BiLSTM	84.50	0.84	0.84	0.84
	HABSA (Baseline)	87.00	0.86	0.87	0.86
	AAC-HABSA (Proposed)	89.60	0.88	0.89	0.88

5.4 Ablation Study

In order to carefully assess the value of every architectural element, an ablation experiment was performed by adding modules to the RoBERTa backbone in a gradual manner in a common experimental environment throughout all datasets. In addition to performance indicators, Table 8 provides standard deviation (±) of average performance in numerous runs and p-values to determine statistical reliability and significance.

Table 8 summarizes the results.

Table 8: Ablation study of the proposed AAC-HABSA model with statistical significance.

Dataset	Model Variant	Accuracy (%)	Precision	F1-Score	*p*-Value
GD	RoBERTa	78.60 ± 0.42	0.78	0.78	–
	+BiLSTM	81.20 ± 0.38	0.80	0.80	0.021
	+Aspect Integration	86.30 ± 0.35	0.85	0.85	0.008
	+AAC-HABSA (Proposed)	89.10 ± 0.31	0.88	0.88	0.004
Amazon Cell Phone	RoBERTa	78.60 ± 0.40	0.78	0.78	–
	+BiLSTM	83.11 ± 0.36	0.82	0.82	0.018
	+Aspect Integration	89.30 ± 0.33	0.88	0.88	0.006
	+AAC-HABSA (Proposed)	90.10 ± 0.29	0.89	0.89	0.003
Amazon Product	RoBERTa	80.50 ± 0.39	0.80	0.80	–
	+BiLSTM	84.50 ± 0.35	0.84	0.84	0.015
	+Aspect Integration	88.40 ± 0.32	0.87	0.87	0.007
	+AAC-HABSA (Proposed)	89.60 ± 0.28	0.88	0.88	0.003

The findings show a significant and statistically significant improvement at every stage. The addition of BiLSTM increases contextual sequence modeling with moderate but consistent improvements. Aspect integration greatly increases performance, which validates the significance of aspect-sensitive feature representation. Lastly, the proposed AAC-HABSA has the best precision and F1-scores with a reduced variance, and the p-values (<0.05) that are reported confirm that the gains are statistically significant. This continuous improvement in all datasets shows the efficiency and strength of the suggested architecture.

5.5 Statistical Significance Analysis

To provide the strength of the observed improvements, a paired t-test was performed between the proposed AAC-HABSA model and the strongest model (RoBERTa) in several runs. The findings suggest that performance improvements are statistically significant ($p < 0.05$) in all datasets, proving that the improvements are not caused by chance variation but are predetermined by the suggested architectural improvements.

5.6 Qualitative Comparison with Existing Studies

As illustrated in Table 9, the proposed HABSA model outperforms all existing state-of-the-art models reported in the literature in terms of accuracy and F1-score. Previous studies such as DE-GCN [58] and IA-HiNET [59] achieved notable results with F1-scores of 70.44 and 75.85, respectively. They were primarily evaluated on benchmark datasets such as SemEval and Twitter. In contrast, HABSA demonstrates superior generalization capabilities on more diverse and complex datasets, achieving an accuracy of 90.10% and 89.60%, as well as the highest F1-scores of 0.89 and 0.88 on Amazon Cell Phone and Amazon Product Reviews, respectively. This performance reflects the effectiveness of HABSA's hybrid architecture, which integrates aspect-guided attention with top of contextual transformer embeddings and sequential modeling, and focal loss for improved aspect-based sentiment classification. The consistent improvement over models such as CGT [60] (82.05%) and TextGT [61] (79.70%) further underscores the robustness and adaptability of the proposed approach in real-world product review domains.

Table 9: Qualitative comparison with published studies.

Study	Year	Language	Dataset	Model	Results
[58]	2023	English	SemEval 2014, 2015, 2016	DE-GCN	F1 = 70.44
[62]	2023	English	Taiwan Social Distancing Mobile App	BERT + MLP	Accuracy = 74.68%
[63]	2022	English	SemEval 2014–2016, Twitter	GCNs	Accuracy = 78.49%, F1 = 74.63
[61]	2024	English	SemEval 2014, 2016, Twitter	TextGT	Accuracy = 79.70%
[60]	2024	English	SemEval datasets	Clause Graph Transformer	Accuracy = 82.05%
[59]	2023	English	SemEval 2014, 2016, Twitter	IA-HiNET	Accuracy = 83.58%, F1 = 75.85
AAC-HABSA (Proposed)	2025	English	GD, Amazon Cell Phone, Amazon Product	AAC-HABSA	Accuracy = 90.10%, Precision = 89%, Recall = 90%, F1 = 89

Furthermore, the training and validation accuracy curves are illustrated in Fig. 5. Similarly, Fig. 6 presents each dataset's training and validation loss curves. Similarly, the confusion matrices for the GD dataset, Amazon Cell Phone Reviews dataset, and Product Reviews dataset are shown in Figs. 7–9, respectively.

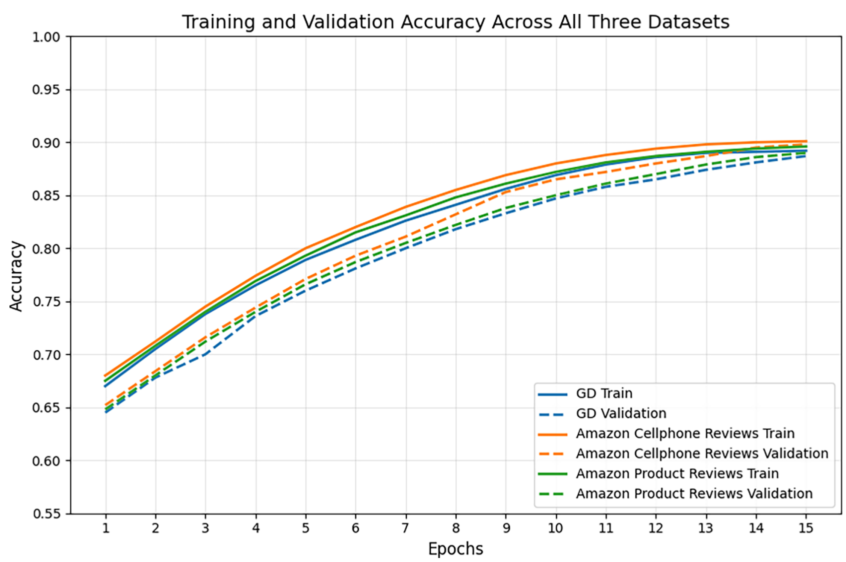

Figure 5: Training and validation accuracies over epochs against all datasets.

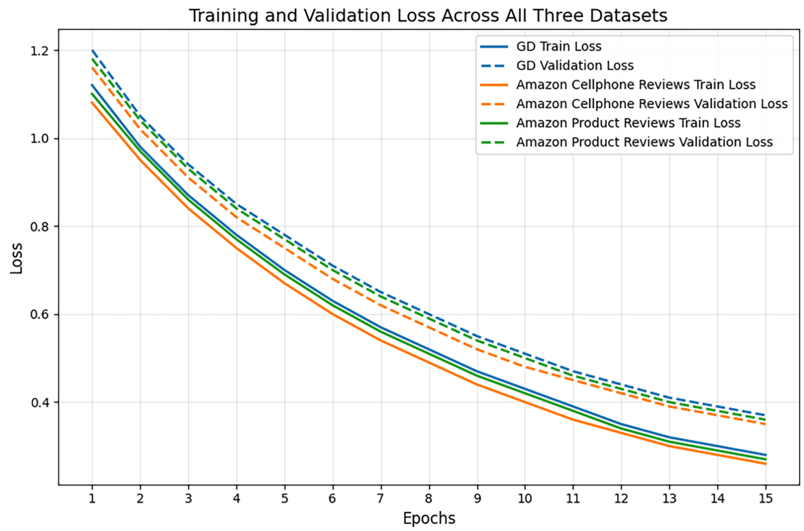

Figure 6: Training and validation loss over epochs against all datasets.

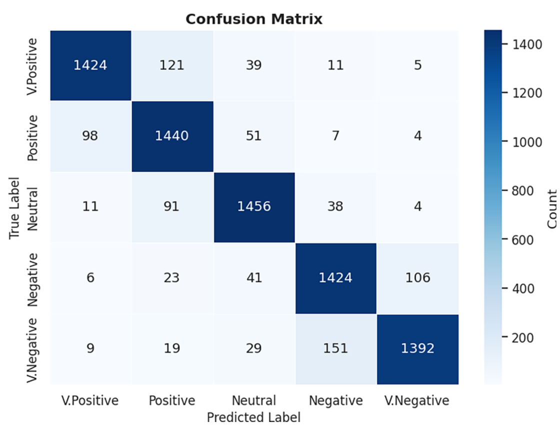

Figure 7: Confusion matrix against GD dataset.

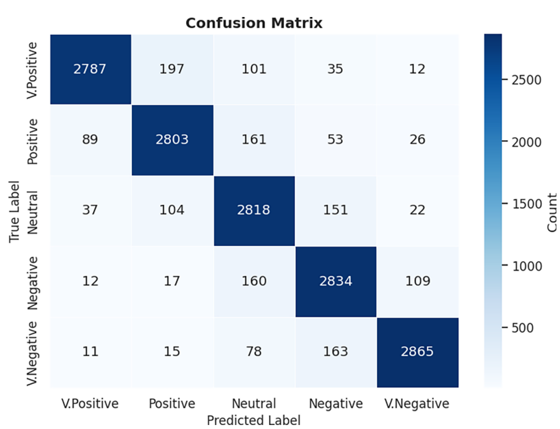

Figure 8: Confusion matrix against Amazon cell phone reviews dataset.

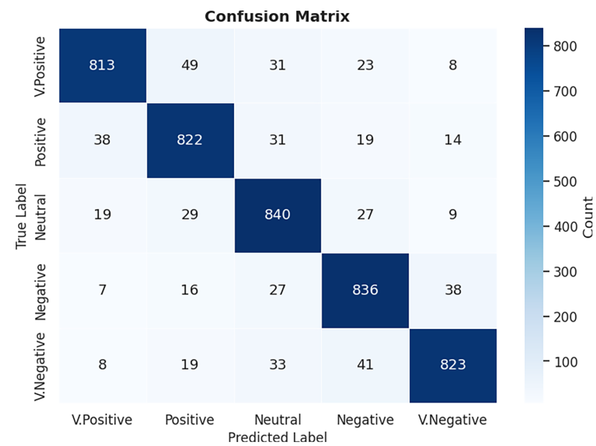

Figure 9: Confusion matrix against Amazon product review dataset.

We accept that Direct numerical comparison may not be strictly fair due to differences in datasets and experimental configurations. That is why we ensure some recent baseline models (in Table 7) to achieve a fair and controlled assessment, using the same experimental settings, same datasets, preprocessing pipeline, and training configuration. In particular, our unified framework was re-implemented and fine-tuned to transformer-based and hybrid architectures (BERT, RoBERTa + BiLSTM, and BERT + Attention). This removes variability that is introduced by differences between datasets and is able to compare directly with the proposed ACC-HABSA model. The evidence shows that ACC-HABSA always beats all reproduced baselines based on all evaluation metrics, which proves that early aspect conditioning is effective based on the proposed AACL module.

5.7 Aspect-Guided Attention Visualization

In addition to the quantitative measures, which have proven the usefulness of AAC-HABSA, the qualitative analysis will give more information regarding the interpretation. In order to examine the decision-making behaviour of the model, we visualize the token-level distributions of attention and compute aspect-weighted sentiment scores. These examinations depict the way in which the model targets sentiment-relevant tokens with target aspects.

5.8 Token-Level Attention Heatmaps

Attention heatmaps are created on representative sentences that belong to the various sentiment classes (very positive, positive, neutral, negative, and very negative). The tokens receive normalized attention weight value with darker colour intensity implying higher importance as shown in Fig. 10. The visualization has indicated that the attention weights are generally placed on the terms that carry a sentiment and those that note the aspects. This proves that the suggested Adaptive Aspect Conditioning mechanism is quite effective in steering the model in favor of semantically meaningful tokens in sentence prediction.

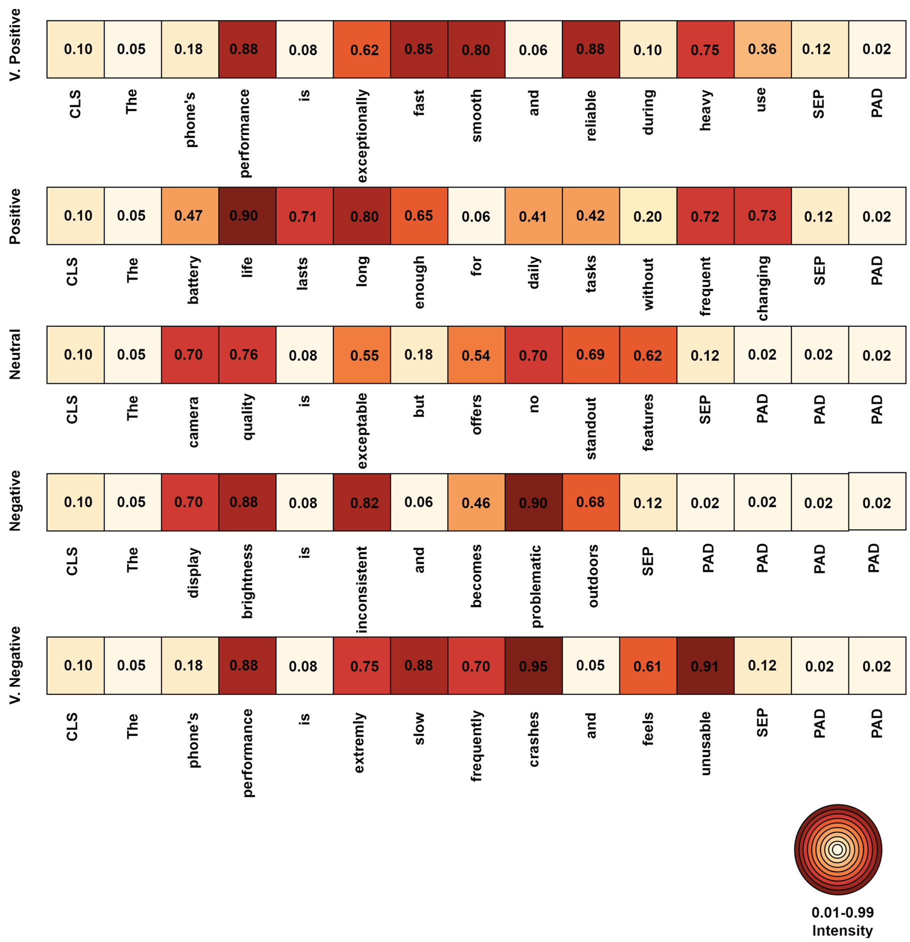

Figure 10: Heatmap of representative sentences from very positive, positive, neutral, negative, and very negative sentiment classes. Color intensity indicates the relative contribution of each token to the model's decision.

5.9 Aspect-Weighted Sentiment Score

To measure the aspect polarity on a quantitative basis, an attention-weighted sentiment aggregation mechanism is adopted. Each token i is assigned

$$\text{AspectScore}(a) = \sum_{i=1}^{n} \alpha_i \cdot s_i \tag{25}$$

where

- α_i = normalized attention weight of token i
- S_i = sentiment contribution of token i
- n = numbers of token related to aspect

This equation will make sure that the tokens that are given more attention will make relatively more contributions to the final aspect sentiment measure.

5.10 Polarity Mapping Strategy

The calculated aspect score is transformed into discrete sentiment categories by applying predefined ranges of threshold values. This conversion guarantees a consistent method of converting continuous sentiment values into categorical values of polarity (e.g., very negative to very positive). The threshold design ensures that polarity is equally assigned to various sentence structures as illustrated in Table 10.

Table 10: Aspect-level sentiment score to polarity mapping.

Aspect Score Range (s)	Polarity Class
$s \geq +0.75$	Very positive
$+0.40 \leq s < +0.75$	Positive
$-0.25 < s < +0.40$	Neutral
$-0.60 \leq s \leq -0.25$	Negative
$s < -0.60$	Very negative

5.11 Worked Example

As an example of the process, the attention weights at the token level are multiplied with their respective sentiment scores as shown in Table 11. The contribution of these contributions is the weighted sum to get the aspect sentiment score. The last polarity is decided by using the threshold-based strategy of mapping. The correspondence of the distribution of attention, token sentiment contributions, and the expected polarity proves the interpretability and reliability of a suggested AAC-HABSA framework.

$$\begin{aligned} \text{AspectScore} = {} & (0.78 \times -0.85) + (0.90 \times -1.00) \\ & + (0.88 \times -0.95) + (0.15 \times -0.40) \\ & + (0.22 \times -0.30) + (0.18 \times -0.20) \end{aligned} \tag{26}$$

$$\text{AspectScore} \approx -0.94 \tag{27}$$

Mapped Polarity: Very Negative

Table 11: Token-level attention weights and sentiment scores for a very negative sentence.

Token	Attention Weight	Sentiment Score
Performance	0.12	0.00
Extremely	0.15	−0.40
Slow	0.78	−0.85
Frequently	0.22	−0.30
Crashes	0.90	−1.00
Feels	0.18	−0.20
Unusable	0.88	−0.95

5.12 Quantitative Interpretability Metrics

Here, we use two quantitative metrics to complement qualitative attention maps and measure attention reliability and faithfulness. As can be observed in Table 12.

Table 12: Quantitative interpretability evaluation of AAC-HABSA.

Dataset	Attention Entropy↓	Faithfulness Score↑
GD	1.72	0.41
Amazon Cell Phone	1.65	0.45
Amazon Product	1.68	0.43

5.12.1 Attention Entropy (Focus Measurement)

We use entropy to quantify the sharpness of the attention distribution:

$$H(\alpha) = -\sum_{i=1}^{n} \alpha_i \log(\alpha_i) \tag{28}$$

A lower value means a more focused attention.j

5.12.2 Attention Faithfulness (Ablation-Based Test)

We measure faithfulness by ablating most attentive tokens and examining changes in predictions:

$$\Delta = |f(x) - f(x \backslash \text{Top-}k(\alpha))| \tag{29}$$

where athe larger Δ indicates higher reliance on important tokens.

5.13 Interpretability Summary

The Qualitative heatmaps are complemented with quantitative interpretability measures. The entropy of the model's attention indicates that the model focuses on a small number of sentiment-relevant tokens, and the faithfulness indicates that these tokens play an important role in the model's predictions. The agreement between attention visualisation and quantitative evaluations shows that the proposed AAC-HABSA model offers faithful and interpretable aspect-level sentiment reasoning.

6 Theoretical and Practical Implications

6.1 Theoretical Implications

This study contributes to the knowledge of aspect-based sentiment analysis (ABSA) more on the insights of architectural design and empirical validation, but not to theoretical development. The suggested ACC-HABSA model indicates the success of incorporating transformer-based contextual encoding into a unified framework of the BiLSTM-based sequential model and adaptive aspect-conditioned attention. The hybrid structure allows modeling the global semantic context and fine-grained aspect-wise sentiment dependencies better. Moreover, the addition of adaptive aspect representations to the attention system is empirically demonstrated to increase aspect sentiment alignment. It is also demonstrated by the use of focal loss that imbalance-aware optimization is important in multi-class sentiment. Taken together, these results offer useful design concepts and tested modeling techniques that can be used to create more effective and understandable ABSA systems.

6.2 Practical Implications

The proposed AAC-HABSA model is a sound solution in application-wise when applied to opinion mining in a real-life scenario, as well as sentiment analytics on a large scale. The framework is accurate at

90.10% on the Amazon Cell Phone data set and 89.60% on Amazon Product reviews and performs similarly and reliably across domains. Fine-grained sentiment intensity allows for analyzing customer feedback in a more detailed way than coarse polarity classification does. This is especially useful in e-commerce websites, customer relationship management, brand awareness systems, and smart recommendations. Furthermore, the release of the GD dataset gives a multi-class benchmark that is balanced and can be used in the future to evaluate experiments and comparative studies in ABSA.

7 Conclusion and Future Work

7.1 Conclusion

The paper introduces a conceptually and architecturally consistent model of Aspect-Based Sentiment Analysis based on the models of HABSA and AAC-HABSA. In contrast to traditional hybrid architectures, which vaguely interconnect contextual encoders and sequential networks, the proposed framework provides a mathematically unified processing pipeline where aspect information is injected before deep contextual encoding through the Adaptive Aspect Conditioning Layer (AACL). The model conditions the token representations and then contextualizes them with a transformer to guarantee that aspect relevance is directly integrated in semantic feature learning and not done post hoc via generic attention weighting. The proposed approach should be viewed as an incremental yet meaningful extension to hybrid ABSA architectures rather than a fundamentally new modeling paradigm.

The architectural flow-Tokenization, Token and Positional Embedding, AACL, Transformer Encoder, BiLSTM, Aspect-Guided Attention, Fully Connected projection, and Focal Loss optimization eliminates the redundancy of the embedding, continuity of the representations, and ensures interpretability through the analysis of structured attention. BiLSTM integration allows sequential narrowing of contextual embedding, whereas the aspect-guided attention ensures that the decision made on polarity is pegged to aspect-specific semantic units. Focal loss is also adopted to achieve greater robustness by diversifying the imbalance in classes and focusing on hard samples.

Empirical comparison of a recently built balanced dataset of some 10,000 annotated reviews, in combination with traditional benchmarks, proves that HABSA and AAC-HABSA are consistent in their accuracy, Macro-F1, robustness, and interpretation analysis compared to transformer-only and conventional hybrid baselines. Ablation experiments demonstrate the utility of the AACL and that the conditioning of aspects at an early phase dramatically enhances semantic-polarity alignment.

In general, this contribution to ABSA is that it creates a structurally integrated, computationally efficient, and interpretable architecture that mediates between learning contextual representation and explicit aspect modeling. The suggested framework offers a theoretically conceptualized, yet scalable, solution to fine-grained sentiment knowledge of real-world opinion mining systems.

The proposed AAC-HABSA framework showed great promise; however, this research has its own limitations. The model has been mainly tested on text-based datasets, which may need to be validated on real-time industrial environments based on multilingual and domain adaptation datasets.

7.2 Future Work

Although the suggested ACC-HABSA framework shows good results in text-based ABSA, there are multiple promising extensions to consider. One of the directions is the adaptation of the Adaptive Aspect Conditioning Layer (AACL) to multimodal conditions, where aspect representations can be co-conditioned with respect to textual, visual or acoustic features. As an example, AACL can be modified to incorporate aspect-aware cues of multiple modalities prior to deep encoding to allow a more holistic sentiment comprehension in practice in real-world contexts like video reviews or social media posts.

In cross-linguistic and low-resource settings, AACL provides a rich interface to the external knowledge, e.g., multilingual embeddings or aspect representations translated into other languages, to enhance aspect sentiment alignment with scarce labeled data. Moreover, AACL variants that are lightweight can be developed to increase performance in resource-limited applications. These guidelines demonstrate the flexibility of the suggested architecture and open-ended opportunities to expand ACC-HABSA to more generalized and scalable sentiment analysis systems.

Acknowledgement: None.

Funding Statement: The authors extend their appreciation to the Deanship of scientific research, Islamic University of Madinah, Saudi Arabia, for funding this research work.

Author Contributions: Conceptualization, Mahander Kumar; methodology, Mahander Kumar and Lal Khan; software, Mahander Kumar; validation, Lal Khan; formal analysis, Lal Khan, Mohammad Zubair Khan and Ibrahim Aljubayri; investigation, Mahander Kumar; resources, Lal Khan, Ibrahim Aljubayri and Mohammad Zubair Khan; data curation, Mahander Kumar; writing—original draft preparation, Mahander Kumar and Lal Khan; writing—review and editing, Lal Khan, Mohammad Zubair Khan; supervision, Lal Khan, Mohammad Zubair Khan and Ibrahim Aljubayri; project administration, Lal Khan and Mohammad Zubair Khan; funding acquisition, Mohammad Zubair Khan and Ibrahim Aljubayri. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The data sets that were used and analyzed in the present study are available to the public and can be obtained from the Kaggle platform and can be accessed through the following links: Dataset 1 (Amazon Cell Phones Reviews): https://www.kaggle.com/datasets/grikomsn/amazon-cell-phones-reviews?select=20191226-reviews.csv; Dataset 2 (Sentiment Analysis: Amazon Product Reviews): https://www.kaggle.com/datasets/miriamodeyianypeter/sentiment-analysis-amazon-product-reviews; Dataset 3: (GD dataset): https://www.kaggle.com/datasets/soul81/gd-dataset/data.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Hussain A, Cambria E. Semi-supervised learning for big social data analysis. Neurocomputing. 2018;275(5):1662–73. doi:10.1016/j.neucom.2017.10.010.
2. Aziz K, Ji D, Chakrabarti P, Chakrabarti T, Iqbal MS, Abbasi R. Unifying aspect-based sentiment analysis BERT and multi-layered graph convolutional networks for comprehensive sentiment dissection. Sci Rep. 2024;14(1):14646. doi:10.1038/s41598-024-61886-7.
3. Levy O, Goldberg Y. Neural word embedding as implicit matrix factorization. Adv Neural Inf Process Syst. 2014;27:2177–85.
4. Lai S, Liu K, He S, Zhao J. How to generate a good word embedding. IEEE Intell Syst. 2016;31(6):5–14. doi:10.1109/mis.2016.45.
5. Li M, Lu Q, Long Y, Gui L. Inferring affective meanings of words from word embedding. IEEE Trans Affective Comput. 2017;8(4):443–56. doi:10.1109/taffc.2017.2723012.
6. Amjad A, Khan L, Ashraf N, Mahmood MB, Chang HT. Recognizing semi-natural and spontaneous speech emotions using deep neural networks. IEEE Access. 2022;10(6):37149–63. doi:10.1109/access.2022.3163712.
7. Khan L. Urdu sentiment analysis using deep learning methods [dissertation]. Taoyuan, Taiwan: Chang Gung University; 2022.

8. Kumar M, Khan L, Chang HT. Evolving techniques in sentiment analysis: a comprehensive review. PeerJ Comput Sci. 2025;11(3):e2592. doi:10.7717/peerj-cs.2592.
9. Kaelbling LP, Littman ML, Moore AW. Reinforcement learning: a survey. J Artif Intell Res. 1996;4:237–85. doi:10.1613/jair.301.
10. Kumar M, Khan L, Choi A. RAMHA: a hybrid social text-based transformer with adapter for mental health emotion classification. Mathematics. 2025;13(18):2918. doi:10.3390/math13182918.
11. Wang Y, He J, Wang D, Wang Q, Wan B, Luo X. Multimodal transformer with adaptive modality weighting for multimodal sentiment analysis. Neurocomputing. 2024;572(6):127181. doi:10.1016/j.neucom.2023.127181.
12. Sun H, Niu Z, Wang H, Yu X, Liu J, Chen YW, et al. Multimodal sentiment analysis with mutual information-based disentangled representation learning. IEEE Trans Affective Comput. 2025;16(3):1606–17. doi:10.1109/taffc.2025.3529732.
13. Li Y, Lan X, Chen H, Lu K, Jiang D. Multimodal PEAR chain-of-thought reasoning for multimodal sentiment analysis. ACM Trans Multimedia Comput Commun Appl. 2024;20(9):1–23. doi:10.1145/3672398.
14. Liu Z, Zhou B, Chu D, Sun Y, Meng L. Modality translation-based multimodal sentiment analysis under uncertain missing modalities. Inf Fusion. 2024;101:101973. doi:10.1016/j.inffus.2023.101973.
15. Cho K, van Merrienboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, et al. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP); 2014 Oct 25–29; Doha, Qatar. p. 1724–34. doi:10.3115/v1/d14-1179.
16. Wang Y, Huang M, zhu X, Zhao L. Attention-based LSTM for aspect-level sentiment classification. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing; 2016 Nov 1–5; Austin, TX, USA. p. 606–15. doi:10.18653/v1/d16-1058.
17. Ma D, Li S, Zhang X, Wang H. Interactive attention networks for aspect-level sentiment classification. arXiv:1709.00893. 2017.
18. Devlin J, Chang MW, Lee K, Bert TK. BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers); 2019 Jun 2–7; Minneapolis, MN, USA. p. 4171–86.
19. Radford A, Narasimhan K, Salimans T, Sutskever I. Improving Language Understanding by Generative Pre-Training. 2018 [cited 2026 Jan 1]. Available from: https://cdn.openai.com/research-covers/languageunsupervised/language_understanding_paper.pdf.
20. Raffel C, Shazeer N, Roberts A, Lee K, Narang S, Matena M, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. J Mach Learn Res. 2020;21(140):1–67. doi: 10.48550/arxiv.1910.10683.
21. Sun C, Huang L, Qiu X. Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence. arXiv:1903.09588. 2019.
22. Rogers A, Kovaleva O, Rumshisky A. A primer in BERTology: what we know about how BERT works. Trans Assoc Comput Linguist. 2020;8(5):842–66. doi:10.1162/tacl_a_00349.
23. Chatterjee A, Gupta U, Chinnakotla MK, Srikanth R, Galley M, Agrawal P. Understanding emotions in text using deep learning and big data. Comput Hum Behav. 2019;93(3):309–17. doi:10.1016/j.chb.2018.12.029.
24. Rezaeinia SM, Rahmani R, Ghodsi A, Veisi H. Sentiment analysis based on improved pre-trained word embeddings. Expert Syst Appl. 2019;117(19):139–47. doi:10.1016/j.eswa.2018.08.044.
25. Liu G, Guo J. Bidirectional LSTM with attention mechanism and convolutional layer for text classification. Neurocomputing. 2019;337(5):325–38. doi:10.1016/j.neucom.2019.01.078.
26. Kumar M, Khan L, Khan MZ, Ali Alhussan A. AGF-HAM: adaptive gated fusion hierarchical attention model for explainable sentiment analysis. Mathematics. 2025;13(24):3892. doi:10.3390/math13243892.
27. Pang B, Lee L, Vaithyanathan S. Thumbs up? Sentiment classification using machine learning techniques. In: Proceedings of the ACL Conference on Empirical Methods in Natural Language Processing (EMNLP); 2002 Jul 6–7; Philadelphia, PA, USA. p. 79–86.

28. Wang Y, Chen Q, Shen J, Hou B, Ahmed M, Li Z. Aspect-level sentiment analysis based on gradual machine learning. Knowl Based Syst. 2021;212(3):106509. doi:10.1016/j.knosys.2020.106509.
29. Wang J, Xu B, Zu Y. Deep learning for aspect-based sentiment analysis. In: Proceedings of the 2021 International Conference on Machine Learning and Intelligent Systems Engineering (MLISE); 2021 Jul 9–11; Chongqing, China. p. 267–71. doi:10.1109/mlise54096.2021.00056.
30. Xu H, Shu L, Yu PS, Liu B. Understanding pre-trained BERT for aspect-based sentiment analysis. arXiv:2011.00169. 2020.
31. Luo W, Zhang W, Zhao Y. A survey of transformer and GNN for aspect-based sentiment analysis. In: Proceedings of the 2021 International Conference on Computer Information Science and Artificial Intelligence (CISAI); 2021 Sep 17–19; Kunming, China. p. 353–7. doi:10.1109/cisai54367.2021.00074.
32. Jagtap V, Pawar K. Analysis of different approaches to sentence-level sentiment classification. Int J Sci Eng Technol. 2013;2(3):164–70.
33. Noble WS. What is a support vector machine? Nat Biotechnol. 2006;24(12):1565–7. doi:10.1038/nbt1206-1565.
34. Webb GI. Naïve Bayes. In: Encyclopedia of machine learning. Boston, MA, USA: Springer; 2011. p. 713–4. doi:10.1007/978-0-387-30164-8_576.
35. Varghese R, Jayasree M. Aspect based sentiment analysis using support vector machine classifier. In: Proceedings of the 2013 International Conference on Advances in Computing, Communications and Informatics (ICACCI); 2013 Aug 22–25; Mysore, India. p. 1581–6. doi:10.1109/icacci.2013.6637416.
36. Hegde R, Seema S. Aspect based feature extraction and sentiment classification of review data sets using incremental machine learning algorithm. In: Proceedings of the 2017 Third International Conference on Advances in Electrical, Electronics, Information, Communication and Bio-Informatics (AEEICB); 2017 Feb 27–28; Chennai, India. p. 122–5. doi:10.1109/aeeicb.2017.7972395.
37. Manek AS, Shenoy PD, Mohan MC, R VK. Aspect term extraction for sentiment analysis in large movie reviews using Gini Index feature selection method and SVM classifier. World Wide Web. 2017;20(2):135–54. doi:10.1007/s11280-015-0381-x.
38. Qi Y, Shabrina Z. Sentiment analysis using Twitter data: a comparative application of lexicon- and machine-learning-based approach. Soc Netw Anal Min. 2023;13(1):31. doi:10.1007/s13278-023-01030-x.
39. Naeem S, Logofatu D, Muharemi F. Sentiment analysis by using supervised machine learning and deep learning approaches. In: Advances in computational collective intelligence. Cham, Switzerland: Springer; 2020. p. 481–91. doi:10.1007/978-3-030-63119-2_39.
40. Pathik N, Shukla P. Aspect based sentiment analysis of unlabeled reviews using linguistic rule based LDA. J Cases Inf Technol. 2022;24(3):1–19. doi:10.4018/jcit.20220701.oa3.
41. Shi J, Li W, Bai Q, Ito T. BeeAE: effective aspect term extraction with artificial bee colony. J Supercomput. 2022;78(16):17969–91. doi:10.1007/s11227-022-04579-0.
42. Ray P, Chakrabarti A. A mixed approach of deep learning method and rule-based method to improve aspect level sentiment analysis. ACI. 2022;1(2):163–78. doi:10.1016/j.aci.2019.02.002.
43. Xu G, Zhang Z, Zhang T, Yu S, Meng Y, Chen S. Aspect-level sentiment classification based on attention-BiLSTM model and transfer learning. Knowl Based Syst. 2022;245(2):108586. doi:10.1016/j.knosys.2022.108586.
44. Khan L, Amjad A, Ashraf N, Chang HT, Gelbukh A. Urdu sentiment analysis with deep learning methods. IEEE Access. 2021;9:97803–12. doi:10.1109/access.2021.3093078.
45. Tenney I, Das D, Pavlick E. BERT rediscovers the classical NLP pipeline. arXiv:1905.05950. 2019.
46. Rani S, Jain A. Aspect-based sentiment analysis of drug reviews using multi-task learning based dual BiLSTM model. Multimed Tools Appl. 2024;83(8):22473–501. doi:10.1007/s11042-023-16360-3.
47. Koroteev MV. BERT: a review of applications in natural language processing and understanding. arXiv:2103.11943. 2021.
48. Khan L, Qazi A, Chang HT, Alhajlah M, Mahmood A. Empowering Urdu sentiment analysis: an attention-based stacked CNN-Bi-LSTM DNN with multilingual BERT. Complex Intell Syst. 2024;11(1):10. doi:10.1007/s40747-024-01631-9.

49. Alhadlaq A, Altheneyan A. Distilroberta2gnn: a new hybrid deep learning approach for aspect-based sentiment analysis. PeerJ Comput Sci. 2024;10(6):e2267. doi:10.7717/peerj-cs.2267.
50. Meškelė D, Frasincar F. ALDONAr: a hybrid solution for sentence-level aspect-based sentiment analysis using a lexicalized domain ontology and a regularized neural attention model. Inf Process Manag. 2020;57(3):102211. doi:10.1016/j.ipm.2020.102211.
51. Huang X, Li J, Wu J, Chang J, Liu D, Zhu K. Flexibly utilizing syntactic knowledge in aspect-based sentiment analysis. Inf Process Manag. 2024;61(3):103630. doi:10.1016/j.ipm.2023.103630.
52. Xiao L, Wu X, Yang S, Xu J, Zhou J, He L. Cross-modal fine-grained alignment and fusion network for multimodal aspect-based sentiment analysis. Inf Process Manag. 2023;60(6):103508. doi:10.1016/j.ipm.2023.103508.
53. Chen P, Sun Z, Bing L, Yang W. Recurrent attention network on memory for aspect sentiment analysis. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing; 2017 Sep 7–11; Copenhagen, Denmark. p. 452–61. doi:10.18653/v1/d17-1047.
54. Tang D, Qin B, Liu T. Aspect level sentiment classification with deep memory network. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing; 2016 Nov 1–5; Austin, TX, USA. p. 214–24. doi:10.18653/v1/d16-1021.
55. Xue W, Li T. Aspect based sentiment analysis with gated convolutional networks. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); 2018 Jul 15–20; Melbourne, Australia. p. 2514–23. doi:10.18653/v1/p18-1234.
56. Xu H, Liu B, Shu L, Yu PS. BERT post-training for review reading comprehension and aspect-based sentiment analysis. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers); 2019 Jun 2–7; Minneapolis, MN, USA. p. 2324–35.
57. Huang B, Ou Y, Carley KM. Aspect level sentiment classification with attention-over-attention neural networks. In: Social, cultural, and behavioral modeling. Cham, Switzerland: Springer; 2018. p. 197–206.
58. Zhao M, Yang J, Shang F. Dependency-enhanced graph convolutional networks for aspect-based sentiment analysis. Neural Comput Appl. 2023;35(19):14195–211. doi:10.1007/s00521-023-08384-5.
59. Gu T, Zhao H, Li M. Effective inter-aspect words modeling for aspect-based sentiment analysis. Appl Intell. 2023;53(4):4366–79. doi:10.1007/s10489-022-03630-0.
60. Su Z, Gao B, Pan X, Liu Z, Ji Y, Liu S. CGT: a clause graph transformer structure for aspect-based sentiment analysis. Data Knowl Eng. 2024;153:102332. doi:10.1016/j.datak.2024.102332.
61. Yin S, Zhong G. TextGT: a double-view graph transformer on text for aspect-based sentiment analysis. Proc AAAI Conf Artif Intell. 2024;38:19404–12.
62. Nuha U, Lin CH. Aspect-based sentiment analysis with semi-supervised approach on Taiwan social distancing app user reviews. In: Proceedings of the 2023 International Conference on Artificial Intelligence in Information and Communication (ICAIIC); 2023 Feb 20–23; Bali, Indonesia. p. 444–7. doi:10.1109/icaiic57133.2023.10067048.
63. Dai A, Hu X, Nie J, Chen J. Learning from word semantics to sentence syntax by graph convolutional networks for aspect-based sentiment analysis. Int J Data Sci Anal. 2022;14(1):17–26. doi:10.1007/s41060-022-00315-2.