Aspect Based Sentiment Analysis using Pre-trained models

MSc Research Project
MSc Artifical Intelligence

Madhuri Chowdappa Madhugiri
Student ID:X22200622

School of Computing
National College of Ireland

Supervisor: Kislay Raj

National College of Ireland
Project Submission Sheet
School of Computing

Student Name:	Madhuri Chowdappa Madhugiri
Student ID:	X22200622
Programme:	MSc Artifical Intelligence
Year:	2024
Module:	MSc Research Project
Supervisor:	Kislay Raj
Submission Due Date:	12/08/2024
Project Title:	Aspect Based Sentiment Analysis using Pre-trained models
Word Count:	5933
Page Count:	16

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	
Date:	12th August 2024

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	☐
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	☐
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	☐

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Aspect Based Sentiment Analysis using Pre-trained models

Madhuri Chowdappa Madhugiri
X22200622

Abstract

Aspect Based Sentiment Analysis is the important task in natural language processing focusing on identifying the sentiments expressed towards specific aspects in the text. This project introduces an advanced ABSA methodology by integrating the pre trained BERT model with Hierarchical Attention Network (HAN) and the Knowledge Graph . The proposed approach enhances the model's ability to comprehend and analyze context-specific nuances in customer reviews by using the external knowledge and attention mechanisms at both word and sentence levels. The methodology involves the systematic process of data pre-processing, aspect extraction, and sentiment classification and followed by model interpretation using LIME. Experimental results demonstrate that the combined HAN and BERT model significantly improves the sentiment classification accuracy for positive sentiments by offering the detailed insights into the sentiment analysis process. Although the model performs exceptionally well overall there is a slight limitation in detecting negative sentiments which can be addressed in future work. This project contributes to the development of context-aware sentiment analysis models with applications in customer feedback analysis, product reviews and other domains requiring detailed sentiment insights.

Keywords: Aspect-Based Sentiment Analysis, Hierarchical Attention Network, Knowledge Graph, BERT, Natural Language Processing, Sentiment Classification, LIME, Data Pre-processing, Model

1 Introduction

Sentiment analysis focuses on identifying and categorizing the opinions in the text. Sentiment analysis in field of Artificial intelligence is the combination of natural language processing, machine learning and psychology which determines the sentiment in the text and categorize the text into positive, negative or neutral. Since 2000 there is huge amount of textual data created by social media, online reviews and blogs and because of the increase in data day by day, sentiment analysis has gained popularity in recent years. Analyzing this data manually is infeasible so making automated sentiment analysis tools is important. Sentiment analysis is used in various domains such as business intelligence, social media, market research and customer feedback analysisBhattacharyya (2013).

Traditional sentiment analysis methods assign a single sentiment label to an entire piece of text and provides the high-level overview of the sentiment but it fails to capture the nuances and specific opinions about different aspects.

Aspect-Based Sentiment Analysis (ABSA) is the advanced technique in the field of sentiment analysis which addresses the limitations of traditional sentiment analysis methods. ABSA identify the specific aspects mentioned in the text and determines the sentiment associated with each aspect instead of just assigning an overall sentiment to a piece of text. This fine-grained analysis provides the deeper understanding of the text which gives the detailed insights of various components. For example consider a product review that says "The battery life of thr phone is good but the build quality is poor." Traditional sentiment analysis classify this review as neutral or mixed and it fails to provide the actionable insights. In ABSA "battery life" and "build quality" are identified as the distinct aspects and determine that the sentiment towards "battery life" is positive and the sentiment towards "camera quality" is negative. This details gives the insights for businesses to improve specific features of the products or services based on feedback. Traditional aspect based sentiment analysis methods provides the effective model for determining overall sentiment but miss the subtleties associated with individual aspects within a review. This limitation has driven research towards more advanced models that can better understand context and aspect-specific sentiments.

A promising approach that involves the integration of the Hierarchical Attention Network with the Knowledge Graph and the pre-trained BERT model. The combination of these technologies allows the more nuanced understanding of text by using both the external knowledge and deep semantic comprehension. The Knowledge Graph provides structured, domain-specific information that increases the model's ability to capture the relationships between different aspects and the BERT model which is fine-tuned for ABSA tasks which offers deep understanding of language nuances.

The importance of this research lies in addressing the limitations of the previous models that struggle with the context-specific nuances and complex aspect relationships. By using the Knowledge Graph the model can access the external and structured knowledge which is useful in domains where understanding the relationships between various aspects is important. This study also aims to explore how the synergy between HAN, Knowledge Graphs and BERT model can lead to the improvements in sentiment classification accuracy in challenging scenarios where context is the key.

2 Related Work

2.1 Traditional method for Sentiment analysis

Sentiment analysis has traditionally been approached using various machine learning techniques such as Naive Bayes classifiers, Support Vector Machines (SVMs), and Decision Trees. These methods are foundational in the development of sentiment analysis. Naive Bayes models assume that the presence of a particular feature in a class is independent of other features. This method has shown the effectiveness in text classification tasks due to its simplicity and the fast training times. But the strong independence assumption does not hold true in real-world text data which leads to the suboptimal performance in complex situations. For example in the movie review domain the sentiment classification can be challenging because sentiment is expressed subtly and contextually Pang et al. (2002). SVMs are effective in the high-dimensional spaces and they are used in sentiment analysis due to their robustness in handling large feature sets and effectiveness in binary classification tasks Joachims (1998). But SVMs require high computational costs with large datasets and it requires the tuning of parameters such as the regularization term and kernel functions to achieve optimal performance. For sentiment analysis decision tree creates the tree structure in which each leaf node represents a sentiment class (e.g., positive, negative, neutral). Random Forests is used enhance classification accuracy by combining multiple decision trees to reduce overfitting and improve generalization. A comprehensive survey by Safavian and Landgrebe highlights the applications and advantages of DTCs Safavian and Landgrebe (1991). While DTCs offer interpretability and are capable of handling large datasets but they can be prone to overfitting with noisy data. Random Forests mitigate this issue by averaging the results of multiple trees by leading to better generalization and robustness.social media sentiment analysis on platforms like YouTube where user comments can be highly variable and noisy the Decision Tree and Random Forest algorithms have proven effective with this dataset. The Aufar et al. (n.d.) study analyzed public comments on Nokia's products from YouTube. They classified these comments into positive, negative, and neutral sentiments using both Decision Tree and Random Forest algorithms. The results showed that the Decision Tree algorithm performs better with the accuracy of (89.4%). But traditional methods requires the extensive feature engineering and cannot capture the context of words. They struggle with polysemy which is multiple meanings of a word and synonymy which is different words with similar meanings leads to the lower performance in complex sentiment analysis tasks. For example traditional methods might treat "bank" as a single feature without understanding whether it refers to a financial institution or the side of a riverKatić and Milićević (2018).

2.2 Sentiment Analysis Using Pre-trained Models

Traditional machine learning models treat words as atomic units without inherent similarity which is robust but limited in tasks like automatic speech recognition and machine translation where data is constrained. Recent machine learning advancements enable training of complex models on larger datasets which outperforms the simpler models. The distributed representations like Word2Vec and GloVe have significantly improved over traditional models. Word2Vec, introduced by Mikolov et al. (2013) includes Continuous Bag-of-Words (CBOW) and Continuous Skip-gram architectures by capturing complex word relationships and analogies. GloVe by Pennington et al. (n.d.) constructs the word vectors based on global word-word co-occurrence matrices by effectively capturing global statistical information. But Both models lack context sensitivity by providing a single representation for each word regardless of context which can lead to misunderstandings in sentiment analysis. For instance, "apple" would have the same embedding in "apple pie" and "Apple Inc.," which confuses the sentiment models that rely on context to determine sentiment polarity.

Transformers are considered as the backbone of modern NLP models with architectures like BERT and GPT leading the way. BERT's bidirectional approach allows it to consider context from both directions by making it handle the polysemous words by providing different embeddings for "bank" in "river bank" and "financial bank." GPT also excels in generating coherent, context-aware text, proving invaluable in dialogue systems. Based on these architecture RealFormer introduces a Residual Attention Layer technique that significantly enhances the performance of Transformers. This technique not only stabilizes training but also results in sparser attention by improving model robustness and generalization. RealFormer achieves these improvements by adding residual attention scores from previous layers by creating a direct path for propagating attention through the network without additional parameters or complex modifications. This advancement shows the potential of simple yet effective modifications to existing Transformer architectures in pushing the boundaries of NLP performance He et al. (2023).

Models like Word2Vec and GloVe treat words as atomic units without inherent similarity which leads to the context insensitivity and limited handling of polysemy and homonymy. These limitations restrict their effectiveness in the aspect-based sentiment analysis where understanding of the specific context of each aspect is important. Advanced models like BERT and GPT overcome these issues by considering bidirectional context and generating coherent, context-aware text but they come with high computational costs and complexity which is highlighted by Devlin et al. (2019) and Radford et al. (2018). Despite of the advancements, interpreting decisions made by these models remains challenging and they still struggle with fine-grained aspect-level sentiment detection. RealFormer introduced He et al. (2023) enhances transformers through a Residual Attention Layer by improving stability and generalization. But it lead to sparser attention, potentially missing intricate dependencies, and poses challenges in integration and scalability.

2.3 Aspect-Based Sentiment Analysis

Aspect-Based Sentiment Analysis (ABSA) is the specialized form of the sentiment analysis which focuses on identifying the sentiments expressed about specific aspects or features within the text. Traditional methods for ABSA include techniques such as Latent Dirichlet Allocation (LDA), rule-based systems and machine learning approaches like Conditional Random Fields (CRFs) and Support Vector Machines (SVMs).

LDA is the generative probabilistic model that is used for topic modeling in ABSA. It helps to identify the topics within a set of the documents. This method is useful in identifying aspects in the e-commerce user reviews, as demonstrated by Wahyudi and Kusumaningrum (n.d.). By applying LDA we can uncover hidden topics or aspects in large corpora of text data by making it possible to associate sentiments with these topics effectively. But one limitation of LDA is its difficulty in capturing the nuanced and context-specific nature of certain aspects which can lead to less accurate sentiment analysis in complex sentences.

Rule based systems uses the predefined linguistic rules and lexicons to extract aspects and determine the corresponding sentiments. These systems operate by identifying patterns in text, such as syntactic dependencies or frequent co-occurrences of specific words to infer relationships between aspects and sentiments. Rule based systems are highly interpretable and effective for initial analyses or datasets with relatively simple and consistent language but they face significant challenges when dealing with the variability inherent in natural language. For example, these systems often struggle with complex sentences, multi-word aspects, and implicit expressions and also require the extensive manual rule adjustments for different contexts and domains. The scalability and adaptability of rule-based systems are limited when applied to diverse or evolving datasetsRuskanda et al. (2018).

Traditional methods like Latent Dirichlet Allocation (LDA) and rule-based systems are useful but they have limitations when it comes to handling the complexity and nuance of the natural language. Pretrained models address the issue of the traditional method for ABSA tasks. Recent advancement in Aspect-Based Sentiment Analysis (ABSA) with the integration of more advanced models such as Hierarchical Attention Networks (HAN) and BERT address these limitations. For example HAN models are designed to capture the hierarchical structures in text which applies the attention mechanisms at the word and sentence levels which allow the model to focus on important words and sentences by improving the extraction of relevant information and sentiment analysis. This approach is particularly effective in contexts where the textual comments is involved as demonstrated in the researchZhou et al. (2024) . By using the hierarchical structures in text (HAN) and the contextual embeddings provided by BERT, ABSA methods can offer more precise and contextually aware sentiment analysis. These advanced models not only improve the accuracy of sentiment analysis but also reduce the need for extensive manual feature engineering, making them more adaptable to different domainsNg et al. (2024). In summary while traditional methods like LDA and rule-based systems have laid the groundwork for ABSA by enabling the discovery of aspects and initial sentiment classifications but the advent of neural network-based approaches has significantly enhanced the granularity and accuracy of sentiment analysis. The hybrid approach combining the HAN and BERT model represents the significant leap forward by integrating the strengths of both traditional and modern techniques.

3 Methodology

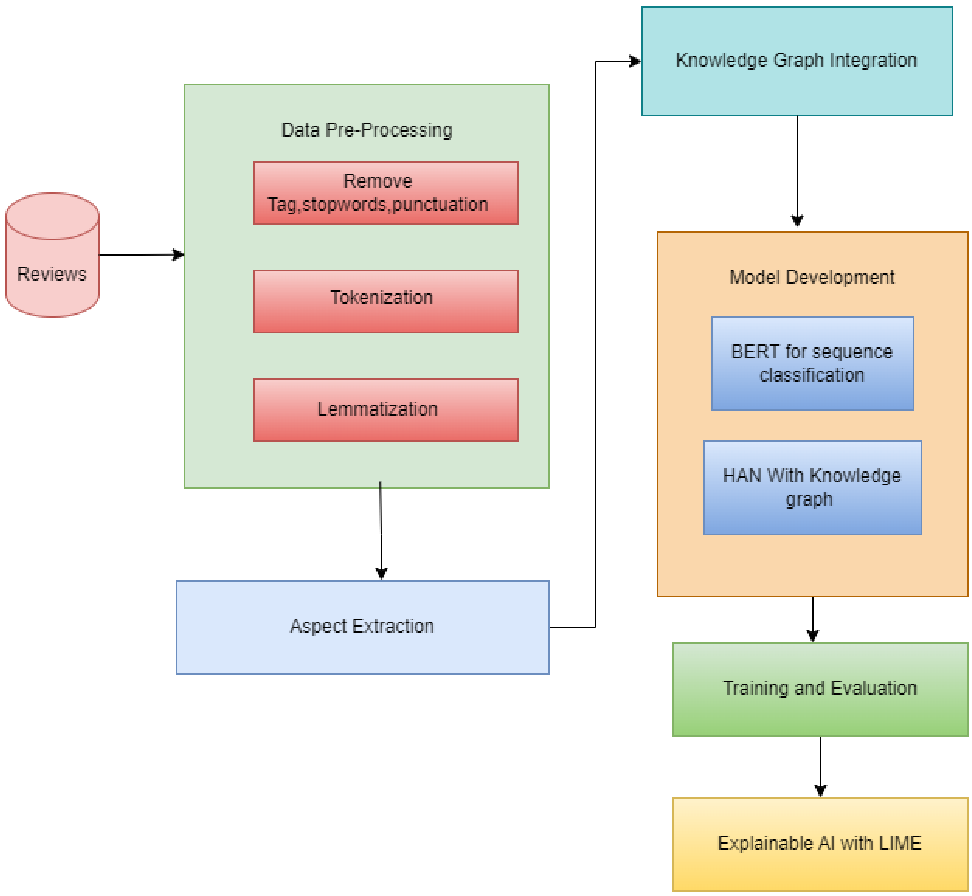

Figure 1: Overview of the proposed methodology for Aspect-Based Sentiment Analysis (ABSA) using Hierarchical Attention Network (HAN), Knowledge Graph, and pre-trained BERT model.

3.1 Data Collection

In data collection reviews from three sources are combined together: Yelp, Trip Advisor and Amazon. Using the Yelp API we gathered reviews on various businesses such as restaurants, shops, and services and gathered the structured data including review text, ratings, business information, and reviewer details. Similarly accessed the Trip Advisor's API to collect reviews related to travel, hotels, restaurants and attractions to get review content, ratings, and other metadata. Additionally performed web scraping on Amazon to extract product reviews, capturing review text, ratings, product details, and reviewer information directly from the web pages. Data processing is focused on standardizing and cleaning the data. The reviews from Yelp, Trip Advisor, and Amazon are combined into unified dataset ensuring a consistent structure across all sources. This involved removing duplicates, handling missing values by either filling in gaps or removing incomplete entries and normalizing data formats, such as standardizing rating scales and text encoding. Furthermore we derived a sentiment score for each review based on its rating categorizing them into "Positive," "Neutral," and "Negative" sentiments to enrich the dataset with sentiment information.

3.2 Data Pre-Processing

The dataset contains 40,490 reviews collected from multiple sources which contains three primary attributes including Rating, Review, and Sentiment. The Rating attribute which is an integer and ranges from 1 to 5 which has mean value of 3.91 and a standard deviation of 1.33 indicating that the reviews are generally positive but exhibit some variability. The Review attribute contains the textual content of the reviews. Sentiment attribute is derived from the ratings which indicates whether the review is positive (1) or negative (0). Approximately 72% of the reviews are positive as reflected by the mean sentiment score of 0.72 and a standard deviation of 0.45. This shows the imbalance in the sentiment attribute. Resampling is applied to address the issue of attribute imbalance.

3.2.1 Text Preprocessing

Text preprocessing is the important step in preparing the raw textual data for the analysis by converting it into the format that can be processed by the machine learning models effectively. This process involves several key sub-steps.Tokenization where the text is split into tokens which allows the detailed word-level analysis. Normalization which standardizes the text by converting all the characters to lowercase and removing the special characters to maintain the uniformity and consistency. Stop Words Removal where the non-informative words such as "and" "the" and "is" is removed to focus on more meaningful words and Lemmatization which reduces the words to their base or root forms. It minimizes the number of unique words and treating variations of the same word as a single entity.

3.2.2 Aspect Extraction

The second step is to extract key components or entities from the text which are referred as the aspects. This process involves identifying the elements within the reviews that contribute more to the sentiment analysis. We utilize spaCy's advanced NLP capabilities for this purpose which includes Named Entity Recognition (NER). spaCy's NER capabilities are used to identify and classify named entities in the text such as product names, brands, places, and other relevant terms. These entities o represent the key aspects that users mention in their reviews. Noun Chunking is used which involves identifying and extracting noun phrases from the text. . For example in the phrase "great battery life," "battery life" is a noun chunk. These aspects are important as they represent the primary components that users discuss in their reviews. By isolating these key components we can perform targeted sentiment analysis which focuses on specific parts of the text that are relevant to the user's experience and opinions.

3.2.3 Sentiment Labelling

After extracting the aspects the next step is to the assign sentiment labels to each aspect. This involves determining whether the sentiment expressed towards each aspect is positive or negative. We have used pretrained BERT model to assign to assign the sentiments of each aspect. Each aspect is assigned the binary sentiment label whether it is positive or negative. This binary classification simplifies the task and makes it easier for the model to learn and predict sentiments accurately. For example, aspects like "battery life" can be labeled as positive if the user mentions it favorably and negative if the user complains about it.

3.3 Topic modelling with Latent Dirichlet Allocation (LDA)

Topic modeling with LDA is used to uncover the hidden sentiment structure in the text data. LDA helps to identify the topics that are prevalent in the corpus which can provide the additional context for the aspect extraction and sentiment analysis. Each review is analyzed to determine the distribution of topics which are then used to understand the text and its sentiment. By integrating LDA we can understand the context in which aspects are mentioned which leads to more accurate sentiment predictions.

3.4 Knowledge Graph Integration

Integrating the external knowledge into the sentiment analysis process can improve the model's understanding and performance. We have used pre-trained spaCy model such as "en_core_web_lg" to obtain vector representations of named entities in the text. These entity vectors encapsulate rich semantic information about the entities mentioned in the reviews. When multiple entities are present their vectors are averaged to create a single and comprehensive knowledge embedding for each text segment. This embedding enriches the textual data with external knowledge by providing the broader context that improves the model's ability to understand and analyze sentiments.

3.5 Model Development

3.5.1 BERT for Sequence Classification

For aspect-based sentiment analysis we have used the BERT model. The Variant 'BertForSequenceClassification' from Hugging Face's Transformers library is used. This model was initialized with the pre-trained 'bert-base-uncased' version and fine-tuned for the sequence classification tasks. BERT was chosen over the simpler models because of its ability to the generate deep contextual embeddings, capturing the nuances of language and understands the context of the words. The training arguments are configured with the number of epochs, batch size, weight decay, evaluation strategy and the logging configuration. These parameters ensure that the model is optimized for the performance and effectively learn from the training data.

3.5.2 Hierarchical Attention Network (HAN) with Knowledge Graph

To enhance the model's performance Hierarchical Attention Network (HAN) is integrated with the knowledge graph embeddings. The HAN model is designed to process text at both the word and sentence levels. Attention mechanisms is applied to identify and emphasize the most relevant words and sentences. Word-level attention focuses on the importance of individual words within the sentence and the sentence-level attention highlights the key sentences in the document. The output from this hierarchical structure is then concatenated with knowledge graph embeddings which gives the additional context and structured external knowledge. The knowledge graph embeddings are generated using spaCy's advanced NLP capabilities which extract named entities from the text and represent them as vectors. This combination of textual and external knowledge results in the more contextually aware and accurate sentiment analysis model.

3.6 Training and Evaluation

3.6.1 Custom Dataset Class

custom dataset class is defined to manage the reviews, aspects, sentiments and knowledge embeddings efficiently. This class handles the tokenization of the text and retrieval of knowledge embeddings and prepares the data for model training. The class make sure that the each review is correctly formatted and includes all the necessary components required for the smooth and effective training.

3.6.2 Training Loop

The training process involve using the binary cross-entropy loss and the Adam optimizer to minimize the prediction errors and update the model parameters. The training performs the forward and the backward passes and also optimizes the model parameters. This iterative process adjusts the model to predict the sentiments better based on the training data and improves its accuracy and performance.

3.6.3 Evaluation

Evaluation metrics such as accuracy, F1 score, precision and recall is used after training the model. The evaluation process involves applying the model to the validation dataset, computing predictions and calculating the evaluation metrics. This step is important for understanding the model's effectiveness and also to identify the areas for improvement.

3.7 Explainable AI with LIME

3.7.1 LIME Explainer Initialization

To ensure the transparency and interpretability of the model's predictions LIME (Local Interpretable Model-agnostic Explanations) is used. LIME helps to understand the complex models by approximating them locally with simpler, interpretable models. LimeTextExplainer is initialized to provide explanations for the sentiment analysis model's predictions.

3.7.2 Prediction Function for LIME

prediction function which is compatible with LIME is defined which tokenizes the input text, performs inference on the CPU and returns the predicted probabilities. This function ensures that LIME can generate explanations based on the model's predictions making the complex sentiment analysis model more understandable.

3.7.3 Aspect Explanation

Aspect is selected from the dataset for explanation. The aspect keys are concatenated to form a single text input which is then used by the LIME explainer to generate explanations. LIME perturbs the input text and observes changes in model predictions and identifies the parts of the text that contribute most to the prediction. This process provides the insights into the model's decision-making and highlights the influential words that led to the sentiment prediction.

3.7.4 Visualization

Finally the LIME explanations are displayed in the notebook allowing to visualize which parts of the input text are most influential in the model's prediction. This visualization helps in validating the model's behavior and understanding how it arrives at its predictions and also ensuring that the sentiment analysis process is transparent and interpretable.

4 Design Specification

4.1 BERT

The DistilBERT model used for the sentiment analysis which is the faster version of the original BERT model. The process starts with the tokenization where the special tokens are added to mark the beginning and end of the sentences. These tokens are then turned into the vectors in the embedding layer which helps the model to understand their meanings and positions. The core of DistilBERT consists of the six transformer layers. Each layer has two main part which is multi-head self-attention and the feed-forward neural network. The self-attention mechanism allow the model to understand the context and relationships between the words. The feed-forward network process this information to enhance understanding. Each transformer layer includes the normalization steps to stabilize the learning process and the improve performance. The final layer outputs the summary of the entire text which is then passed to the classification layer. This layer makes the final decision about the sentiment of the text by determining if it is positive or negative. The model is trained using the set of predefined parameters that control aspects like the number of training cycles, batch sizes and how often the model should be evaluated. DistilBERT's architecture allows it to quickly and efficiently analyze text for sentiment using fewer resources than the original BERT model and also provides the higher accuracy.

4.2 Hierarchical Attention Network (HAN) with Knowledge Graph integration

The architecture of the Hierarchical Attention Network (HAN) with Knowledge Graph integration for the aspect-based sentiment analysis begin with the input text which undergoes the tokenization and the embedding is done through BERT embeddings which includes the word, position and token type embeddings. This processed text is then passed through the 12 layers of BERT's transformer blocks each utilizing the multi-head self-attention and dense layers to capture the deep contextual relationships. The output from BERT's [CLS] token is pooled and processed through a dense layer with Tanh activation to generate the comprehensive representation. word-level attention is also applied to the BERT output to focus on the most relevant words within each sentence which is also followed by sentence-level attention to highlight the most important sentences within the document. External knowledge from sources like spaCy is integrated by extracting named entities and averaging their vector representations to create a comprehensive knowledge embedding. This enriched sentence-level attention output is then concatenated with the knowledge embedding which forms the combined representation that uses both the text and external knowledge. This combined vector is processed through a fully connected layer and sigmoid activation function to generate probabilities for sentiment classification. The final output provides sentiment predictions (positive or negative) based on the enriched and processed representations which effectively combines the BERT's deep contextual understanding, hierarchical attention mechanisms and external knowledge for robust aspect-based sentiment analysis

5 Implementation

The aspect based sentiment analysis is implemented through the series of methodical steps where each producing specific outputs and using a variety of tools and libraries to ensure comprehensive analysis and robust model performance. The implementation involved data preparation, exploratory data analysis, model selection using Lazy Predict, and the development of advanced models such as LDA, HAN, and BERT. LIME was also used for model interpretation.

The first step begins with the data preparation where the raw text data containing reviews, aspects, and sentiment labels were cleaned to remove noise such as special characters, numbers, and stop words. The cleaned text was tokenized using the DistilBERT tokenizer and added special tokens like [CLS] and [SEP] for for classification and separating the sentences.

Exploratory data analysis is done to understand the sentiment distribution within the dataset. Fig. 2 shows the imbalance in sentiment distribution with 71.8% positive and 28.2% negative reviews. After resampling as shown in the right pie chart the sentiment distribution is balanced with 50% positive and 50% negative reviews. This ensures more unbiased and effective dataset for training sentiment analysis models.

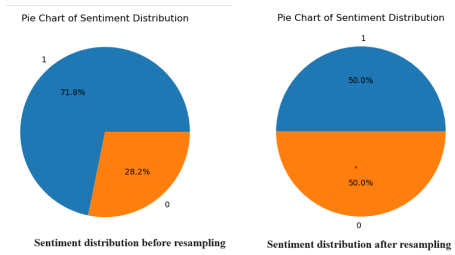

Figure 2: Sentiment distrubution before and after the resampling

In next step various machine learning models were compared using Lazy Predict to assess their performance. The text data was transformed into feature vectors using techniques such as TF-IDF (Term Frequency-Inverse Document Frequency) and Lazy Predict was used to benchmark the performance of different models.

LDA was applied to discover the underlying topics within the text data. Topics are visualized using word clouds by highlighting the most relevant words for each topic and providing insights into the common themes present in the reviews. The Fig. 3shows the output generated by pyLDAvis by providing an interactive and detailed analysis of the topics discovered by the LDA model. The Intertopic Distance Map (left panel) represents the relationships between topics using a two-dimensional plane where each bubble corresponds to the topic. The distance between these bubbles shows how distinct the topics are from one another with overlapping or closely positioned bubbles suggesting similar topics. The size of each bubble shows the proportion of that topic within the entire corpus which helps to identify the most dominant topics. The red bars shows the overall frequency of these terms within the entire corpus and blue bars show their estimated frequency within the selected topic. This dual representation allows for a nuanced understanding of how terms contribute to each topic with the relevance metric (adjustable via the slider) balancing term frequency and exclusivity to the topic. The Marginal Topic Distribution (bottom left) shows the distribution of topics across the corpus as a percentage by providing insights into the prevalence of each topic. This combination of visual elements enables a comprehensive interpretation of the LDA results and the identification of key themes and their importance within the dataset.

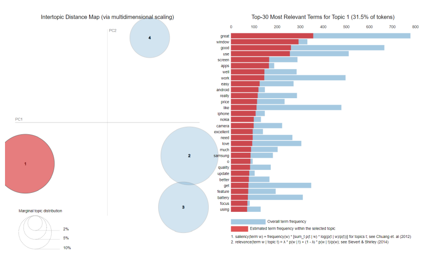

Figure 3: LDA generated by pyLDAvis

A pre-trained DistilBERT model was fine-tuned for sentiment analysis. The model was initialized and trained using the Hugging Face Trainer API managing the training loop, evaluation, and saving of the best model based on accuracy and F1 score. To further enhance the the model Hierarchical Attention Network (HAN) was developed using word-level and sentence-level attention mechanisms. The HAN model was trained on the tokenized and padded text data by producing a trained model with performance metrics that highlighted its ability to understand and classify sentiments based on the hierarchical relationships within the text. Model Explanation was done using LIME to interpret the model's predictions by analyzing the impact of different parts of the input text. This enhanced the interpretability of the model's decisions.

The Fig. 4 shows the LIME output which provides the detailed explanation of the model's sentiment prediction for the review. The left side shows the prediction probabilities by showing a slight preference for a negative sentiment (0.56) over a positive (0.44). Key words contributing to each sentiment are listed with their respective weights: negative words like "ringer," "died," "bother," and "vibrating," and positive words like "reliable" and "purchased." On the right, the review text is highlighted to indicate word influence with orange highlighting positive words such as "reliable," "purchased," and "cover," and blue highlighting negative words like "died," "snake," "ringer," and "bother." This visualization helps to understand which specific words influenced the model's prediction by providing transparency and insight into the model's decision-making process.

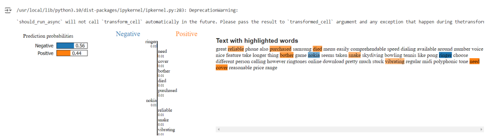

Figure 4: Model interpretation using LIME

5.1 Tools and Languages used

Throughout the project various tools and libraries are used. Python was the primary programming language with key libraries including Hugging Face Transformers for handling BERT models, PyTorch for deep learning frameworks, pandas and numpy for data manipulation, matplotlib for plotting visualizations, sklearn for machine learning tasks, Lazy Predict for rapid model comparison, and LIME for model interpretability. The development was primarily carried out in Google Colab and Jupyter Notebooks by providing an interactive environment for code execution and visualization.

6 Evaluation

6.1 Case study 1: Aspect-Based Sentiment Analysis (ABSA) using Pre-trained BERT Model

The aspect-based sentiment analysis using the BERT pretrained model shows the robust performance which is seen in the Fig. 5 a) shows the high overall accuracy of 94% and strong metrics across the board in the classification report. The model exhibited a precision of 95% and the recall of 98% for the positive sentiment class which shows that model is highly effective in identifying the positive sentiments within the reviews. The negative sentiment class shows the slightly less robust with a recall of 79% but maintained the precision of 92%. This suggests that the model is particularly adept at recognizing positive sentiments, which may be due to the nature of the training data or inherent biases in the pretrained model. The training and validation loss graph shown in the Fig. 5 b) supports these findings by showing the consistent decrease in loss without significant overfitting implying that the model was well-tuned during training. But the confusion matrix that some challenges in misclassification of 31 negative reviews as positive. This indicates that while the model generally performs well but it struggles with certain nuances in negative sentiment detection which leads to incorrect classifications when the sentiment is more context-dependent.

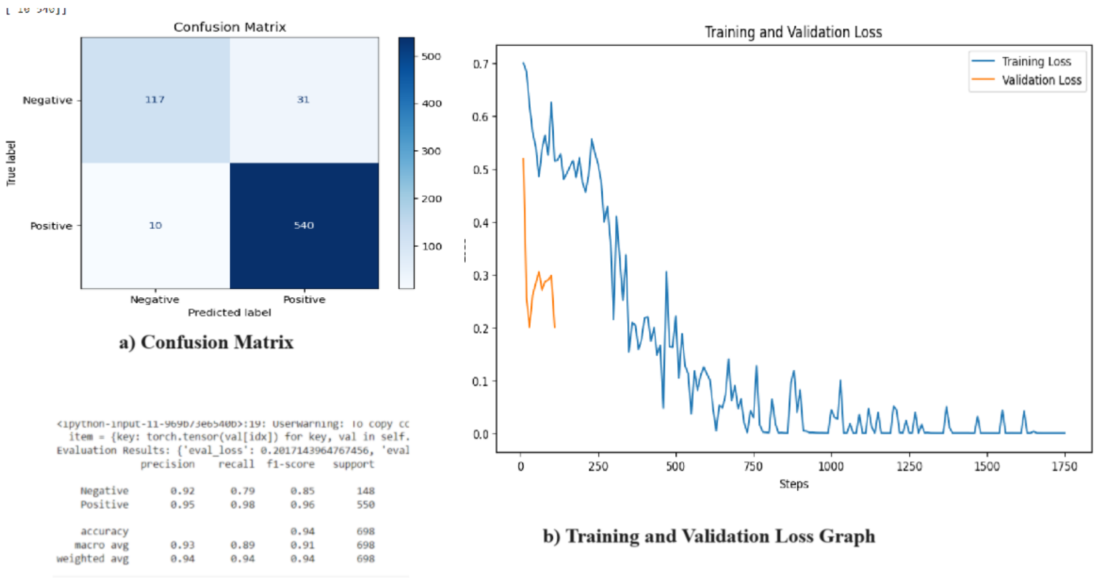

Figure 5: Confusion matrix and Training and validation loss graph for BERT

6.2 Case Study 2: Aspect-Based Sentiment Analysis (ABSA) using Hierarchical Attention Network (HAN) with Knowledge Graph and Pre-trained BERT Model

The Hierarchical Attention Network (HAN) model with integrated Knowledge Graph and pre-trained BERT for Aspect-Based Sentiment Analysis (ABSA) shows the exceptional performance across various evaluation metrics. The model achieves a precision of 0.92 for the negative class and 0.99 for the positive class which shows that the model is highly effective in correctly identifying the sentiments associated with different aspects. The recall values are 0.69 for the negative class and 1.00 for the positive class showing that the model captures almost all positive sentiments while being slightly less effective at identifying all negative sentiments. The overall accuracy of the model is 0.99 which reflects its strong performance in accurately classifying sentiments on the large dataset. The confusion matrix in the Fig. 6 further supports these findings with only a small number of negative examples being misclassified as positive and a very low number of positive examples being misclassified as negative. This high level of performance indicates that the model is robust in handling the nuances of aspect-based sentiment analysis in the scenarios with a large proportion of positive reviews.

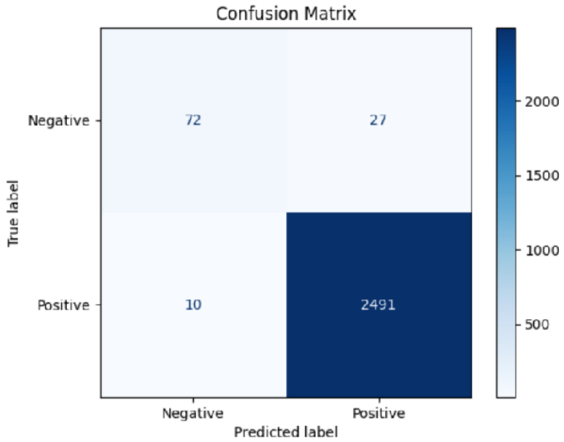

	precision	recall	f1-score	support
Negative	0.92	0.69	0.79	99
Positive	0.99	1.00	0.99	2501
accuracy			0.99	2600
macro avg	0.95	0.84	0.89	2600
weighted avg	0.99	0.99	0.98	2600

Figure 6: Confusion matrix and performance evaluation of HAN and BERT model

6.3 Discussion

The results from the aspect-based sentiment analysis using the BERT pretrained model are largely positive by showcasing the model's strong capability in handling complex sentiment tasks. The high accuracy and strong F1-scores suggest that the model correctly captures the sentiment of various aspects within reviews which is a significant achievement given the complexity of natural language understanding. The model's ability to correctly classify positive sentiments with high precision and recall makes it the valuable tool for applications where understanding customer satisfaction and positive feedback is important. Despite these strengths there are some limitations. The model's tendency to misclassify certain negative sentiments as positive shows the areas for improvement. These misclassifications occurred in reviews where the sentiment was more nuanced or context-dependent which suggests that the model is proficient in handling clear sentiment expressions but it need further refinement to accurately capture more subtle or complex negative sentiments. This can be addressed by augmenting the training data with more diverse examples of negative sentiment or by fine-tuning the model further on a dataset designed to highlight these nuances.

The combination of the Hierarchical Attention Network (HAN), Knowledge Graph and pre-trained BERT model is highly effective for aspect-based sentiment analysis. By using the Knowledge Graph the model uses the additional structured information that helps it better to understand the context and relationships between different aspects which leads to the more accurate sentiment classification. The fine-tuned BERT model adds a deep understanding of language by allowing the model to interpret and classify sentiments at a detailed level. This combined approach outperforms BERT model used in the previous case especially in complex texts. But the model is slightly less effective in identifying the negative sentiments. Future work can focus on using the techniques to make the model more sensitive to less common negative sentiments. This approach represents the significant improvement in ABSA by offering the high accuracy and precision for positive sentiments making it a valuable tool for understanding customer feedback in detail.

7 Conclusion and Future Work

In conclusion the integration of the Hierarchical Attention Network with Knowledge Graph and pre-trained BERT model shows the advancement in aspect-based sentiment analysis . This hybrid approach combines the contextual understanding of BERT, the hierarchical text processing of HAN, and the structured external knowledge from the Knowledge Graph which results in the good performance in recognizing positive sentiments. The model's ability to use external knowledge increases the capability of the model to capture complex aspect relationships and context-specific nuances which are the challenging for other models. But this combined model exhibited the slightly lower recall for the negative class which indicates the potential areas for improvement such as addressing class imbalance through data augmentation and advanced class weighting techniques. Future work can focus on refining these aspects to improve the detection of less frequent sentiments by enhancing the model's overall robustness and adaptability across diverse application domains.

References

Aufar, M., Andreswari, R. and Pramesti, D. (n.d.). Sentiment analysis on youtube social media using decision tree and random forest algorithm: A case study.

Bhattacharyya, P. (2013). Sentiment analysis, *2013 1st International Conference on Emerging Trends and Applications in Computer Science*, pp. i–i.

Devlin, J., Chang, M.-W., Lee, K. and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding.

He, R., Ravula, A., Kanagal, B. and Ainslie, J. (2023). Realformer: Transformer likes residual attention, *arXiv preprint arXiv:2303.07012* .

Joachims, T. (1998). Text categorization with support vector machines: Learning with many relevant features.

Katić, T. and Milićević, N. (2018). Comparing sentiment analysis and document representation methods of amazon reviews, pp. 000283–000286.

Mikolov, T., Chen, K., Corrado, G. and Dean, J. (2013). Efficient estimation of word representations in vector space.

Ng, H., Chong, W. K., Chia, Y. Z., Vun Yap, T. T., Goh, V. T., Wong, L. K., Teck Tan, I. K. and Cher, D. T. (2024). Aspect-based subjectivity analysis using a bert-based approach.

Pang, B., Lee, L. and Vaithyanathan, S. (2002). Thumbs up? sentiment classification using machine learning techniques, Association for Computational Linguistics.

Pennington, J., Socher, R. and Manning, C. D. (n.d.). Glove: Global vectors for word representation.

Radford, A., Narasimhan, K., Salimans, T. and Sutskever, I. (2018). Improving language understanding by generative pre-training.

Ruskanda, F. Z., Widyantoro, D. H. and Purwarianti, A. (2018). Comparative study on language rule based methods for aspect extraction in sentiment analysis.

Safavian, S. and Landgrebe, D. (1991). A survey of decision tree classifier methodology, *IEEE Transactions on Systems, Man, and Cybernetics* .

Wahyudi, E. and Kusumaningrum, R. (n.d.). Aspect based sentiment analysis in e-commerce user reviews using latent dirichlet allocation (lda) and sentiment lexicon.

Zhou, J., Jin, P. and Zhao, J. (2024). Sentiment analysis of online reviews with a hierarchical attention network.