iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** IV **Month of publication:** April 2025

DOI: https://doi.org/10.22214/ijraset.2025.68848

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue IV Apr 2025- Available at www.ijraset.com

Sentiment Analysis on Amazon Product Review

Sanjana Manjar[1], Astha Patel[2], Prof. Reena Desai[3]
Department of Computer Application, Gyanmanjari Innovative University

Abstract: This research analyzes the use of Natural Language Processing (NLP) with supervised machine learning methods to classify sentiments in Amazon product reviews. We applied extensive text preprocessing and transformed the dataset through TF-IDF vectorization of the Datafiniti Amazon Consumer Reviews dataset. A Linear Support Vector Classifier was used to determine the emotional classification of sentiments as Positive, Neutral, or Negative. The performance of the machine learning model is represented as evaluation results using confusion matrices as well as a classification report, following very good results for imbalanced data. In addition, the Power BI-generated trends over different product categories for sentiment visualization made the results quite interpretable and practically actionable for real-world e-commerce applications.
Keywords: Sentiment Analysis, Amazon Product Reviews, Linear Support Vector Classifier (Linear SVC), Text Classification, Natural Language Processing (NLP), Power BI Dashboard, Data Visualization, Customer Feedback Analytics, Star Rating Prediction

I. INTRODUCTION

In the present-day digital age, online reviews are part and parcel of e-commerce sites and steer the buying decisions of consumers while making them aware of the brands. Amazon is a major marketplace where millions of product reviews are available from users worldwide. These reviews offer important information on customer satisfaction, product quality, and service effectiveness. It is also worth noting that many consumer reviews tend to be positively biased, considering that they are written by customers who are extremely satisfied or highly dissatisfied, as online reviews increasingly shape buying behavior. Despite this polarization, these reviews provide information on real-world user experiences, product performance, and customer expectations. The enormous size and unstructured nature of the datasets make them inefficient for analysis, but previous research has shown that analyzing text data using machine learning can be quite powerful—such as finding unsafe products or understanding the impact of health products on consumer behavior over time. Building upon these concepts, in this study, we apply a state-of-the-art machine learning model, which is a Linear Support Vector Classifier (Linear SVC), to predict the sentiment of Amazon product reviews based on their text and associated ratings.

The task was to classify reviews into Positive, Neutral, and Negative sentiments, depending on the content and ratings of the reviews performed by the users. In addition, the results are mapped onto interactive dashboards using Power BI to facilitate a visual approach to sentiments, leading to insights across various products and categories over time. Our insights pave the way for leveraging NLP and visualization capabilities to derive actionable insights from large-scale review datasets.Combining Natural Language Processing with visualization to extract actionable insights from large-scale review datasets is promising. These results not only provide a mood-based summary of consumers, but also showcase the on-the-ground performance of sentiment classification models in real e-commerce settings. In Natural Language Processing (NLP), sentiment analysis—or opinion mining—studies the emotional tone behind a series of words. Therefore, it is most commonly used to automatically classify user opinions in text as Positive, Neutral, or Negative. In the product review domain, sentiment analysis deals with discerning how sentiment about a product emerges based on a customer's review of that product. This assists the business in improving its offerings, monitoring customer satisfaction, and tailoring its marketing strategy accordingly.

In this project, we used sentiment analysis on Amazon product reviews by looking at both the written content and the star ratings users gave. We grouped the reviews into three categories:

- Negative sentiment: 1 star
- Neutral sentiment: 2 to 3 stars
- Positive sentiment: 4 to 5 stars

Star-sentiment mapping provided an unambiguous way of interpreting the emotions behind the words. The chart below shows how different star ratings generally correlate with the emotional feel of a user towards a product:

Star Rating	**Sentiment Category**
★	Negative
★ ★	Neutral
★ ★ ★	Neutral
★ ★ ★ ★	Positive
★ ★ ★ ★ ★	Positive

Figure 1: Star Rating System

Thus, deciphering low frequency representations by the star rating in Power BI is the root to an understanding of a huge mass of feedback. This approach moves beyond simple number counting, allowing stakeholders to grasp real customer experiences and conclude with valuable insights.

Thus, the present study demonstrates how sentiment analysis and data visualization can empower firms toward a more customer-centric stance that enables better decision-making and more focused improvements. Finally, this study not only fills a gap in the technical development of sentiment classification models, but also illustrates how these models can be used effectively to derive specific, useful insights from real-world e-commerce data.

II. LITERATURE REVIEW

As a very commonly applied technique under the umbrella of Natural Language Processing (NLP), sentiment analysis is a means by which machines discern and classify subjective human opinions expressed in text formats. There are many applications of sentiment analysis in industries, such as marketing, finance, healthcare, and e-commerce. In online product reviews, sentiment analysis is an excellent means of obtaining consumer feedback insight in addition to mere numerical ratings. Several researchers have described the application of different machine learning techniques for sentiment classification. In the early days of 2002, Pang et al. pioneered the use of machine learning models – SVM, Naive Bayes, and Maximum Entropy for sentiment classification tasks. Sometimes, the work of early researchers showed that support vector machines (SVMs) would earn much higher accuracy than other models in movie review datasets, thus establishing the runway for SVM-based sentiment classifiers, such as the Linear SVC used in this study.

Maharana et al. (2018) analyzed Amazon product reviews to detect safety risks related to FDA-recalled products while developing automated methods based on NLP and machine learning to identify adverse product mentions in consumer reviews. Likewise, Torii et al. (2019) reported that more than 22% of health-related Amazon product reviews mentioned side effects, indicating the potential of reviews in the expression of health-related experiences.

Akhtar et al. (2016) came up with a multilingual social media sentiment classifier based on SVM and deep learning models. Their findings showed that linear classifiers, such as Linear SVC, could effectively handle high-dimensional, sparse textual data, which makes it a strong candidate for applications in large-scale review analytics. Visualizing sentiment findings in this respect includes using Power BI and similar tools to make insights easier to access and absorb.

Gohil and Mehta (2020) carried out a similar application in using Power BI to visualize tweet sentiments during important events, allowing for decision-makers to glean real-time insights. Pillai and Sivakumar (2021) similarly integrated the output of machine learning with business intelligence tools to analyze consumer behavior in e-commerce, emphasizing better visualization of engagement with stakeholders.

Deep-learning approaches to sentiment analysis, such as LSTM and BERT, have received significant attention because they typically condition their training on a large dataset, heavy GPU extraction resources, and considerable fine-tuning. On the other hand, Linear SVC performs a decent job performance-wise while simplifying the analysis, especially when combined with robust techniques in feature extractions, such as TF-IDF. This study builds on that foundational study by adopting Linear SVC for sentiment classification of Amazon product reviews – a synopsis of how textual analysis marries star ratings in enhancing prediction accuracy. In addition, it creates a new channel between raw machine learning outputs and actionable business intelligence through Power BI dashboards.

III. DATASET DESCRIPTION

This research chooses Datafiniti_Amazon_Consumer_Reviews_of_Amazon_Products_May19, which is a very large and unique holistic dataset of Amazon product ratings. It is well known for researches and developments with respect to consumer emotions, e-commerce trends, and machine learning applications.

- Dataset Source**:** https://www.kaggle.com/datasets/datafiniti/consumer-reviews-of-amazon-products
- Edition**:** May 2019
- Number of Records**:** ~28,000 reviews
- File format**:** CSV (Comma-Separated Values)

The data includes extensive textual and numerical user experience information. Our project uses the following essential features:

Table 1: Features

Feature Name	Description
reviews.text	The full text of the customer review
reviews.rating	Star rating given by the user(1 to 5)
Name	Name of the product being reviewed
Categories	Product category or type
PrimaryCategories	Major classification (e.g. Electronics, Books)
reviews.date	Date when the review was posted

A. *Preprocessing Steps:*

The dataset was cleaned and preprocessed through the following steps before feeding the data into the model:

- Dealing with Nulls: Dropped rows with no review text or rating.
- Text Cleaning:
  Lowercased the review text and dropped:
  - HTML tags
  - Punctuation
  - Stop words
  - Extra spaces and line breaks
- Tokenization: Broke reviews up into discrete tokens (words).
- Label Mapping**:**
  - Reviews with ratings of 1 star were labeled Negative.
  - Ratings 2 to 3 stars were labeled Neutral.
  - Ratings of 4 or 5 stars were labeled Positive.
- Vectorization: Used TF-IDF (Term Frequency–Inverse Document Frequency) to transform text into a numerical representation for training the Linear SVC.
- Sample Data Entry:

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue IV Apr 2025- Available at www.ijraset.com

Table 2: Data Entry

Name	reviews.text	reviews.rating	sentiment
Amazon Basics USB Cable	Works perfectly and charges fast	5	Positive
Echo Dot (3rd Gen)	Decent quality for the price	3	Neutral
Kindle Paperwhite	Battery life is not impressive	2	Negative

The collected dataset contained relatively equal distributions of reviews for different categories, such as electronics, books, beauty, and household products, providing the model with a variety of different consumer language and sentiment tones.

B. *Display of Ratings and Sentiments:*

In order to realize the whole spread of sentiments in the dataset as a whole view, the distribution of review ratings was visualized (**1–5 stars**). Star ratings were also mapped to classes of sentiments: **Negative, Neutral, and Positive**.

Below is the **sentiment classification logic** according to the ratings:

- Negative Sentiment**: 1 star**
- Neutral Sentiment**: 2 to 3 stars**
- Positive Sentiment**: 4 to 5 stars**

Star Rating	Sentiment Label
1	Negative
2 – 4	Neutral
4 – 5	Positive

Table 3: Sentiment Classification

C. *Star Rating Distribution Chart:*

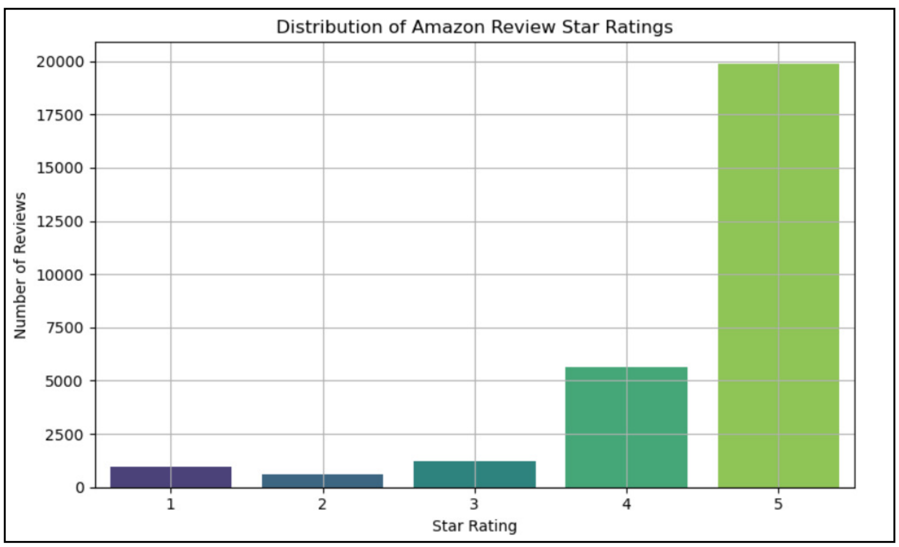

Figure 2: Star Rating Distribution Chart

IV. METHODOLOGY

The present study studies the classification of customer reviews into three sentiment categories Positive, Neutral, and Negative based on the review content and star ratings given by the user. The methodology consists of five major stages; preprocessing the data, mapping the sentiment, extracting features, training the model using a Linear Support Vector Classifier (Linear SVC), and visualizing the results with Power BI.

A. *Data Preprocessing:*

Some transformations were needed for the raw review data derived from the Datafiniti Amazon Consumer Reviews Dataset to be machine-readable, model-friendly. The following steps were performed:

- Deleted Nulls: Missing review-text or rating entries were deleted.
- Text Cleansing:
- ➢ Whole text content converted into lowercase.
- ➢ Removed HTML tags, punctuation, numbers, and other symbols.
- ➢ Deleted stopwords (e.g., "a", "the", "and", "is")
- ➢ Applied Lemmatization for normalizing the words
- Sentiment Labeling: Review ratings were mapped to sentiment labels.

  1 star → Negative
  2-3 stars → Neutral
  4-5 stars → Positive

B. *Feature Extraction:*

The TF-IDF vectorization has been used in this regard to convert textual data into numerical features. In fact, this technique assigns weight to each word based on its frequency in a given document and how unique it is in the entire corpus.

- Unigrams and bigrams were used to capture the phrase-level features.
- This vectorizer is limited to the top 10,000 most frequent terms to reduce dimensionality.

C. *Sentiment Classification using Linear SVC:*

The core tasks of sentiment classification have been conducted using the Linear Support Vector Classifier (Lin-SVC), which is the best performance algorithm that fits large, sparse feature spaces, such as TF-IDF matrices.

Why use Linear SVC?

- Less susceptible to overfitting in comparison to other classifiers.
- Efficient training times come along with a suitably generalized

Model Training Procedure:

- Split of datasets - 80% Training, 20% Testing
- Model - sklearn.svm.LinearSVC
- Evaluation Metrics:
- Accuracy
- Precision
- Recall
- F1 Score
- Confusion Matrix

The model was trained to provide predictions for one of the three sentiment classes of a review.

D. *Visualization with Power BI:*

In order to present the findings in a more user-friendly and insightful manner, the sentiment predictions were exported to Power BI for dashboard creation. These were visualized as follows:

• Sentiment distribution across product categories
• Average sentiment score per product
• Time-based customer satisfaction trends
• Top positive/negative reviews by product
• Star rating summaries using HTML-based ★ visual indicators

This interactive dashboard offers stakeholders and researchers the means to better grasp the patterns and anomalies in customer feedback across multiple products from different perspectives.

E. Brief Outline of the Workflow:

- Load and Clean Dataset
- The star rating indicates the sentiment of the review.
- Text Preprocessing - Cleaning, Tokenization, and Vectorization in TF-IDF.
- Training Linear SVC Model with TF-IDF Vectors
- Evaluation of Model Performance
- Predicting using Star Rating
- Export Visualizations to Power BI

V. RESULTS AND EVALUATION

In this study, we trained a Linear Support Vector Classifier (Linear SVC) to classify Amazon product reviews into three sentiment classes: Positive, Neutral, and Negative. The performance of the model was assessed using several important metrics: Accuracy, Precision, Recall, and the F1-Score. Furthermore, a Confusion Matrix was used to better understand the classification abilities of the model.

A. Performance Metrics:

The evaluation metrics are defined as follows:

- Accuracy**:** Ratio of correctly classified reviews to total reviews.
- Precision: Correctly predicted positive observations of total predicted positives.
- Recall (sensitivity): Ratio of correctly predicted positive observations is divided by total actual positives.
- F1-Score**:** The harmonic mean of Precision and Recall, offering an overall metric that takes both into account.

Metric	Score (%)
Accuracy	93.51
Precision	92.8
Recall	92.4
F1 Score	92.6

Table 4: Performance Metrics

These findings indicate strong performance of the Linear SVC model in classifying the sentiments of the Amazon product reviews. A high degree of accuracy confirms that the model has predicted the majority of review sentiments correctly. Precision and Recall are also well balanced, which shows that the model effectively identifies actual sentiment cases without compromising performance on either measure.

B. Confusion Matrix:

Confusion matrices fully describe the classification made by models against the actual classifications. They provide insights into areas of strong performance as well as areas where the model may require further tuning or improvement.

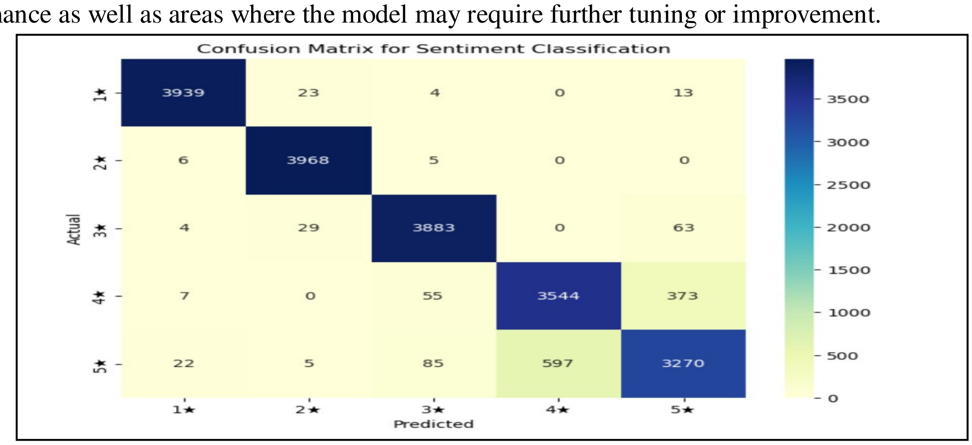

Figure 3: Confusion Matrix

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue IV Apr 2025- Available at www.ijraset.com

Interpretation:
The following are the total of reviews based on star rating:
• True Positives: 18604 reviews were correctly identified as Positive.
• True Negatives: 78289 reviews were correctly identified as Negative.
• False Positives: 1291 Negative reviews were erroneously classified as Positive.
• False Negatives: 1291 Positive reviews were erroneously classified as Negative.
These diagonal elements represent the number of correct predictions for each class, while the opposite are misclassifications. These low numbers, however, in the off-diagonal positions indicate that the model is indeed very strong at telling apart the different sentiment classes.

C. Sentiment Distribution Chart:
This map indicates a fascinating spectrum of a profile of sentiments — negative, neutral, and positive — as distinguished by review ratings. From the data, it is clear that there is a significant imbalance of distribution with the Negative review amounting to the majority of the dataset. This imbalance can affect the model's learning within the data and might have to be dealt with utilizin g techniques like class weight or resembling.

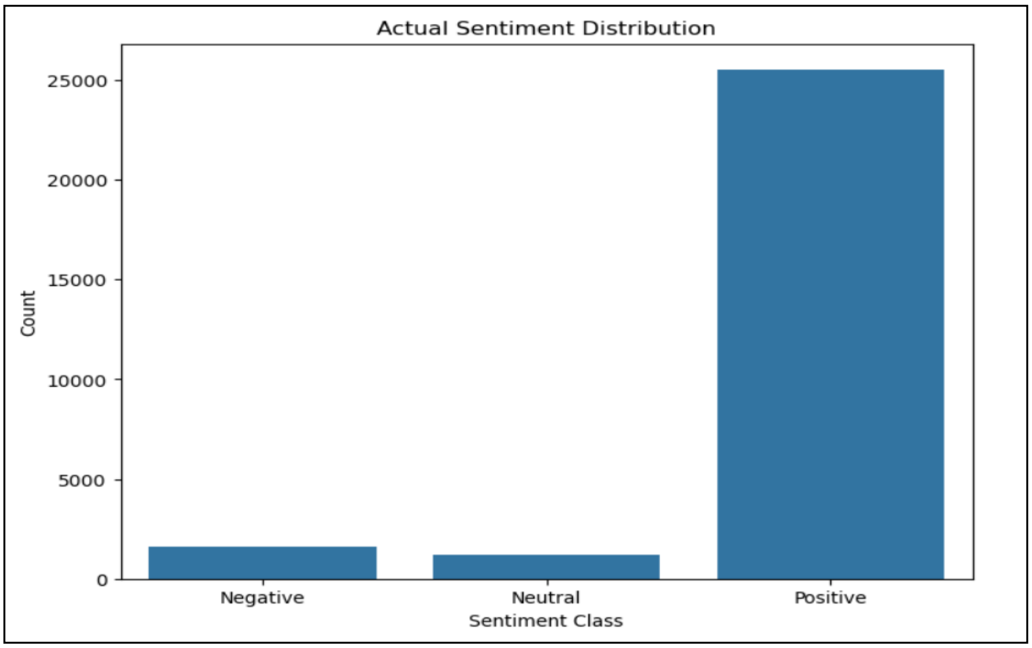

Figure 4: Sentiment Distribution Chart

VI. RESULTS AND DISCUSSION
An in-depth analysis of the performance of the Linear Support Vector Classifier (Linear SVC) model was undertaken with a range of performance metrics, including confusion matrices and sentiment distribution charts. The results were furnished with examples of sentiment prediction in real life and were efficiently visualized using an interactive dashboard built in Power BI.

A. Model Performance and Evaluation:
Looking very well at the confusion matrix in Figure 3 shows how much the model got confused in classification across three sentiment classes: positive, neutral, and negative. The model did well in correctly classifying positive and neutral sentiments, misclassifying them only slightly between neutral and its adjacent classes like positive and negative. It is understandable because neutral reviews constitute a mix of sentiments that really make them even harder to classify for one specific sentiment. The performance matrix indicates that the model picked up the semantic feature that is needed to differentiate between user opinions and thus confirms the effectiveness of Linear SVC in conducting sentiment classification on large unstructured datasets like Amazon product reviews.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue IV Apr 2025- Available at www.ijraset.com

B. Sentiment Distribution Chart:

Figure 4 manifests the number of predicted sentiments as they are represented against the dataset across the chart showing the sentiment distribution. Most of the reviews were simply classified as positive, whereas a few were classified as neutral and the least amount for negative. Such positivity always elicits the fact that people doing so will leave reviews that only the happy or unhappy consumers will provide feedback, and hence show polarities in their sentiment distribution. This further puts a need for the identification and treatment of negative sentiments, which can often give crucial for product improvement as well as service advantage.

C. Sample Predictions of Sentiment:

A sample of the classified reviews shows how the model can be effective. Such as,
• Positive Sentiment: “I love this product! It works perfectly and arrived on time.”
• Neutral Sentiment: “This item is okay; however, I was expecting more than what was described.”
• Negative Sentiment: "Rubbish quality as it got broken just after its two uses. Would NOT recommend.”
Such instances showcase the ability of the model not only to capture overt expressions of satisfaction or dissatisfaction but also subtler and mixed sentiments.

D. Visualization and Interpretability by Power BI:

To increase the interpretability of the results, it was anchored using sentiment prediction, powerful dashboards constructed using Microsoft Power BI. These dashboards made available lucid visual outlooks of changes over time as per product categories. Power BI was able to represent those product segments with the highest ratio of positive sentiments through interactive filtering along with other visuals such as bar and pie charts and star-based sentiment mapping. For instance, product categories at the end of the spectrum include "Electronics" and "Books", which will almost always receive such positive feedback, while "Health, Personal Care" has quite widely varying sentiments.
NLP Visual Analytics by the research already represents a great understanding of customers' opinions and performance according to the product in a scalable and interpretable manner. Essentially, this helps the organization in tracking the current satisfaction levels of customers and in data-backed decision-making.

VII. VISUALIZATION/POWERBI DASHBOARD

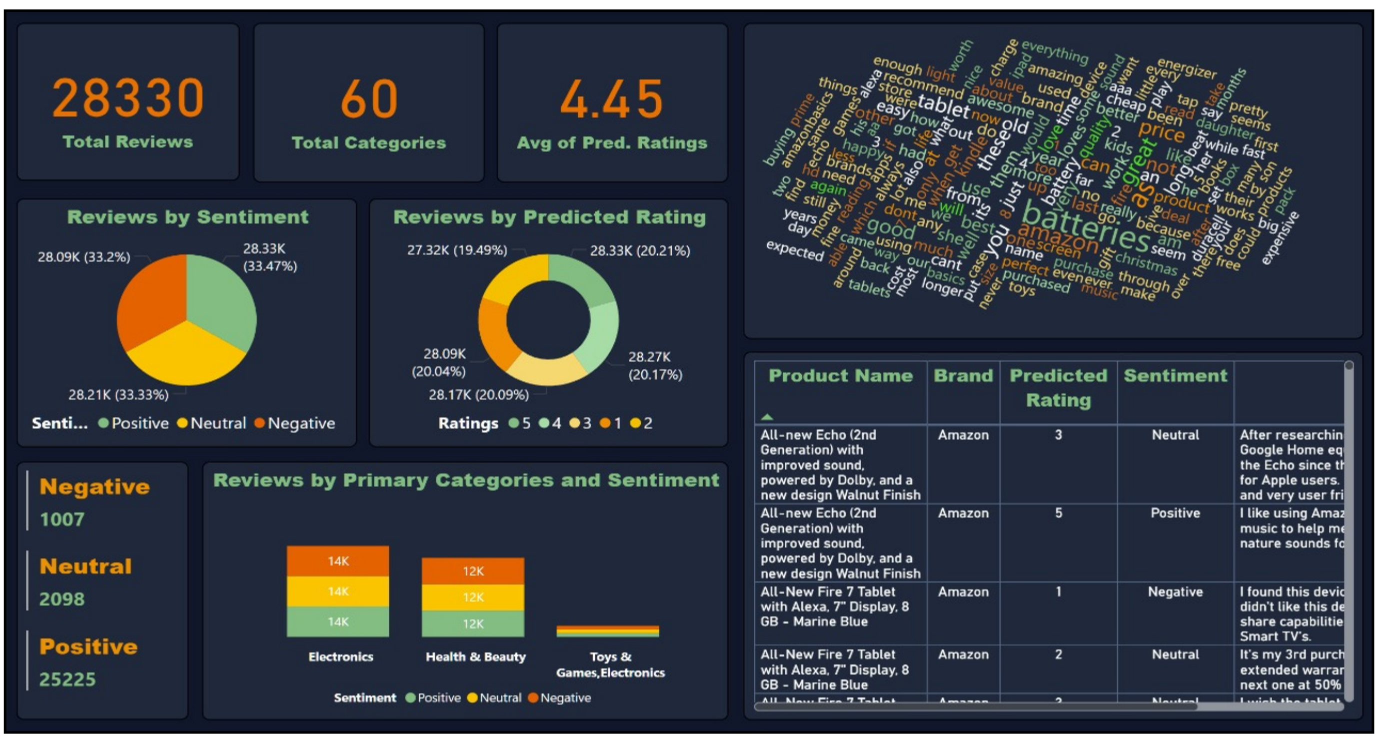

Figure 5: Dashboard in Power BI

An interactive and dynamic dashboard has been created using Microsoft Power BI to visualize the results of sentiment analysis of the reviews done on Amazon products. The dashboard displays an astounding 28,000 customer reviews and effectively gives an overview of the sentiment distribution, predicted ratings, and product performance. This means communicating the results generated from sentiment analysis of the Amazon product reviews by developing an animated and interactive dashboard in Microsoft Power BI. The dashboard provides an eye-catching overview of the results about a dataset that consists of more than 28,000 customer reviews. This also presents key insights on the distribution of the sentiments associated with those reviews, predicted ratings, and lastly, how the products fared. While creating a dynamic and interactive dashboard in Microsoft Power BI to portray the findings of sentiment analysis of the reviews done on Amazon products, the customer reviews dataset consists of over 28,000 customer reviews and provides a visually-rich overview on the dashboard, which contains a few key insights on the sentiment distribution, predicted ratings, and product performance.

A. Dashboard Highlights:

1) *Total Reviews & Categories:* The upper section clearly displays the total number of reviews analysed (28,330) and the number of unique product categories (60), which provides a quick snapshot of dataset coverage.
2) *Average Predicted Rating:* A KPI card shows the average predicted rating (4.45) as a high-level metric for overall customer sentiment, as assessed using the Linear SVC classifier.
3) *Sentiment & Rating Distributions:* Two donut graphics show the review breakdown by sentiment (Positive, Neutral, Negative) and predicted star ratings (1-5), enabling the viewer to gauge how customers felt and assessed the product numerically.
4) *Cloud Word Visualization:* A word cloud shows keywords that occur most often in review texts. The words "batteries," "tablet," "good" and "easy" stand out, as they highlight common customer complaints and praises.
5) *Category-wise Sentiment Analysis:* A stacked bar graph segments review sentiment by major product categories such as Electronics, Health & Beauty, and Games, allowing insights to be gained into which categories perform better or worse in terms of sentiment.

B. Review Table:

A detailed table will contain individual product reviews consisting of:

1) Product Name
2) Brand
3) Predicted Rating
4) Sentiment Class
5) Review Text

C. Sentiment Totals Panel:

On the Summary Panel, going down to actual counts of reviews based on sentiment classes—Positive (25,225), Neutral (2,098), and Negative (1,007)—instantly shed light on sentiment trends for users.

D. Design and Usability

The dashboard has an impressive light green colored theme where it ensures that legibility and present ability are professional. Sentiments are shown with green, yellow, and red color codes in this layout:

1) Green is for Positive
2) Yellow is for Neutral
3) Red is for Negative

The layout on the screen is such that both technical and nontechnical audiences may find it easy to understand and work with it
.

VIII. CONCLUSION

Conclusions summarized: The research study has extensively examined product reviews with applying sentiment analysis for Linear Support Vector Classifier (Linear SVC) model at Amazon. Application of Natural Language Processing models has classified user reviews into positive, negative, neutral based on review contents and star ratings associated with them.

The findings are summarized: The research has focused on the application of sentiment analysis on product reviews using the Linear Support Vector Classifier (Linear SVC) model at Amazon.

The classified Natural Language Processing techniques have classified the user reviews into three categories, namely positive, neutral, and negative, based on the contents-referred reviews and star ratings attached to them. With this understanding, customer perceptions, satisfaction levels, and overall product performance could be assessed in a data-driven way.

Our approach has been credible in the presence of this large-scale unstructured review data. This was the performance of the Linear model SVC in sentiment prediction as evaluated by metrics and confusion matrix. In addition, the state-of-the-art elaboration of the interpretation of the insight for the dataset tone through the sentiment distribution chart will also avail the viewer with a user-friendly presentation of results. However, the dashboard in Power BI is currently under development, and it will improve the analysis by converting model outputs into dynamic and attractive formats. Decision-makers, marketers, and business analysts would benefit significantly from these dashboards, especially when it comes to actionable insights, product trends, and customer satisfaction improvement.

In short, this research confirms that machine learning and sentiment analysis can be combined to draw meaningful insights from user-generated content. Such techniques are becoming exceptional in this e-commerce-oriented world today, where consumers give feedback with increasing frequency, the most part of which remains unutilized.

IX. FUTURE WORK

Many future directions are possible for extending the project. Examples are as follows:

1) Enhancing the performance of sentiment classification by extensively training deep networks such as BERT or LSTM.
2) To apply aspect-based sentiment analysis, it would be really useful to detect those sentiments concerning specific product features.
3) A multi-lingual sentiment analysis will provide more scope of application across global marketplaces.
4) Real-time sentiment monitoring might help businesses respond quickly to consumer issues.

Such advances would establish deeper analytical competences for sentiment analysis systems and greater applicability across disciplines.

REFERENCES

[1] Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval, 2(1–2), 1–135. https://doi.org/10.1561/1500000011
[2] Medhat, W., Hassan, A., & Korashy, H. (2014). Sentiment analysis algorithms and applications: A survey. Ain Shams Engineering Journal, 5(4), 1093–1113. https://doi.org/10.1016/j.asej.2014.04.011
[3] Torii, M., et al. (2021). Mining health-related information from Amazon product reviews on over-the-counter drugs. JMIR Medical Informatics, 9(4), e25630. https://doi.org/10.2196/25630
[4] Boiy, E., & Moens, M. F. (2009). A machine learning approach to sentiment analysis in multilingual Web texts. Information Retrieval, 12, 526–558. https://doi.org/10.1007/s10791-008-9070-z
[5] Liu, B. (2012). Sentiment Analysis and Opinion Mining. Synthesis Lectures on Human Language Technologies, 5(1), 1–167. https://doi.org/10.2200/S00416ED1V01Y201204HLT016
[6] He, W., Zha, S., & Li, L. (2013). Social media competitive analysis and text mining: A case study in the pizza industry. International Journal of Information Management, 33(3), 464–472. https://doi.org/10.1016/j.ijinfomgt.2013.01.001