ISSN: 2454-9940

INTERNATIONAL JOURNAL OF APPLIED SCIENCE ENGINEERING AND MANAGEMENT

E-Mail :
editor.ijasem@gmail.com
editor@ijasem.org

www.ijasem.org

INTERNATIONAL JOURNAL OF APPLIED SCIENCE ENGINEERING AND MANAGEMENT

ISSN 2454-9940

www.ijasem.org

Vol 19, Issue 1, 2025

Public Relations Insights: Text Classification and Sentiment Analysis

T. Shalini [1], Gorre Himabindu[2], Bugatha Swetha [3], Mathi Bhavyasritulasi[4], Ankam Mounika [5]

[1] Assistant Professor, Dept. of Computer Science & Engineering, Vijaya Institute of Technology for Women, Enikepadu, Vijayawada-521108

[2,3,4,5] Students, Dept. of Computer Science & Engineering, Vijaya Institute of Technology for Women, Enikepadu, Vijayawada-521108

Email id: shalinitammana09@gmail.com[1], himabindugorre7965@gmail.com[2], bugataswetha456@gmail.com[3], bhavyasritulasimathi@gmail.com4, mounikaankam57@gmail.com5,

Abstract:

The aim of the project to perform sentiment analysis, a crucial task in natural language processing, to understand the sentiments expressed in customer reviews of Amazon Alexa products. Sentiment analysis involves determining whether a piece of text expresses positive or negative. Leveraging Natural Processing Language (NLP) along with machine learning algorithms such as Naïve Bayes and Logistic Regression. These algorithms are commonly used in sentiment analysis tasks due to their effectiveness in classifying text data into different sentiment categories. The dataset consists of customer reviews of Amazon Alexa products, which likely contain a mix of positive and negative sentiments. By training the machine learning models on this dataset, we are categorizing reviews as either positive or negative by using sentiment. To evaluate the performance of the sentiment analysis model, we are using various metrics including accuracy, precision, recall, and F1-score. Accuracy measures the overall correctness of the model's predictions, while precision measures the proportion of correctly predicted positive reviews out of all reviews classified as positive. Recall measures the proportion of correctly predicted positive reviews out of all actual positive reviews, and F1-score is the harmonic mean of precision and recall, providing a balanced assessment of the model's performance. These metrics together will help assess the effectiveness of the sentiment analysis model in accurately capturing the sentiments expressed in Amazon Alexa product reviews.

Keywords – Sentiment Analysis, Amazon Alexa Reviews, Machine Learning, NLTK, word Cloud.

1.Introduction

In recent years, the proliferation of voice-controlled virtual assistants has revolutionized the way individuals interact with smart devices. Among these, Amazon Alexa has emerged as a prominent and widely adopted virtual assistant, powering an extensive range of smart home devices and applications. As the popularity of Amazon Alexa continues to soar, the importance of understanding customer sentiments and opinions towards the product becomes paramount for businesses and developers.

Customer reviews play a pivotal role in shaping product perception and influencing purchasing decisions. Analyzing these reviews provides valuable insights into customer satisfaction, identifies potential pain points, and offers opportunities for product enhancements. However, manually processing a large volume of customer feedback is a daunting and time-consuming task.

Sentiment analysis, a subfield of natural language processing (NLP), presents an automated solution to this challenge by automatically classifying text into positive, negative, or neutral sentiments. This project aims to apply sentiment analysis techniques to customer reviews of Amazon Alexa products

using machine learning algorithms. The primary objective is to build a sentiment analysis model capable of accurately classifying customer sentiments expressed in the reviews. The research leverages state-of-the-art machine learning models and NLP techniques to extract valuable insights from textual data.

In this methodology consists of several key steps. First, a dataset of Amazon Alexa customer reviews is collected from a reputable online source. The dataset is then pre-processed to remove irrelevant information, correct spelling and grammar errors, and convert all text to lowercase. Next, feature extraction techniques such as Word Cloud, Tokenization and Removing Stop words are applied to convert the textual data into numerical vectors that can be input into machine learning models.

2.Literature Review

Sentiment analysis has gained substantial attention in the field of natural language processing (NLP) and machine learning due to its applications in understanding human emotions, opinions, and attitudes expressed in textual data. They have contributed significantly to the development of sentiment analysis techniques and methodologies. Zhang et al. (2018) explored sentiment analysis using a machine learning approach, focusing on the classification of online reviews into positive and negative sentiments. Their study utilized Naïve Bayes Algorithm and Logistic Regression to capture both local and global textual features for improved sentiment classification accuracy. The results demonstrated the effectiveness of machine learning models in handling complex linguistic patterns for sentiment prediction. In a study by Kim (2014), a simple yet effective approach to sentiment analysis was proposed using Convolutional Neural Networks (CNNs) applied to sentence-level classification tasks. The research showed that CNNs can automatically learn hierarchical features from textual data, leading to state-of-the-art performance in sentiment classification tasks. The application of CNNs in this context highlighted their adaptability and robustness to varying lengths of input text. Traditional machine learning algorithms have also contributed significantly to sentiment analysis. Pang and Lee (2008) investigated the application of Support Vector Machines (SVMs) to sentiment classification, emphasizing the importance of feature selection and preprocessing techniques. Their research emphasized the significance of domain-specific features and the effectiveness of supervised learning algorithms in handling sentiment analysis tasks. Furthermore, prior work has explored feature engineering techniques to enhance sentiment classification accuracy. Turney (2002) introduced the concept of using pointwise mutual information to identify and extract sentiment-bearing words. This approach showcased the potential of feature engineering in sentiment analysis, providing insights into the significance of domain-specific lexicons for accurate sentiment prediction. The literature suggests a diverse range of techniques, including deep learning models, traditional machine learning algorithms, and feature engineering methods, to address sentiment analysis tasks. While deep learning approaches have demonstrated impressive performance, traditional algorithms continue to offer valuable insights into feature extraction and model interpretability.

Sentiment Analysis for Amazon Reviews

Sentiment analysis has emerged as an increasingly popular research area in computational linguistics and text mining. Our study aims to investigate the relationship between Amazon product reviews and the corresponding product ratings provided by customers. We have employed various conventional machine learning algorithms like the Naïve Bayes approach, Support Vector Machines, and the K-nearest neighbor method, along with advanced deep neural networks like Recurrent Neural Networks (RNNs) to achieve our objectives. Through a comparative analysis of the results obtained from these techniques, we have obtained valuable insights into their respective advantages and disadvantages. Additionally, our discoveries can also serve as a useful supplement to other fraud detection approaches.

Amazon Review Classification and Sentiment Analysis

On Amazon, reviews aren't just about the product – they also pertain to the service rendered to customers. When the clients are provided with a clear segregation of product reviews and service reviews, decision-making is rendered simpler. In this document, we put forth a scheme that facilitates customer review classification by seeking out the review sentiment. We also execute a regulation-grounded extraction of product feature sentiment. Additionally, we furnish a display for summating our outcomes.

Sentiment Analysis of Product Reviews – A Survey

With the enlarging expansion of the web network, the clients produce an abundance of feedback in the form of social media posts, online portals, appraisals, evaluations, and reviews on a wide range of subjects like books, individuals, products, research, events, etc. While intended to be helpful, most of this user-generated content necessitates the use of sentiment analysis techniques. Sentiment Analysis is a discipline that deals with the study of feelings, opinions, and subjectivity of sentiments. This paper is a comparable exploration of several recently proposed algorithms' advancements and various sentiment analysis applications, primarily for product reviews.

Sentiment analysis on large scale Amazon product reviews

In this modern era, the world is experiencing digitalization at an increasing rate. E-commerce has gained dominance in this digital world by bringing products to the customers' fingertips, eliminating the need to go out of their houses. As people increasingly rely on online shopping, the significance of product reviews has increased exponentially. A customer may need to sift through thousands of reviews to make an informed purchase decision. However, with the advent of machine learning, a model can be employed to analyze and classify these reviews, making the process much more efficient. In this study, we utilized a supervised learning approach on a large dataset of Amazon reviews to polarize them and obtain a satisfactory level of accuracy.

Sentiment Analysis on Amazon reviews

Opinion intel's vital for biz and makers. They want to know what folks think of their goods and services pronto. But sifting through every post on a site by hand is not feasible. There's just too much data. That's where sentiment analysis comes in, allowing for cost-effective, large- scale data processing. To better understand the potential and drawbacks of this technology in the context of business, the author delves into the topic in this paper. Using Python and machine learning techniques like Naïve Bayes and logistic regression, a model was built to predict the sentiment of Amazon Alexa reviews. To deal with the unbalanced dataset, SMOTE was utilized, and AUC/ROC was used to compare and evaluate the effectiveness of the methods.

3.System Analysis

Existing System

The existing system seamlessly integrates the General French Language Model with sophisticated analytical tools like logistic regression, naïve Bayes, and natural language processing (NLP) techniques. With a primary focus on sentiment analysis, it caters to French language data, aiming to decode the nuances of opinions and emotions embedded within textual information. By leveraging this blend of established models and cutting-edge NLP methodologies, the project strives to extract actionable insights from vast troves of text data. These insights could fuel various applications, ranging from deciphering customer feedback to conducting in-depth market research, ultimately empowering businesses to make informed decisions and enhance their strategies in the French-speaking domain.

Disadvantages

- It was developed only based on the reviews that are in French Language. It will not work in English Language.
- It has some components which are well advanced and requires high-end computers or machines to work efficiently.

Proposed System

The primary objective of this project is to delve into sentiment analysis within Amazon reviews through the lens of computer programs. Leveraging logistic regression and Naïve Bayes methods, the project seeks to unravel the sentiments underlying customer feedback, enabling businesses to gain valuable insights into consumer perceptions. In its initial stages, the project meticulously addresses data preprocessing tasks. This involves thorough checks for missing data and an in-depth assessment of data distribution. Additionally, we can do advanced visualization techniques such as heatmaps, histograms, and word clouds to gain a comprehensive understanding of the dataset. These visualizations offer valuable insights into data patterns, highlighting frequently used word. Common preprocessing steps include the removal of punctuation and stop words, which helps streamline the text and focus on the essential content. The project integrates a variety of techniques to facilitate accurate sentiment analysis. Through the combined use of logistic regression, Naïve Bayes methods, and comprehensive data preprocessing, the project aims to do sentiment patterns embedded within Amazon reviews. By deciphering the sentiments expressed by customers, companies can refine their products, services, and marketing strategies to better meet consumer needs and preferences.

Advantages

- By using powerful and effective machine learning algorithms like Naïve Bayes and Logistic Regression in this project made it more accurate than other existing models.
- This model prediction is better than existing ones.
- It has better performance than existing models

4. Material And Methods

Data were collected from Kaggle.com and some open sources. The data sources are corporate with the amazon alexa reviews which can include rating, date, variations, verified reviews and feedback. In this, we are verifying the positive and negative feedback by using word cloud. By using NLP (Natural Processing Language) is a branch of artificial intelligence that focuses on the interaction between computers and human languages. Its primary goal is to enable computers to understand, interpret, and generate human language in a manner that is both meaningful and contextually appropriate. NLP encompasses a wide range of tasks, including text analysis, speech recognition, machine translation, sentiment analysis, and more. By leveraging algorithms, statistical models, and linguistic rules, NLP enables computers to process and analyze vast amounts of textual data, unlocking valuable insights and facilitating communication between humans and machines.

INTERNATIONAL JOURNAL OF APPLIED SCIENCE ENGINEERING AND MANAGEMENT
ISSN 2454-9940
www.ijasem.org
Vol 19, Issue 1, 2025

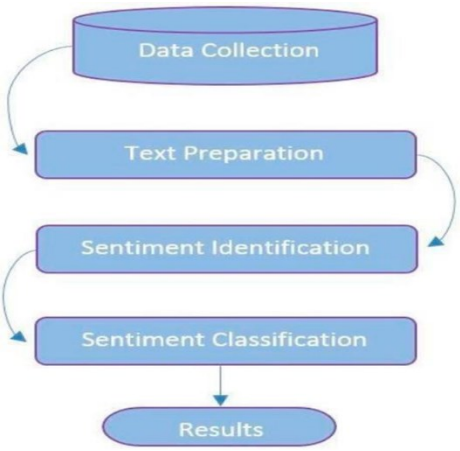

Figure 1: Flow Chart for Proposed Model

Datasets

The dataset used for this project is sourced from Kaggle, a well-known platform for sharing and exploring datasets. The dataset focuses on customer reviews of Amazon Alexa products, specifically the Amazon Echo and Echo Dot devices. This dataset is publicly available on Kaggle and is a valuable resource for sentiment analysis tasks. The dataset comprises customer reviews and ratings for Amazon Alexa products. Each review entry contains the following attributes:

- Rating: The rating given by the customer, ranging from 1 to 5.
- Date: The date when the review was posted.
- Variation: The specific model or variation of the Amazon Alexa product.
- Verified Reviews: The main text content of the review, where customers share their opinions and experiences regarding the product.
- Feedback: From this we can identify 0 and 1 that means positive and negative feedback. From that 0 and 1,we can say that positive has high feedback than negative feedback.
- Prior to analysis, the dataset underwent several preprocessing steps to ensure data quality and uniformity. These steps included:
- Handling missing values: Any missing values in the ‘Rating’ and ‘Verified Reviews’ columns were addressed.
- Text pre-processing: The textual content of the reviews underwent tokenization, removal of punctuation, conversion to lowercase, and other text processing techniques to facilitate further analysis.

The primary objective of using this dataset is to conduct sentiment analysis on customer reviews of Amazon Alexa products.

By leveraging machine learning algorithms, we aimed to categorize the reviews into positive and negative sentiments based on the customers’ experiences and feedback. The dataset includes customer reviews along with their corresponding ratings (ranging from 1 to 5). We preprocess the text data by removing punctuation, converting text to lowercase, and tokenizing the text using the Natural Language Toolkit (NLTK) library.

INTERNATIONAL JOURNAL OF APPLIED SCIENCE ENGINEERING AND MANAGEMENT
ISSN 2454-9940
www.ijasem.org
Vol 19, Issue 1, 2025

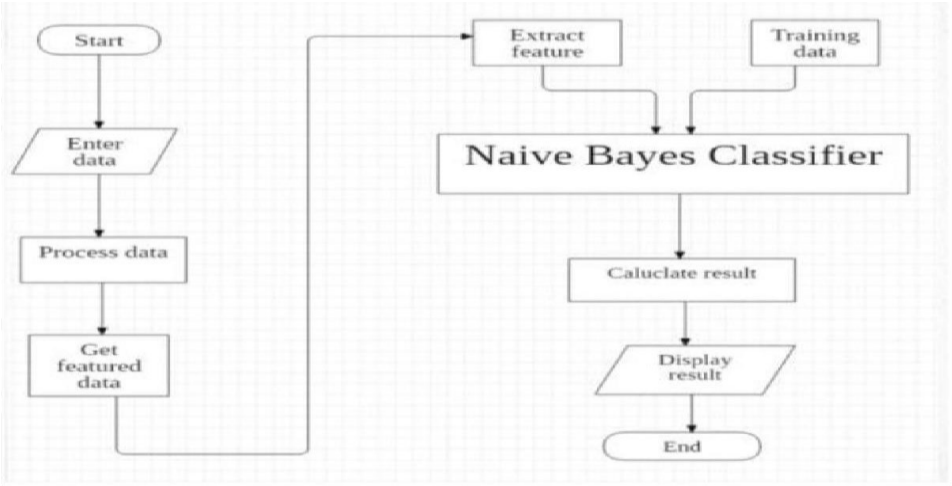

Figure 2: Naïve Bayes Classifier

Input Design

Designing an input system for Amazon Alexa reviews with sentiment analysis involves several key components to ensure accuracy and efficiency. The goal is to create a seamless user experience while accurately capturing and analyzing sentiments expressed in reviews. Firstly, the input system should allow users to submit their reviews through voice commands to Alexa. Users can simply speak their reviews, and Alexa will transcribe the spoken words into text. This text input will then be processed for sentiment analysis.

To ensure accuracy in sentiment analysis, the input system should employ natural language processing (NLP) techniques. This includes tokenization, part-of-speech tagging, and dependency parsing to break down the text into its constituent parts and understand the relationships between them. Sentiment analysis algorithms such as supervised machine learning models like Naïve Bayes Algorithm, Logistic Regression. From that we are getting positive and negative sentiment.

Output Design

Designing the output system for Amazon Alexa reviews with sentiment analysis involves presenting the sentiment analysis results to users in a clear and informative manner. The goal is to provide users with valuable insights into the sentiments expressed in their reviews, helping them understand the overall tone and reception of their feedback.

- The output system should begin by displaying the sentiment classification of the review, indicating whether it is positive or negative. This provides users with an immediate understanding of the overall sentiment conveyed in their review.
- To enhance user understanding, the output system can also provide a sentiment score or confidence level associated with the classification. This score indicates the degree of certainty with which the sentiment analysis model assigned the sentiment label to the review. Higher scores indicate greater confidence in the classification, while lower scores may indicate uncertainty or ambiguity.

5. Result and Discussions

The proposed model was tested and compared with two standard algorithms, including Naïve Bayes Algorithm, Logistic Regression. The test examined how accurate the tested algorithms predict for the

INTERNATIONAL JOURNAL OF APPLIED SCIENCE ENGINEERING AND MANAGEMENT
ISSN 2454-9940
www.ijasem.org
Vol 19, Issue 1, 2025

amazon alexa reviews. From that naïve bayes we are using Multinomial Naïve Bayes is a variant of the Naive Bayes algorithm specifically designed for text classification tasks where the features are word counts or term frequencies. It assumes that the features (word counts) are generated from a multinomial distribution. It works by calculating the probabilities of each class given the presence or absence of words in the document, taking into account the frequency of each word. Despite its simplicity, Multinomial Naive Bayes often performs well for text classification tasks, especially when dealing with sparse and high-dimensional feature spaces like text data.

After analysing the problem and applying visualization techniques and the above algorithms using certain parameters. we came to the following conclusions:

- Positive reviews are very common.
- Positive reviews are shorter.
- Longer reviews are more helpful.
- Despite being more common and shorter, positive reviews are found more helpful.

We build a confusion matrix and computed various results which include Fake Positive Reviews predictive value, Fake Negative Reviews, predictive value, Real Positive Reviews predictive value, Real Negative Reviews predictive value, accuracy and Precision.

The sentiment analysis output is visually represented through charts that provide insights into the model's performance:

Positive Values Chart:

This chart visually represents the distribution of correctly classified positive sentiment instances. This can would like to extract the reviews with rating 5 points and with 1 point, so that we can look at the best the worst reviews at the same time. it can be seen that the outstanding words like “love”, “echo”, “great”, “alexa”, “music”, “like”, “use”, “sound”, “new” “easy” and so on are in the rating 5. For the reviews of rating 1, we can see the words “amazon”, “device”, “work”, “disappointed”, “echo”. It says the people who gave rating 1 will describe the facts and the feelings together. And the world could be just a big picture of the words, we need to do more work to find more valuable things.

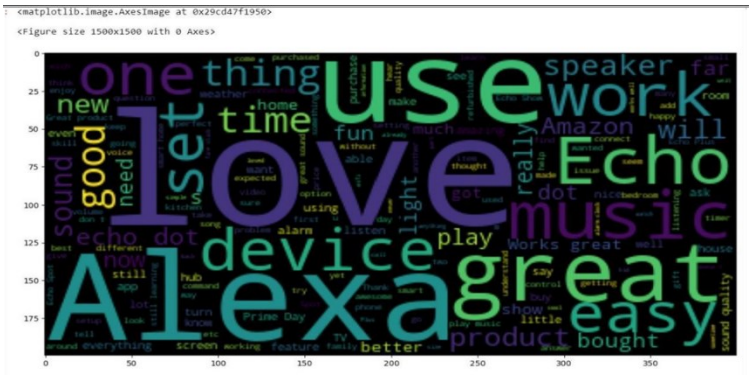

Figure 3: Popular Positive Word Cloud

INTERNATIONAL JOURNAL OF APPLIED SCIENCE ENGINEERING AND MANAGEMENT
ISSN 2454-9940
www.ijasem.org
Vol 19, Issue 1, 2025

Figure 4: Popular Negative Word Cloud

Discussions:

Sentiment analysis, a subfield of natural language processing, involves the classification of text into positive, negative, or neutral categories based on the underlying sentiment expressed. In the context of reviews, sentiment analysis plays a crucial role in extracting insights from large volumes of user-generated content. Positive reviews typically convey satisfaction, appreciation, or enthusiasm towards a product, service, or experience, highlighting its strengths, benefits, and positive attributes. On the other hand, negative reviews express dissatisfaction, disappointment, or criticism, pointing out flaws, issues, or areas for improvement. Analyzing both positive and negative reviews provides valuable feedback to businesses, allowing them to identify patterns, trends, and pain points among customers. By understanding the sentiments conveyed in reviews, companies can make informed decisions to enhance product quality, improve customer service, and address concerns effectively. Moreover, sentiment analysis enables businesses to monitor brand reputation, track customer satisfaction levels, and tailor marketing strategies to better meet consumer preferences and expectations.

6. Implementation

The dataset consists of 3150 rows and 5 columns, each representing different aspects of amazon alexa reviews for generating positive and negative sentiments. Analyzing 2018 amazon alexa reviews for generating the analysis

Rows: The dataset contains 3150 rows, each data set can have based on the reviews.
Columns
Column 1: Rating – This column have the rating from 1 to 5 based on the reviews given by the customers.
Column 2: Date - The date when the review was posted.
Column 3: Variation- The specific model or variation of the Amazon Alexa product.
Column 4: Verified Reviews- The main text content of the review, where customers share their opinions and experiences
Column 5: Feedback – The feedback can give based on the verified reviews as 0 and 1.

Data Preprocessing:

Data preprocessing plays a pivotal role in ensuring the accuracy and effectiveness of the sentiment analysis model. We performed a series of preprocessing steps on the textual data to prepare it for analysis. This involved removing any extraneous characters, symbols, or numbers that might distort the analysis. Additionally, the text was converted to lowercase to ensure uniformity, thus mitigating the case sensitivity of natural language. Tokenization, carried out using the Natural Language Toolkit

INTERNATIONAL JOURNAL OF APPLIED SCIENCE ENGINEERING AND MANAGEMENT
ISSN 2454-9940
www.ijasem.org
Vol 19, Issue 1, 2025

(NLTK), divided the reviews into individual words or subwords, aiding in the extraction of meaningful features. We used the following techniques to preprocess the data.

Text Cleaning:

- Remove Special Characters: Remove non-alphanumeric characters, punctuation marks, and symbols that do not contribute to the analysis.
- Remove stop words: Remove common stop words (e.g., "a," "an," "the", “is”, “are”) that do not carry significant meaning for the reviews.

Tokenization:

- Word Tokenization: Split text into individual words.
- Sentence Tokenization: Split into individual.

System Testing

The purpose of testing is to discover errors. Testing is the process of trying to discover every conceivable fault or weak or less in a work product. It provides a way to check the functionality of components, sub-assemblies, assemblies and/or a finished product. It is the process of exercising software with the intent of ensuring that the software system meets its requirements and user expectations and does not fail in an unacceptable manner. There are various types of test. Each test type addresses a specific testing requirement.

Figure 5: Plotting of Positive Reviews Using Word cloud

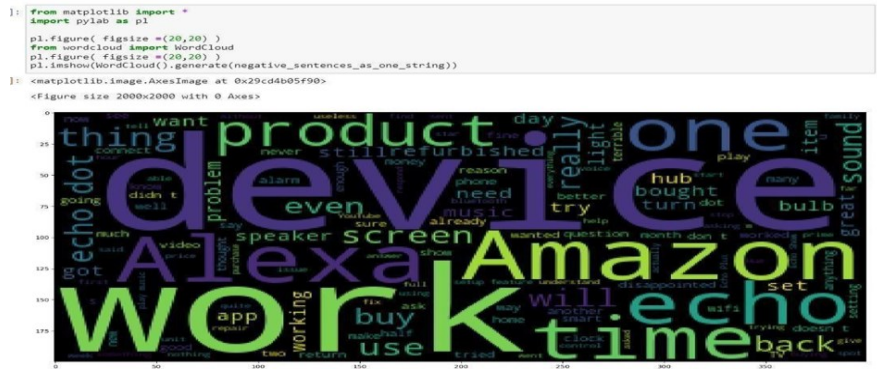

Figure 6: plotting of negative reviews using word cloud

INTERNATIONAL JOURNAL OF APPLIED SCIENCE ENGINEERING AND MANAGEMENT
ISSN 2454-9940
www.ijasem.org
Vol 19, Issue 1, 2025

```
[55]: from sklearn.metrics import classification_report,confusion_matrix

[57]: y_predict_test = NB_classifier.predict(X_test)
      cm = confusion_matrix(y_test,y_predict_test)
      sns.heatmap(cm,annot = True)

[57]: <Axes: >
```

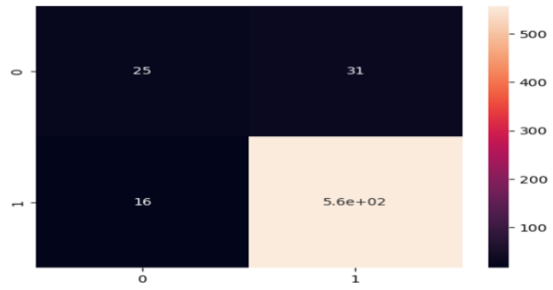

Figure 7: Naïve Bayes Confusion Matrix

```
]: print(classification_report(y_test,y_predict_test))
```

	precision	recall	f1-score	support
0	0.61	0.45	0.52	56
1	0.95	0.97	0.96	574
accuracy			0.93	630
macro avg	0.78	0.71	0.74	630
weighted avg	0.92	0.93	0.92	630

Figure 8: Accuracy Of Naïve Nayes Classification

```
]: cm = confusion_matrix(y_test,y_pred)
   sns.heatmap(cm,annot = True)

]: <Axes: >
```

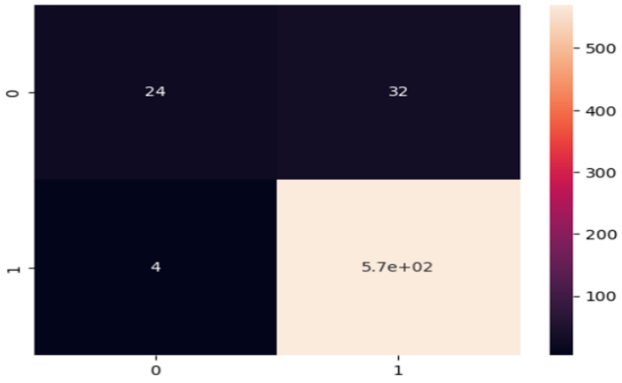

Figure 9: Logistic Regression Confusion Matrix

```
]: print(classification_report(y_test,y_predict_test))
```

	precision	recall	f1-score	support
0	0.61	0.45	0.52	56
1	0.95	0.97	0.96	574
accuracy			0.93	630
macro avg	0.78	0.71	0.74	630
weighted avg	0.92	0.93	0.92	630

Figure 10: Accuracy of Logistic Regression

Conclusion:

In delving into sentiment analysis within the realm of Amazon Alexa product reviews, our project has revealed the intricate layers of consumer sentiment and its implications for businesses and consumers alike. By harnessing the capabilities of machine learning, we've sifted through vast amounts of data to uncover the nuanced expressions of satisfaction, frustration, and everything in between. This journey has not only provided us with insights into how customers perceive Alexa products but has also underscored the pivotal role sentiment analysis plays in modern decision-making processes. Our findings serve as a testament to the value of sentiment analysis in decoding the voice of the customer. Businesses can leverage this understanding to refine their products, tailor marketing strategies, and enhance overall customer satisfaction. By discerning patterns within the sentiment data, companies can pinpoint areas of improvement, address pain points, and capitalize on strengths. Moreover, consumers stand to benefit from this analysis by gaining a deeper understanding of product features, performance, and reliability, empowering them to make informed purchasing decisions. As sentiment analysis continues to evolve, its significance as a strategic asset for businesses cannot be overstated. It enables companies to stay attuned to the ever-shifting landscape of consumer sentiment, adapt swiftly to changing preferences, and foster meaningful connections with their customer base. In essence, our exploration of sentiment analysis applied to Amazon Alexa product reviews underscores its transformative potential in shaping the future of customer-centric decision-making processes, paving the way for more informed choices and enriched experiences for both businesses and consumers alike.

Future scope:

Furthermore, the integration of sentiment analysis into Alexa's review system presents significant opportunities for businesses to gain actionable insights from customer feedback. By analyzing sentiment trends across a multitude of reviews, companies can identify recurring themes, sentiments, and patterns, enabling them to make data-driven decisions to enhance product development, marketing strategies, and customer service initiatives. For instance, businesses can quickly detect emerging issues or trends that may require immediate attention, enabling proactive problem-solving and mitigating potential reputational risks. Moreover, sentiment analysis can aid in benchmarking performance against competitors, allowing companies to gain a comparative understanding of their standing in the market and identify areas of competitive advantage or weakness.

References:

1. Kaggle. (2020). Amazon Alexa Reviews Dataset. Kaggle. https://www.kaggle.com/amazon-alexa- reviews/amazon alexa-reviews
2. Scikit-learn: Machine Learning in Python. (n.d.). Scikit-learn Documentation. Retrieved from https://scikitlearn.org/stable/documentation.html
3. NLTK 3.6 documentation. (n.d.). NLTK Documentation. Retrievedfrom https://www.nltk.org/nltk_data/
4. Pandas Documentation. (n.d.). pandas: powerful data analysis tools for Python. Retrieved from https://pandas.pydata.org/docs/
5. "Amazon Alexa Reviews" Dataset.(n.d.). Kaggle. Retrieved from https://www.kaggle.com/sid321axn/amazon-alexareviews
6. Bird, S., Klein, E., & Loper, E. (2009). Natural Language Processing with Python. O'Reilly Media. Zhang, X., Zhao, J., & LeCun, Y. (2018). Character-level Convolutional Networks for Text Classification. In Advances in Neural Information Processing Systems (NIPS) (pp. 6493-6502).

7. Kim, Y. (2014). Convolutional Neural Networks for Sentence Classification. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP) (pp. 1746-1751).
8. Pang, B., & Lee, L. (2008). Opinion Mining and Sentiment Analysis. Foundations and Trends in Information Retrieval, 2(1-2), 1-135.
9. Turney, P. D. (2002). Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. In Proceedings of the 40th Annual Meeting on Association for Computational Linguistics (ACL) (pp. 417-424).
10. Boothroyd, A., 2018. The benefits (and limitations) of online sentiment analysis tools [WWW Document]. Typely Blog. URL https://typely.com/blog/making-use-of-sentiment- analysis/ (accessed 5.17.19).
11. Brownlee, J., 2017. A Gentle Introduction to the Bag-of-Words Model [WWW Document]. URL https://machinelearningmastery.com/gentle-introduction-bag-words-model/ (accessed 5.15.19).