IJNRD.ORG ISSN : 2456-4184
INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD) | IJNRD.ORG
An International Open Access, Peer-reviewed, Refereed Journal

MONITORING OF ONLINE FAKE REVIEWS THROUGH MACHINE LEARNING

[1]Vijendra S. N, [2]Vaishnavi Patil, [3]Amrutha B M, [4]Jyothi R , [5]Vandana V,

[1]Assistant Professor, [2]Student , [3]Student, [4]Student, [5]Student
[1]Computer Science and Engineering
Impact College of Engineering and Applied Sciences, Bangalore , India

Abstract : The escalating trend of online shopping has witnessed a significant surge, with a growing number of individuals preferring to procure products from virtual stores. Consequently, the utilization of the Internet and online marketing has attained widespread popularity. The online marketplace now offers an extensive array of millions of products and services, thereby generating a vast volume of information. Consequently, locating the most suitable services or products that align with specific requirements has become a daunting task. In such a vast landscape, customers heavily rely on reviews and opinions shared by others, which are based on their personal experiences, to make informed decisions. The spread of bogus reviews has emerged as a critical concern in today's world of intense competition. Companies are increasingly turning to employing people to create positive reviews for their own goods or services while writing unfairly critical comments about their rivals, which is a worrying trend. These dishonest actions have serious repercussions because they deceive prospective buyers who rely on these ratings to make wise purchasing decisions. A strong system that is capable of identifying and removing bogus reviews is thus urgently needed. In order to identify false reviews, this study will examine a variety of supervised, unsupervised, and semi-supervised machine learning algorithms. Multiple classifiers, including Support Vector Machine (SVM), Nave Bayes, and Random Forest Classifier, are used to assess each approach. Additionally, the study takes into consideration the effectiveness of n-gram techniques and linguistic models to enhance the evaluation process. The proposed methodology demonstrates a noteworthy accuracy rate of 87% in effectively identifying fake reviews.

***Index Terms* - Fake reviews detection, Machine learning Algorithms, Text Classification, Natural Language Processing, Bigrams, Term Frequency and Inverse Document Frequency.**

I. INTRODUCTION

In the current landscape of the trade market, product reviews have become indispensable in shaping consumers' online shopping behavior. Prior to making a purchase, a significant number of individuals rely on these reviews to gather insights and make informed decisions. These reviews, whether positive or negative, wield substantial influence over consumers' purchasing choices. Positive reviews tend to captivate customers' attention more effectively than negative ones, ultimately impacting the success or failure of businesses. Consequently, product reviews possess the potential to significantly impact financial outcomes for companies. However, alongside authentic reviews, a proliferation of counterfeit reviews further complicates the purchasing process, posing challenges for consumers seeking genuine feedback on products.

It could end up being a waste for you if you aren't careful when purchasing that item. However, modern consumers now have a different source of buying. From the internet shops of many brands, you can purchase goods. Without viewing or inspecting the original product, you must place the order here. When purchasing a product, you read the reviews. The product reviews are therefore your only source of information. It's possible that these reviews are fraudulent or authentic. The buyer wants to purchase a trustworthy, authentic goods, which is only achievable if you have access to the actual user reviews for that item. 6 billion dollars were spent on Black Friday sales in the United States in 2018, according to research.

In the present era, an impressive number of e-commerce stores, ranging from 12 to 24 million, are actively operating on a global scale. Recent research has shed light on the prevalence of fake reviews, particularly within the Electronics Category of the prominent platform, Amazon. Astonishingly, a study revealed that approximately 61% of the reviews within this category are suspected to be fake. Acknowledging the significance of combating this issue, several websites have emerged with the purpose of detecting and identifying fraudulent reviews. One notable example is Fake Spot, an online platform that employs sophisticated algorithms to analyze suspicious patterns and reviewer behavior, thereby assisting consumers in differentiating genuine feedback from fabricated ones. By leveraging these tools and resources, consumers can make more informed decisions when navigating the complex landscape of online product reviews.

E-commerce is a sector that is expanding very quickly. In general, e-commerce gives customers the option to post feedback about its services. The method that uses a review's language and rating information to identify bogus product reviews. Data mining techniques will be used to determine the honesty value and measure of a false review. An algorithm might be used to monitor customer reviews by extracting subjects and sentiment from online reviews, and it would also filter out fraudulent reviews.

II. LITERATURE SURVEY

The primary objective of conducting a literature review is to acquire comprehensive knowledge and insights into the existing research and ongoing discussions pertaining to a specific topic or area of study. This process involves thoroughly examining scholarly works and scholarly debates, subsequently presenting the gathered information in a well-structured written report. By engaging in a literature review, individuals can enhance their understanding and expertise within their respective fields of study, contributing to the expansion of knowledge in their chosen disciplines.

Fake reviews, fake news, and fraudulent social media IDs have all been the subject of much research. In order to promote or denigrate the seller's product, fake reviewers cooperate in groups to write evaluations on e-commerce websites. This practice is known as "spotting groups of fake reviewers" [7]. Developers search the groups of fictitious reviewers using the "Frequent Itemset Mining (FIM)" technique. The system employs relational models and behavioral models to determine the relationships between the "spammer groups"—also known as fake reviewer groups—that are being used. Since there was no labelled dataset available prior to this work, they used professional human judges to create one in order to test their strategy. This system employs a cutting-edge relation-based model called "GSRank" to identify fake reviewers and the relationships between spammer groups. In this method, each transaction serves as the reviewer's set of IDs, and the set of items is the reviewer's collection of Ids. The system then uses the FIM method to identify the groups of reviewers who have given reviews of a variety of products collectively.

The detection of fake reviews on platforms such as Yelp has emerged as an increasingly critical necessity. In response, a system has been developed to effectively distinguish between fake and genuine reviews. This system employs a classifier that carefully analyzes various components, including review text and additional contextual information, to assess the reliability of the reviews. To conduct this analysis, a dataset obtained from Yelp.com, originally utilized by Rayana and Akoglu, was employed. The dataset encompasses 16,282 reviews, which were divided into distinct subsets for training (70%), development (20%), and testing (10%) purposes. Extracting informative and predictive features from the reviews posed notable challenges during the project. Two primary types of features were extracted: review-centric features and reviewer-centric features. Initially, the percentages of each unigram and bigram token were computed for both fake and non-fake reviews. The top 100 unigrams and bigrams exhibiting the most significant differences in occurrence between fake and non-fake reviews were selected. However, the subsequent approach, which involved considering all unigrams and bigrams, demonstrated superior performance. Multiple machine learning algorithms were evaluated, with the Neural Networks algorithm achieving the highest accuracy of approximately 71.92%. Although the system showcases effectiveness in identifying fake Yelp reviews, there is still room for further improvement to enhance the accuracy of the review filtering process [11].

By employing Sentiment analysis, the fake review monitoring system [13] concentrates on identifying spam and fraudulent reviews and eliminates any that contain profanity or other offensive language. The data on the website is scraped in the suggested system using a web crawler. The data is preprocessed into the necessary format, and then the mixture of authentic and spam reviews is cleaned up of the bogus reviews. The Fake Review Detector can recognize fake reviews. The classifier that determines a review's sentiment score must approve each review before it may be published. The metric for this kind of similarity is cosine similarity. The review is regarded as a fraudulent review if the cosine value that must be calculated is larger than 0.5. Out of 300 reviews, the developed algorithm discovered that 111 were false. The model cannot accurately detect the suspicious patterns due to the tiny size of the data set used to train it.

The objective of the study was to develop an effective method for detecting deceptive and genuine reviews. To achieve this, both supervised and unsupervised techniques were employed, leveraging labeled and unlabeled data. A comprehensive set of features was carefully selected to build the detection model, and sentiment analysis was integrated into the process. For the labeled dataset, well-established classifiers were applied to maximize performance. Additionally, for the unlabeled data, clustering was utilized after extracting relevant attributes to identify spam. Notably, there is a significant likelihood that spam reviewers contribute to content pollution in multimedia social networks, as many users now provide reviews using their social network logins. The incorporation of sentiment analysis in the detection process was crucial for achieving optimal results. Furthermore, the study suggests the possibility of extending the research to uncover suspicious accounts responsible for sharing fabricated multimedia content across various social networks.

Previous research has used machine learning techniques, opinion mining, sentiment analysis, and IP addresses to identify bogus reviews. Some methods employ a relatively small dataset, while others use a few properties connected to the reviews to identify false reviews. The vast dataset of English reviews is used to train the model in the proposed system. By doing this, the machine can more effectively uncover any underlying trends in the evaluations. The accuracy obtained by employing the suggested method is superior to the accuracy of the preceding systems in terms of the English reviews.

III. PROPOSED APPROACH

In the contemporary digital landscape, the ability to identify fraudulent product reviews holds paramount importance. Within the realm of e-commerce platforms, two distinct categories of purchases exist: Verified and Non-verified. A verified purchase signifies that the individual responsible for the review has indeed acquired the item from the online retailer. Consequently, when a customer provides a positive rating and assigns a score of 1 or 2 to the product in question, it becomes apparent that an artificial review has been submitted. To discern the authenticity of a verified purchase, an algorithm is employed, utilizing sentiment analysis. This analytical approach distinguishes sentiment polarity based on the presence of words such as 'good' and 'excellent,' indicating a positive sentiment, or 'bad' and 'poor quality,' indicating a negative sentiment. By leveraging sentiment analysis, the algorithm effectively identifies potentially counterfeit reviews, enhancing the credibility and reliability of the reviewing process.

In the case of non-verified purchases, individuals can submit product reviews on e-commerce platforms without any prior purchase history. This opens up the possibility for both genuine and deceptive reviews to be posted. While it is plausible for someone providing a positive review to also assign a good rating to the product, there exists the risk that such individuals are merely attempting to artificially boost the product's overall rating, indicating spam-like behavior.

In such scenarios, relying solely on sentiment analysis proves insufficient as it may mistakenly classify these reviews as genuine due to the positive sentiment expressed. To address this limitation, the proposed system incorporates an alternative technique utilizing Support Vector Machine (SVM) as a classifier. SVM is a robust machine learning algorithm capable of distinguishing patterns and making accurate classifications. By leveraging SVM as a classifier, the system enhances its ability to identify potentially deceptive reviews and distinguish them from authentic ones. This approach complements sentiment analysis by providing an additional layer of analysis, thereby improving the overall effectiveness and reliability of the review evaluation process.

Details of the suggested strategy are explained in this section. The suggested method consists of three basic steps that can be used to distinguish between genuine and fake reviews. The following is a definition of these phases:

3.1 Collection of Dataset

The first phase of the suggested system is the gathering of datasets. A data set is a collection of connected data. In the case of tabular data, a data set is associated with one or more database tables, where each row alludes to a particular record in the associated data set and each column to a particular variable. The system collected the dataset of English reviews from the kaggle website, which includes both genuine and fake product reviews, and used it to train the model. The Amazon review datasets are the one used in this study. With 12963 rows and 9 columns, the data is organized. The proposed system has an accuracy rate of 87% when applied to the given dataset. Both bogus and real reviews are identified in these datasets. 10501 genuine reviews are on label 1, whereas 2462 fake reviews are on label 2.

3.2 Data Preprocessing

Data preparation is the following stage in the suggested system. This is one of the crucial elements in computerized learning processes. As the world's data is utterly unsuitable for use, data preparation is a crucial task. The raw data of the Kaggle dataset for computational activities have been assembled in this work via a series of preprocessing techniques. The following succinct statement sums it up:

A. Remove Punctuation

Eliminating all of the text's punctuation is the first step in pre-processing. Punctuation, which includes symbols like the full stop, comma, semicolon, and hyphen, is used to demarcate the boundaries of sentences in order to make their intended meanings clear. Each review had these points taken away.

B. Stop Words Removal

Stop words are commonly used words in sentences that serve as connectors. However, in the context of feature extraction for text classification, stop words tend to introduce unnecessary noise. These noise-inducing words, including articles, prepositions, and certain pronouns, should be eliminated before conducting text classification tasks. Consequently, the system undertakes the removal of such stop words, such as 'is,' 'are,' 'am,' 'he,' 'from,' 'who,' and others, from each review. This preprocessing step ensures that the review is streamlined and optimized for subsequent stages of analysis and classification. By eliminating these stop words, the system aims to improve the efficiency and accuracy of the text classification process.

C. Lemmatize Words

An important text classification technique for English text is lemmatization. Its main goal is to reduce words to their most basic forms so that they can be combined with other inflected forms to produce a single word. Finding the lemma, or the base form, of a given word, is a necessary step in the lemmatization process. For instance, the verb "to walk" might be represented by the word "walk," which is its lemma. The system's approach tries to increase the classifier's speed and effectiveness by using lemmatization.
By converting words to their base forms, the system streamlines subsequent preprocessing steps. This approach facilitates improved text classification by providing a more standardized and coherent representation of words. By leveraging lemmatization, the system contributes to faster and more efficient processing, enhancing the overall performance of the text classification process.

D. N-gram Modeling

The approach to feature identification that is most widely used in natural language processing is n-gram modelling. N may take on any value. The word-based or character-based n-gram modelling is crucial for text classification. The system uses this model to generate characteristics for categorizing the reviews. Bigrams (N = 2) is being used in the suggested methodology to determine whether the reviews are accurate or dishonest. Following the conversion of each review into a bigram, the system completes the preprocessing phase and is now prepared to move on to the feature extraction technique. N-gram algorithms or programs can be used to find all continuous neighboring word sequences inside a given set of phrase tokens. It is a feature of Windows that moves windows one position to the left starting from the leftmost word position.

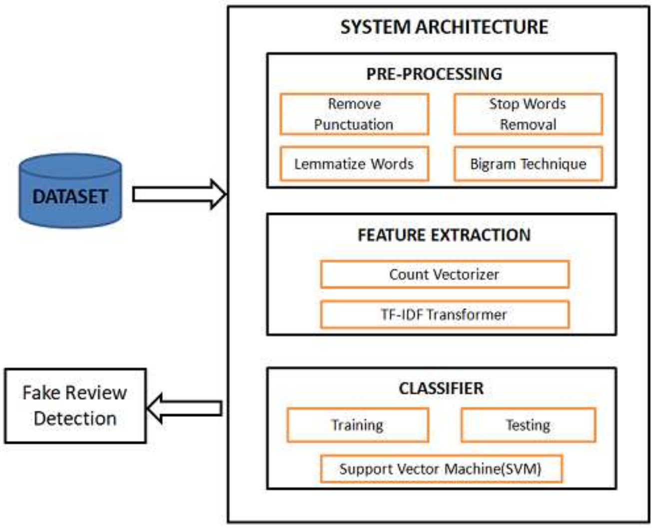

Figure. 1 System Architecture

3.3 Feature Extraction

Among the various steps involved in text classification, feature extraction stands out as a particularly challenging and critical process that significantly contributes to the accuracy of the classifier. Irrelevant data can have a detrimental impact on the performance and precision of the classifier. One well-established technique for feature extraction is textual point analysis, which includes sentiment classification based on the occurrence of positive and negative words (e.g., "good" and "weak") within the review. Additionally, two feature selection techniques, namely term frequency (TF) and term frequency-inverse document frequency (TF-IDF), are employed to identify relevant characteristics within our dataset.
In the proposed approach, the system utilizes Count Vectorizer to convert each review into a bag-of-words representation. This process tokenizes the words present in the reviews, allowing for further analysis. Subsequently, the TF-IDF transformer is applied. This transformation converts the collection of text documents into a matrix representing token counts. The resulting matrix is two-dimensional, with one dimension representing the vocabulary and the other dimension capturing the actual document, as depicted in Table 1.

Table 1: Count Vectorizer Overview

	Word 1	Word 2	…………..	Word N
Review 1	0	2	…………..	1
Review 1	0	1	…………..	1
……………	1	0	…………..	2
Review N	2	1	…………..	0

This matrix has a lot of zero values, which is why it is referred to as a sparse matrix. Depending on how frequently a term appears in the dataset, its significance rises. The TF-IDF score for each word is given. The weight W (d, t) of a word t in a given document d is given as described in

Eq. 1: W(d, t) = TF(t, d) log()..........(1)

N is the total number of documents (reviews) in the dataset, and TF (t, d) represents the frequency of the term t in each document. The term "t" appears in how many documents (reviews) there are. In this research, we use TF-IDF to observe the facts of the contents in two languages models, which are typically bi-gram and tri-gram. After collecting the points that correspond to the customer's behaviors in each language model, we also follow the extended dataset.

3.4 Classification Process

Figure. 2 illustrates the classification process employed by the proposed system to distinguish genuine reviews from fake ones. The process begins with data collection, followed by data pre-processing. During pre-processing, various steps are undertaken, such as removing punctuation and stop words from the text of the reviews. Additionally, the system converts the text to lowercase, lemmatizes words to their base form, and utilizes the bigram technique to represent the text as pairs of consecutive words.
Once the pre-processing stage is completed, feature extraction takes place using Count Vectorizer, which converts each review into a two-dimensional matrix. Subsequently, the TF-IDF transformer is applied to assign weights to each word based on their relevance in the corpus. This step enhances the discriminative power of the features.
After feature selection, the final step in the classification process involves training the classifier. The proposed architecture is evaluated using three different supervised machine learning algorithms: Support Vector Machines (SVM), Naïve Bayes, and Random Forest Classifier. Among these algorithms, SVM demonstrates superior performance, surpassing the others in terms of results and accuracy. Through this systematic approach, the proposed system effectively categorizes reviews into genuine and fake, leveraging the strengths of SVM to achieve notable classification accuracy.

3.5 Predictions and Evaluation

The remaining data in this experiment, which amounts to 70%, is used for testing. The algorithm is then used with testing data to anticipate unseen data in order to determine if it is fake or genuine once the data has successfully been trained. Bigrams feature obtained 87% accuracy in the Amazon reviews dataset using the proposed method. Reviews are categorized into genuine and fake ones in the following Figure. 2.

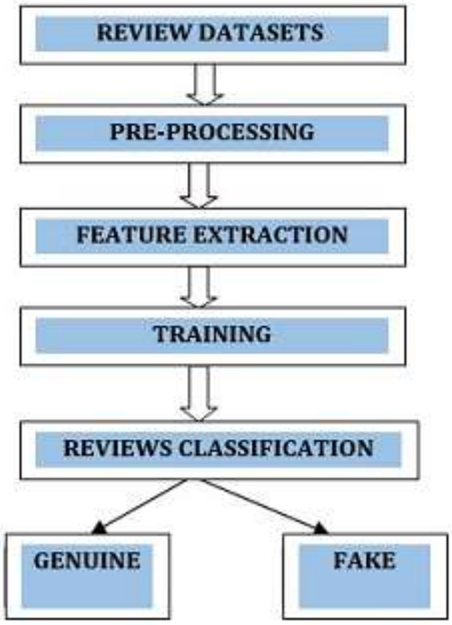

Figure 2 : Classification Process

IV. ALGORITHMS USED

4.1 SENTIMENT ANALYSIS

Understanding the writer's emotions using sentiment analysis or opinion mining is a straightforward undertaking. To identify information that might be subjective, we employ a variety of natural language processing (NLP) and text analysis technologies. To make data classification and use simpler, we must recognise, extract, and quantify such features from the text.

How is sentiment analysis carried out?

Sentiment analysis is based on the fundamental idea of NLP, or natural language processing. Natural language processing, a subset of sentiment analysis, focuses on extracting information of various kinds from a text. Based on numerous criteria, sentiment analysis can be divided into a number of areas.

The use of the data or the purpose of sentiment analysis is another factor that is frequently used to classify data.

1. Straightforward division of text into positive, negative, and neutral categories. Additionally, it might progress to more nuanced responses like "very positive" or "moderately positive".
2. We identify the sentiment and the specific element to which it is related using an approach called aspect-based sentiment analysis. Like determining what features or characteristics were liked or disliked when reading user feedback about various aspects or components of a car.
3. The emotion and an action that go along with it. similar to emails sent to customer service. Recognizing whether it is a question, complaint, proposal, etc.

4.2 SVM-SUPPORT VECTOR MACHINE

One of the most well-liked supervised learning algorithms, Support Vector Machine, or SVM, is used to solve Classification and Regression issues. However, it is mainly employed in Machine Learning Classification issues. The SVM method seeks to define the best line or decision boundary that can divide n-dimensional space into classes in order to quickly categorise fresh data points in the future. This optimal decision boundary is known as a hyperplane. Support vectors are the phrase for these extreme examples, and the SVM technique is named after them. The SVM algorithm is useful for a variety of tasks, including text classification, image classification, and face detection. The SVM algorithm assists in determining the ideal decision boundary, also known as a hyperplane. The SVM algorithm determines which line from each class is closest to the centre. Support vectors are the names given to these points. The margin is the separation between the vectors and the hyperplane. SVM aims to increase this margin. The ideal hyperplane is the one with the greatest margin.

4.3 RANDOM FOREST

Random Forest is a part of the supervised learning methodology. It can be applied to ML issues involving both classification and regression.

It is built on the idea of ensemble learning, which is a method of integrating various classifiers to address difficult issues and enhance model performance. A classifier called Random Forest uses a number of decision trees on various subsets of the data to boost the projected accuracy for the dataset and takes the average to improve the predictive accuracy of that dataset.

4.4 NAÏVE BAYES

The Naïve bayes algorithm is a supervised learning technique used to solve classification issues. With a high-dimensional training dataset, it is mostly utilized in text categorization. One of the most straightforward and efficient classifiers is the naive bayes algorithm, which aids in creating quick machine learning models that can generate predictions quickly. It makes predictions based on the likelihood of an object because it is a probabilistic classifier. Spam filtration, Sentimental analysis, and categorizing articles are a few common uses of the Naive Bayes algorithm.

V. Results and Visualization

Kaggle Reviews : Kaggle reviews dataset consists of 12,693 reviews, 80% data is genuine and remaining 20% is spam. The below given figure 3. Shows the comparison graph between the genuine and spam reviews in the dataset.

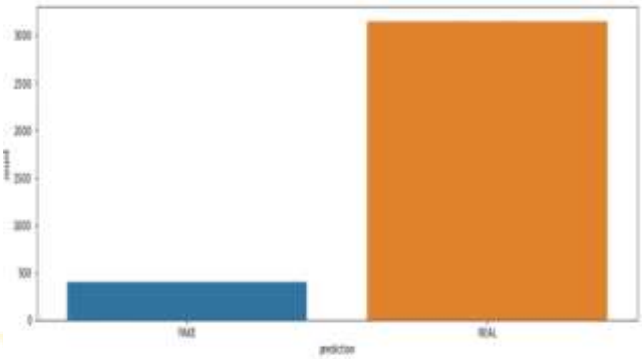

Figure 3: Bar graph showing genuine and spam reviews in the dataset.

Kaggle Dataset : The proposed architecture uses SVM classifier to train the model that consists of 12,963 rows. Confusion matrix of the dataset is given in below Figure 4.

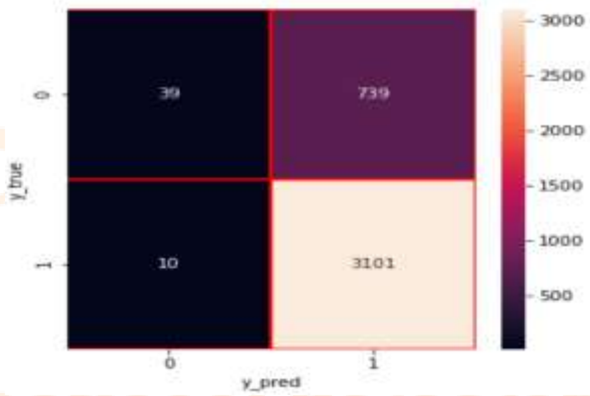

Figure 4 : Confusion Matrix

The confusion matrix is a matrix used to assess how well the classification model is working [16]. The confusion matrix is summarized using count values and classified by class, with the key being the ratio of accurate and incorrect guesses. Despite being easy to understand, the employed parameter is confusing. The situation is as follows:

True Positive (TP): When the model classified the actual value and it is True.

True Positive (TP): When the model correctly identified the real value.

True Negative (TN): When the model identified a negative actual value.

False Positive (FP): When the classifier assumes that the news is accurate when it is not.

False Negative (FN): When the news is genuinely true but the classifier predicts it to be false.

Precision Metrics: Precision metrics demonstrate the proportion of accurately predicted cases that were positive. The model's dependability is assessed using these metrics [16].

$$\text{Precision} = \frac{\text{True Positive (TP)}}{\text{(True Positive (TP) + False Positive(FP))}}$$

Recall Metrics : The number of truly positive cases that the model was able to accurately anticipate is shown by recall metrics [16].

$$\text{Recall} = \frac{\text{True Positive (TP)}}{\text{(True Positive (TP) + False Negative (FN))}}$$

F1 Score: F1 provides an overview of both the precision and recall measurements. This implies that when we attempt to upgrade, the value of precision decreases, and vice versa [16].

$$\text{F1 score} = 2 \times \frac{\text{Precision x Recall}}{\text{(Precision+ Recall)}}$$

Accuracy: Accuracy is measured by the percentage of accurate predictions and the total number of predictions made by the classifiers [16].

$$\text{Accuracy} = \frac{|TP| + |NP|}{|TP| + |TN| + |FP| + |FN|}$$

Table 2: Accuracy of the model

	Precision	Recall	F1-Score	Support
0	0.80	0.05	0.09	778
1	0.81	1.00	0.89	3111
accuracy			0.81	3889
Macro average	0.80	0.52	0.49	3889
Weighted average	0.81	0.81	0.73	3889

The Precision, Recall, F1-score and Accuracy of the whole model is shown in the above table-2.

VI. CONCLUSION

Our software aims to assist users in making informed purchasing decisions by detecting fake product reviews. To address the challenges of manually identifying fake reviews, we propose a machine learning approach. We have developed a comprehensive dataset comprising product reviews to train our system.

The proposed system possesses the ability to discern the nuances of sentiment expressed within reviews and ratings. By analyzing the aspects of product reviews as they are posted, our system can accurately determine the authenticity of each review. This capability allows users to rely on genuine and trustworthy feedback when considering a purchase.

REFERENCES

[1] A. Sinha, N. Arora, S. Singh, M. Cheema, and A. Nazir, "Fake Product Review Monitoring Using Opinion Mining," *vol. 119, no. 12, pp. 13203– 13209, 2018.*

[2] Torbet, Georgina. "U.S. Customers Spent over $6 Billion on Black Friday Purchases." Digital Trends, Digital Trends, www.digitaltrends.com/web/shopping-totals-blackfriday/.

[3] Sterling, Greg. "Study Finds 61 Percent of Electronics Reviews on Amazon Are 'Fake'." *Marketing Land, 19* Dec. *2018,* marketingland*.com/*study-*finds-61-percent-of* electronics-reviews-*on-amazon-are-fake 254055.*

[4] V. Sujitha, V. RajaKumari, P. Hema, M. Tulasi, M. SivaKishore "Fake review Monitoring System." *Vol-6 Issue-3 2020 IJARIIE-ISSN(O)-2395-4396.*

[5] Banerjee, S., & Chua, A. (2014). Applauses in hotel reviews: Genuine or deceptive? In Science and information conference (SAI) (pp. 938–942). IEEE.

[6] Möhring M., Keller B., Schmidt R., Gutmann M., Dacko S. HOTFRED: "A Flexible Hotel Fake Review Detection System". In: Worndl W., Koo C., Steinmetz J.L. (eds) Information and Communication Technologies in Tourism. Springer, Cham. *https://doi.org/10.1007/978-3-030- 65785-7_29.*

[7] Hernandez-Fusilier, D., Guzman-Cabrera, R., Montes-y-Gomez, M., Rosso, P.: Using PU- learning to Detect Deceptive Opinion Spam.

[8] A. Mukherjee, B. Liu, and N. Glance, "Spotting Fake Reviewer Groups in Consumer Reviews," 2012.

[9] Z. Wang, Y. Zhang, and T. Qian, "Fake Review Detection on Yelp Dataset and features," pp. 1–6.

[10] S. Xie, G. Wang, S. Lin, and P. S. Yu, "Review spam detection via temporal pattern discovery," p. 823, 2012.

[11] C. Paper, "Mining millions of reviews : A technique to rank products based on importance of reviews Mining Millions of Reviews : A Technique to Rank Products Based on Importance of Reviews," no. November, 2015.

[12] Rayana and L. Akoglu, "Collective Opinion Spam Detection : Bridging Review Networks and Metadata," Sigkdd, pp. 985–994, 2015.

[13] A. Mukherjee, V. Venkataraman, B. Liu, and N. Glance, "What Yelp Fake Review Filter Might Be Doing?," *Aaai, pp. 409–418, 2013.*

[14] V. K. Madhura N Hegde, Sanjeetha K Shetty, Sheikh Mohammed Anas, "Fake product review monitoring," Int. Res. J. Eng. Technol., *vol. 05, no. 06, p. 4, 2018.*

[15] J. Rout, S. Singh, S. Jena, and S. Bakshi , "Deceptive Review Detection Using Labeled and Unlabeled Data" 2016.

[16] Guide to Confusion Matrix Terminology.

https://www.dataschool.io/simple-guide-to-confusion- matrix-terminology.

[17] T. R. Patil and S. S. Sherekar, "Performance analysis of naive bayes and j48 classification algorithm for data classification," pp. 256–261, 2013.

[18] H. Li et al., "Spotting fake reviews via collective positive-unlabeled learning." 2014.

[19] Brownlee, J. (2020, January 3) How to Calculate Precision, Recall, and F-Measure for Imbalanced Classification Journal of Computer and Communications d-classification.

[20] Ott, Myle & Choi, Yejin & Cardie, Claire & Hancock, Jeffrey. "Finding Deceptive Opinion Spam by Any Stretch of the Imagination", 2011.