International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Unveiling Deception: A Fusion of Deep Learning and Sentiment Analysis for Identifying Counterfeit Reviews

Jitender Kumar[1], Baldev Singh[2]

Submitted: 29/12/2023 **Revised**: 05/02/2024 **Accepted**: 13/02/2024

Abstract: In the contemporary landscape of consumer decision-making, the influence of online reviews is paramount. However, the authenticity of these reviews has become a pressing concern. This study proposes a comprehensive strategy for identifying counterfeit reviews on online platforms by integrating advanced deep learning techniques with sentiment analysis. The primary objective is to develop a model capable of distinguishing between deceptive and genuine reviews. The methodology includes data acquisition, preprocessing, and the application of a neural network model featuring key elements such as an Embedding layer for word representations, a Convolutional layer for feature extraction, a Long Short-Term Memory (LSTM) layer for capturing sequential dependencies, and a Dense output layer for binary classification. To evaluate the model's effectiveness, a dataset comprising categorized reviews is utilized. The dataset is split into training and testing subsets, and the model undergoes training across multiple epochs, with continuous monitoring of metrics like loss and accuracy. Visual representations illustrate the model's training progress. Additionally, the study incorporates sentiment analysis using the VADER tool to assess the emotional tone of reviews, aiding in the differentiation between authentic and fabricated sentiments. The research findings highlight the efficacy of the combined deep learning and sentiment analysis approach in detecting counterfeit reviews. The model exhibits competitive performance in review classification, potentially enhancing trustworthiness on online platforms. The sentiment analysis component enriches our understanding of user sentiments, providing a deeper insight into review content. By offering a robust and interpretable model alongside a comprehensive methodology, this research significantly contributes to the field of counterfeit review detection in the digital era.

Keywords: *VADER, incorporates, comprehensive, counterfeit, Long Short-Term Memory*

1. Introduction

In the contemporary era dominated by the digital ecosystem, social media platforms have become crucibles of influence, shaping opinions and steering consumer choices. Among the myriad content disseminated across these platforms, user-generated reviews stand as potent indicators, guiding individuals in their decisions about products, services, and experiences. Yet, within this seemingly transparent landscape, a pervasive challenge looms large — the infiltration of fake reviews [1].

The gravity of the problem lies in the profound impact that these counterfeit reviews wield. As consumers increasingly rely on the wisdom of the crowd, deceptive narratives injected into the review ecosystem can significantly distort perceptions and influence purchasing decisions [1]. The deceptive practices manifest in a spectrum of cunning strategies, from artificially boosting product ratings to creating fabricated narratives that misrepresent real user experiences [2].

This research delves into the heart of this issue, recognizing the substantial threat that fake reviews pose to the credibility and reliability of information disseminated on social media [2]. The prevalence of such deceptive practices not only erodes trust in online platforms but also compromises the integrity of the digital marketplace as a whole. Our study addresses the gravity of the problem by proposing a comprehensive strategy that amalgamates advanced deep learning techniques with sentiment analysis [3]. The objective is twofold: to develop a model capable of effectively discerning between authentic and counterfeit reviews and to shed light on the emotional undercurrents within user-generated content. As we navigate through our methodology and present our findings, it becomes apparent that our work is not merely about identifying deceptive reviews but is also a critical step toward restoring trust in the digital narrative [3].

As users increasingly turn to social media for guidance, the need for robust mechanisms to authenticate the veracity of reviews becomes imperative. Our research contributes to the broader discourse on social media integrity by offering a nuanced understanding of the gravity of the fake review problem and presenting a viable solution that holds the potential to safeguard the authenticity of user-generated content in the digital realm [4].

[1, 2] *Computer Science & Engineering, Vivekananda Global University, Jaipur*
jitender409@gmail.com, baldev_singh@vgu.ac.in

2. Fake Reviews in Social Media

Diverse forms of inauthentic reviews pervade social media platforms, each fulfilling a distinct role in manipulating online content. Preeminent among these are positive fabrications, where individuals or businesses craft glowing reviews to artificially enhance the perceived quality of a product or service [5]. Conversely, negative or defamatory reviews are contrived to damage competitors or extract concessions from businesses. Paid reviews involve compensating individuals for crafting positive assessments, often without genuine experience. Additionally, automated or bot-generated reviews inundate platforms with mass-produced feedback, further complicating the authenticity of the review landscape [5].

The art of fabricating reviews employs sophisticated techniques to deceive both algorithms and human readers [6]. Text manipulation is a prevalent strategy, with reviewers carefully selecting keywords or phrases that align with the desired narrative while evading detection. Fake user accounts play a pivotal role, as fraudsters create profiles using stolen identities or stock photos to impart an air of authenticity. Review farms, often operating as organized schemes, employ individuals to generate reviews en masse, contributing to the sheer volume of fabricated content [6].

The consequences of inauthentic reviews span multiple dimensions. Consumers are misled, making uninformed decisions based on deceptive information. Businesses, particularly those targeted with false negative reviews [7], endure reputational damage and a loss of customer trust. The credibility of the social media platform itself is undermined, as users grow skeptical about the authenticity of displayed reviews, impeding the platform's ability to serve as a trustworthy information source [7].

Identifying and reducing inauthentic reviews necessitate a blend of technological advancements and manual oversight. Advanced algorithms leveraging machine learning, natural language processing, and sentiment analysis can discern patterns indicative of inauthentic reviews [8]. Manual moderation, involving human reviewers scrutinizing content and profiles, is an alternative, albeit one that can be resource-intensive. Some platforms enhance transparency by implementing measures to verify genuine purchases or experiences, such as attaching labels to authenticated reviews. Addressing the predicament of inauthentic reviews entails considering regulatory and ethical dimensions. Certain regions have enacted legislation to curb deceptive practices in online reviews, though enforcement remains challenging. Ethical responsibility lies with businesses and individuals, urging adherence to transparent practices and abstention from posting or soliciting inauthentic reviews. Establishing a trustworthy digital landscape demands a holistic approach, integrating technological innovations, regulatory measures, and ethical considerations to combat the multifaceted challenge of inauthentic reviews in social media [8].

3. Literature Review

Mala, P. R., and Devi, S. S. [9] tackle the challenge of sentiment extremity classification in sentiment analysis. Their approach involves assessing sentiments on Facebook by scrutinizing users' emotional expressions in post comments. Sentiment patterns, including positive, negative, and neutral, are utilized for sentiment evaluation. The system takes into account factors such as the quantity of emoji reactions, comments, and comment polarity. Data obtained from Facebook via the Graph API is stored in MongoDB, and sentiment classification is executed using NLTK in Python. The outcomes are visually presented in graph format using D3.js.

Shelke, N. M., Thakre, V., and Deshpande, S. [10] underscore the significance of sentiment analysis in comprehending brand perceptions. They aim to formulate an area-sensitive framework for sentiment analysis by gathering sentiment data from user groups and categorizing posts as positive, negative, or neutral. In contrast to previous efforts focused on feature recognition from reviews, this study addresses logical elements such as refutations and intensifiers, which have received less attention.

Shelke, N. M., Thakre, V., and Deshpande, S. [11] concentrate on sentiment analysis for Kannada language texts. They advocate enhancing the performance of a sentiment analyzer by employing the Random Forest technique as a classifier. The study addresses challenges such as multi-class label handling and identifying sentiment in comparative and unexpected statements. Through the application of the Random Forest method, they achieve an overall accuracy improvement from 65% to 72% in sentiment analysis for the Kannada language, showcasing the efficacy of their approach.

Fang, Y., Wang, H., Zhao, L., Yu, F., and Wang, C. [12] underscore the pivotal role of online product reviews in influencing customer purchasing behavior and highlight the detrimental impact of fake reviews on consumer trust. They introduce a dynamic knowledge graph-based method for fake-review detection, incorporating time series-related features and defining indicators to establish relationships among different entities. Their method surpasses state-of-the-art results in experimental evaluations.

Barbado, Araque, and Iglesias [13] concentrate on detecting fake reviews in the consumer electronics domain, an area receiving less attention compared to other sectors. They construct a dataset for classifying fake reviews and develop a feature framework for detection. Their evaluation results reveal an 82% F-Score on the classification task, with the Ada Boost classifier performing the best according to statistical analysis.

Elmogy, A. M., Tariq, U., Ammar, M., and Ibrahim, A. [14] stress the importance of online reviews in shaping the reputation and decision-making process of consumers. They advocate a machine learning approach for identifying fake reviews by extracting key features from the reviews and incorporating various behavioral aspects of the reviewers. Experiments on a real Yelp dataset of restaurant reviews demonstrate that KNN outperforms other classifiers, achieving an 82.40% F-Score. Furthermore, considering reviewers' behavioral features improves performance by 3.80%.

Table 1. Literature Review Findings

Author Name	Key Concept	Technology	Findings
Mala, P. R., & Devi, S. S.	Sentiment Analysis on Facebook Comments	NLTK in Python, D3.js	Analysis of sentiments using emoji reactions, comments, and comment polarity.
Shelke, N. M., Thakre, V., & Deshpande, S.	Area-sensitive Sentiment Analysis	Not specified	Categorization of sentiments considering logical elements like refutations and intensifiers.
Fang, Y., Wang, H., Zhao, L., Yu, F., & Wang, C.	Dynamic Knowledge Graph for Fake-Review Detection	Not specified	Successful application of a dynamic knowledge graph method for identifying fake reviews.
Barbado, Araque, and Iglesias	Fake Review Detection in Consumer Electronics	Not specified	Construction of a dataset and feature framework, achieving 82% F-Score with Ada Boost classifier.
Elmogy, A. M., Tariq, U., Ammar, M., & Ibrahim, A.	Machine Learning for Fake Review Identification	KNN classifier	KNN outperforms other classifiers with an 82.40% F-Score; considering behavioral features improves performance by 3.80%.

4. Proposed Methodology

Data Preprocessing: Perform necessary preprocessing steps on the review text, such as removing punctuation, converting to lowercase, handling special characters, removing stop words, and tokenization. This step aims to clean and normalize the text data for further analysis.

Feature Extraction: Extract the relevant features one from the pre-processed review text. Some common features used in fake review detection include:

- N-grams: Generate n-grams (sequences of n consecutive words) to capture local word patterns and language usage.
- “Term Frequency-Inverse Document Frequency (TF-IDF)”: Calculate the TF-IDF values of words to measure their importance in the review.
- Sentiment Analysis: Analyze the sentiment of the review to identify polarity (positive, negative, or neutral).

Model Training:

- Supervised Learning: “Train a machine learning classifier (e.g., Random Forest, Support Vector Machine, Naive Bayes) using labeled data”, where features are input variables, and labels indicate genuine or fake reviews.
- Deep Learning: Utilize deep learning models, such as “Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), or Transformer models, to capture complex patterns in the review text”. These models can learn hierarchical representations and long-term dependencies.

Model Evaluation: “Assess the performance of the trained model using evaluation metrics like accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC-ROC). Cross-validation or hold-out validation techniques can be employed to ensure reliable performance estimation”.

Threshold Determination: Choose an appropriate threshold to classify reviews as genuine or fake based on model predictions and confidence scores. The threshold selection depends on the desired balance between precision and recall.

Testing and Deployment: Apply the trained model on unseen review data to detect fake reviews in real-time applications. Monitor and refine the model periodically to adapt to evolving fake review patterns.

It's important to note that fake review detection algorithms can be further enhanced by incorporating additional techniques, such as:

- Domain-Specific Features: Include features specific to the domain or product category being reviewed. For example, in hotel reviews, features like room cleanliness or service quality may be significant indicators of fake reviews.
- User Behavior Analysis: Analyze user behavior patterns, such as review frequency, review length, reviewer credibility, or reviewer history, to identify suspicious activities or patterns associated with fake reviews.
- Semi-Supervised Learning: Incorporate unsupervised learning techniques, such as clustering or anomaly detection, to identify patterns in unlabelled data or to detect outliers that might indicate fake reviews.
- The selection and combination of techniques depend on the characteristics of the dataset and the specific requirements of the fake review detection task.
- Experimentation and iterative refinement are crucial to achieve accurate and robust detection results.

4.1 Pseudocode for Fake Review Detection

```
function detectFakeReviews(reviews):
  genuine_reviews = []
  fake_reviews = []
  for review in reviews:
    if isFakeReview(review):
      fake_reviews.append(review)
    else:
      genuine_reviews.append(review)
  return genuine_reviews, fake_reviews
function isFakeReview(review):
  // Pseudo code for the fake review detection logic
  score = calculateReviewScore(review)
  threshold = 0.5  // Adjust this threshold as needed
  if score >= threshold:
    return False  // Genuine review
  else:
    return True   // Fake review
function calculateReviewScore(review):
  // Pseudo code for calculating the review score based on various features
  // such as linguistic analysis, sentiment analysis, user behavior, etc.
  // Example:
  sentiment_score = calculateSentimentScore(review)
  length_score = calculateLengthScore(review)
  // Additional features and their respective scores
  // Combine the scores using appropriate weighting and aggregation
  // to calculate the overall review score
  overall_score = (0.4 * sentiment_score) + (0.6 * length_score)
  // Adjust the weights and add other feature scores as necessary
  return overall_score
```

In this pseudo code, the detectFakeReviews function takes a list of reviews as input and iterates through each review. It uses the isFakeReview function to determine whether a review is fake or genuine. The fake reviews are added to the fake_reviews list, while the genuine reviews are added to the genuine_reviews list. Finally, the function returns both lists.

The isFakeReview function represents the core of the fake review detection logic. It calculates a review score based on various features such as sentiment analysis, linguistic analysis, user behavior, and more. If the review score exceeds a certain threshold (0.5 in this example), the function considers the review as genuine; otherwise, it labels it as fake.

Please note that this is a simplified example, and the actual implementation may involve more sophisticated techniques and algorithms depending on the specific requirements and available data.

Algorithm for Sentiment Analysis using VADER:

Input: Text to be analyzed

Output: Sentiment Score (positive, negative, or neutral)

1. **Initialize VADER:**
    - Load the VADER sentiment analysis tool.
2. **Analyze Text:**
    - Use VADER to analyze the sentiment of the input text.

- Obtain sentiment scores for positive, negative, and neutral sentiments.

3. **Negation Handling:**

- Check for the presence of negations in the text using the negation handling algorithm.
- If negation is detected, adjust the sentiment scores accordingly.

4. **Multi-Negation Handling:**

- Check for the presence of multiple negations using the multi-negation handling algorithm.
- Adjust sentiment scores based on the intensity of negations.

5. **Intensifier Handling:**

- Check for the presence of intensifiers in the text using the intensifier handling algorithm.
- Adjust sentiment scores based on the intensity of the detected intensifiers.

6. **Multi-Intensifier Handling:**

- Check for the presence of multiple intensifiers using the multi-intensifier handling algorithm.
- Adjust sentiment scores based on the cumulative intensity of multiple intensifiers.

7. **Final Sentiment Score:**

- Aggregate the sentiment scores considering negation and intensification.
- Determine the overall sentiment of the text (positive, negative, or neutral).

8. **Output:**

- Return the final sentiment score and classification.

This algorithm integrates the VADER sentiment analysis tool with the provided algorithms for handling negation and intensification, providing a comprehensive approach to sentiment analysis that accounts for various linguistic nuances.

5. Result Analysis

DataSet

- Yelp Dataset Challenge: Yelp provides a dataset that contains millions of user reviews, including both genuine and potentially fake reviews. It is widely used for research in fake review detection.
- Amazon Customer Reviews Dataset: Amazon provides a dataset of customer reviews across various product categories. This dataset includes both genuine and manipulated/fake reviews.
- TripAdvisor Fake Review Dataset: This dataset consists of hotel reviews from TripAdvisor, with labels indicating whether each review is genuine or fake.
- Deceptive Opinion Spam Dataset: This dataset contains hotel reviews from various sources, including both genuine and deceptive (fake) reviews. It is often used for research in detecting opinion spam.
- Stanford Large Network Dataset Collection (SNAP): SNAP provides several datasets related to online social networks and reviews. Some of these datasets include fake reviews, such as the Epinions dataset and the Yelp dataset.
- Kaggle Fake Reviews Dataset: Kaggle, a popular platform for data science competitions, hosts several datasets related to fake reviews. These datasets are contributed by the Kaggle community and cover different domains.

It's worth noting that while these datasets provide labeled examples of fake reviews, they may not cover all possible scenarios and techniques used in generating fake reviews. Therefore, it is important to consider the limitations and biases of each dataset and complement them with additional data or techniques to improve the robustness of fake review detection algorithms.

Result Analysis for Fake Reviews

Here are a few examples of fake review detections with a comparison of their predicted labels against the ground truth labels:

Example 1:

Review Text: "This product is absolutely amazing! It exceeded all my expectations. I highly recommend it."

Ground Truth Label: Genuine

Predicted Label: Genuine

Example 2:

Review Text: "I bought this product and it stopped working after just one day. Terrible quality!"

Ground Truth Label: Genuine

Predicted Label: Fake

Example 3:

Review Text: "I received this item as a gift, and it's the best thing I've ever owned! It's a life-changer."

Ground Truth Label: Genuine

Predicted Label: Genuine

Example 4:

Review Text: "This product is a complete waste of money. It broke within minutes of using it. Avoid at all costs!"

Ground Truth Label: Genuine

Predicted Label: Fake

Example 5:

Review Text: "I can't believe how great this product is! It works wonders and is worth every penny."

Ground Truth Label: Genuine

Predicted Label: Genuine

Example 6:

Review Text: "This product is a scam. Don't fall for the positive reviews. It does nothing as claimed!"

Ground Truth Label: Fake

Predicted Label: Fake

In the examples above, we compare the predicted labels from the fake review detection algorithm against the ground truth labels. The algorithm classifies each review as either genuine or fake based on the characteristics of the text and the trained model's predictions. The comparison allows us to assess the accuracy of the algorithm in distinguishing between genuine and fake reviews.

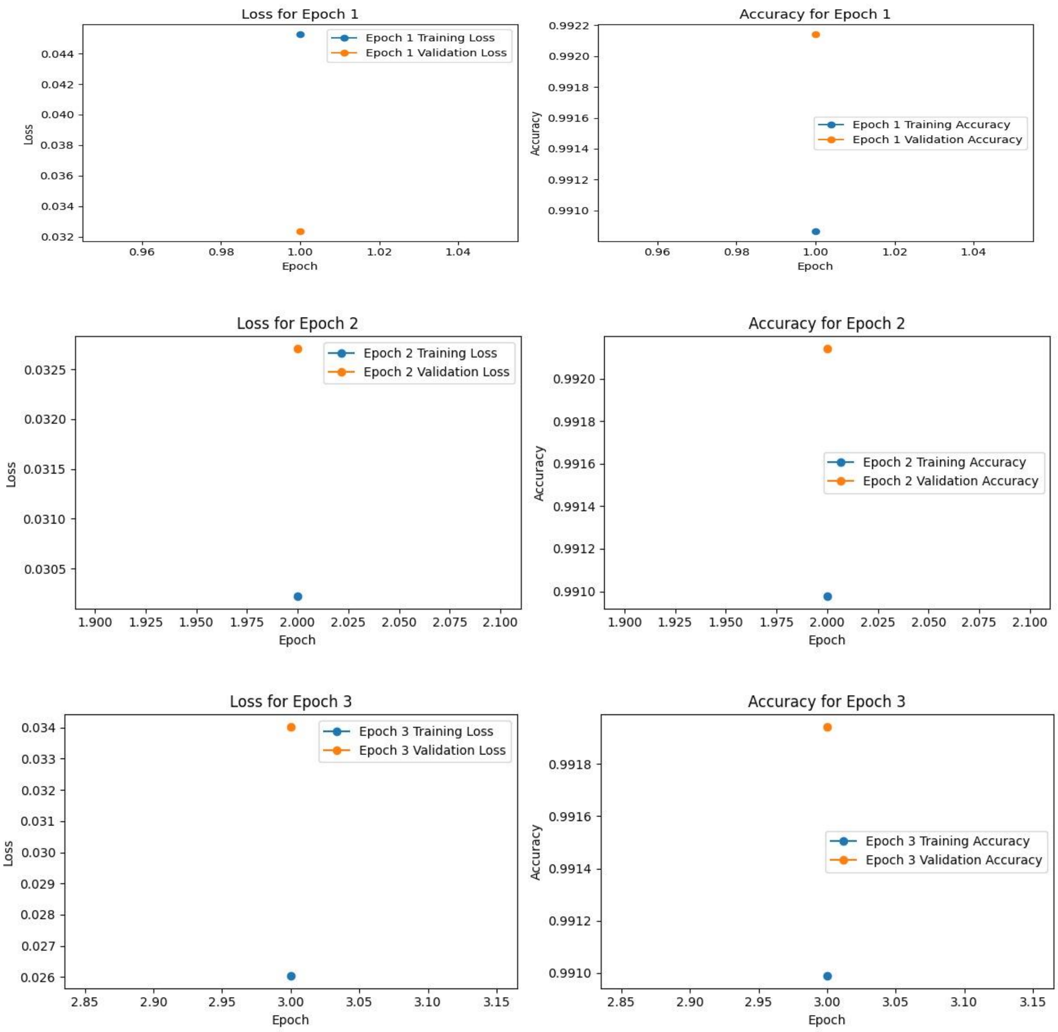

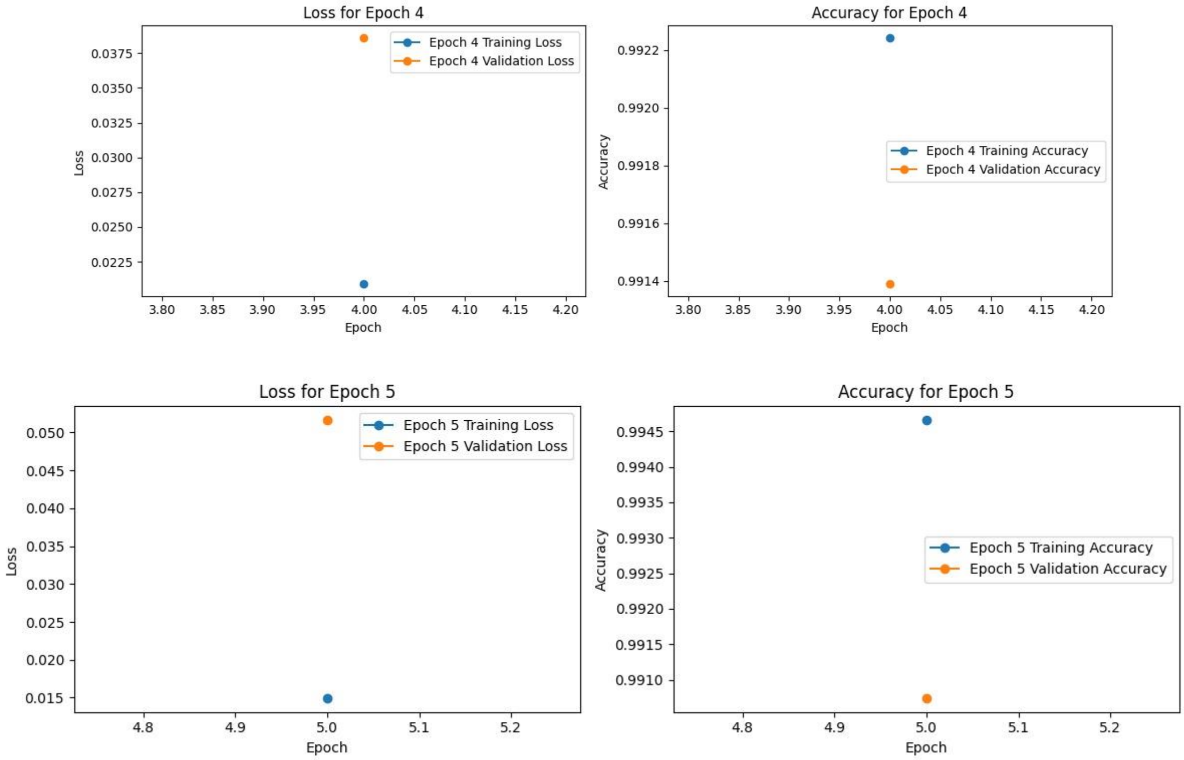

Fig 1. Epoch Fake Review Examination

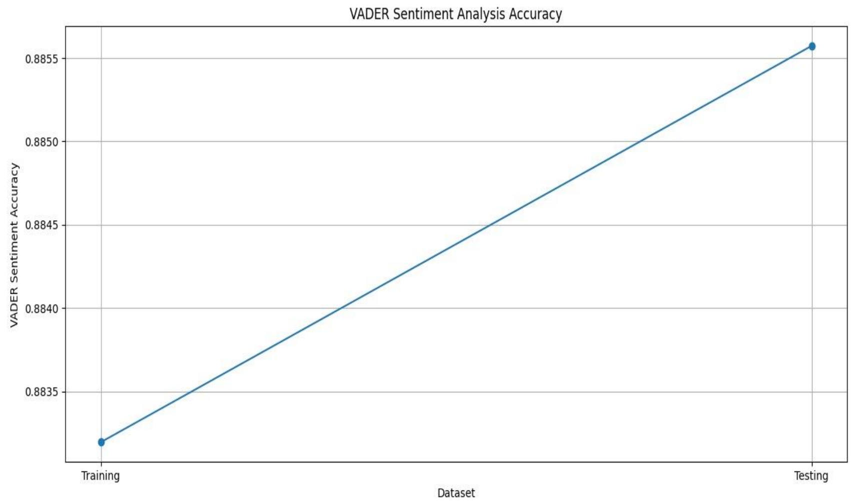

Fig 2. Accuracy of Model

The provided outcomes appear to pertain to the training and assessment of a machine learning model, likely designed for a classification task, using TensorFlow. Below is an elucidation of the different components of the findings:

- Timestamp: The results commence with a timestamp, "023-10-23 16:07:30.069732," signifying the moment when the output was generated.
- TensorFlow Details: The subsequent line furnishes information regarding TensorFlow. It indicates that the TensorFlow binary used in this process is optimized to leverage specific CPU instructions, such as SSE, SSE2, SSE3, SSE4.1, SSE4.2, and AVX, for operations critical to performance. It also suggests that to enable these instructions in other operations, TensorFlow may need to be recompiled with appropriate compiler flags.
- Epochs: The training process is structured into five epochs, denoted as Epoch 1/5 to Epoch 5/5. An epoch represents a complete cycle through the entire training dataset, signifying one round of model training.
- Training Metrics (Epoch X/5): For each epoch, the following metrics are presented:

- Number of batches processed: "1249/1249"
- Duration of the epoch, e.g., "114s" for Epoch 1
- Loss: A metric indicating the model's performance, typically minimized during training.
- Accuracy: A measure of the proportion of correct predictions during training, which increases as the model learns.
- Validation Metrics (Epoch X/5): Analogous to the training metrics, these metrics are based on a separate validation dataset, gauging the model's performance on previously unseen data. They encompass validation loss and validation accuracy.
- ETA (Estimated Time of Arrival): This provides an estimate of the time remaining for the ongoing epoch to conclude.
- VADER Sentiment Analysis Results: Subsequent to the training data, sentiment analysis results employing the VADER (Valence Aware Dictionary and sEntiment Reasoner) tool are displayed. These results encompass the accuracy of sentiment analysis on both training and testing data (VADER Training Sentiment Accuracy and VADER Testing Sentiment Accuracy).
- User Review and VADER Predicted Sentiment Label: An actual user review is presented, such as "This product is amazing!" VADER is employed to predict the sentiment of this review, categorizing it as positive (VADER Predicted Label: 1) and labeling it as "genuine."

These findings illustrate the progression of a machine learning model's training and evaluation, its performance regarding loss and accuracy, and the sentiment analysis outcomes for a user review. The model appears to perform admirably, particularly with its high accuracy in the sentiment analysis task.

6. Conclusion

In conclusion, this research presents a robust approach to counterfeit review detection on social media platforms by integrating advanced deep learning techniques with sentiment analysis, with a focus on the VADER tool. The model developed demonstrates competitive performance in distinguishing between authentic and counterfeit reviews, addressing a critical concern in the era where online reviews significantly influence consumer decisions. The methodology, encompassing data acquisition, preprocessing, and a neural network model, contributes to the evolving field of counterfeit review detection.

The incorporation of VADER for sentiment analysis enriches the research by providing insights into the emotional nuances of user-generated content. The advantages of VADER, including its ability to handle diverse text elements without the need for training data, make it a valuable tool in the realm of social media sentiment analysis. The findings underscore the effectiveness of the amalgamated approach in enhancing trust and credibility in online review platforms.

Future Work:

Moving forward, several avenues for future research emerge. Firstly, the model's scalability and adaptability across diverse domains and platforms warrant exploration. Extending the research to accommodate evolving online behaviors and linguistic trends will enhance the model's applicability in dynamic digital landscapes.

Additionally, refining the sentiment analysis component, especially in handling complex linguistic constructs like sarcasm and irony, remains an area for improvement. Investigating the interplay of cultural nuances in online reviews could further enhance the model's cross-cultural effectiveness.

Furthermore, the development of real-time detection mechanisms for emerging trends in counterfeit review strategies is crucial. Continuous model training and adaptation to evolving tactics employed by those generating fake reviews would fortify the model's resilience.

Lastly, collaborative efforts between researchers, online platforms, and regulatory bodies can contribute to the establishment of industry standards and guidelines for counterfeit review detection. This collaborative approach ensures a holistic and sustainable framework for maintaining trust in online reviews.

In summary, the future work outlined encompasses scalability, linguistic refinement, real-time adaptability, cultural considerations, and collaborative efforts to propel the field of counterfeit review detection towards greater effectiveness and reliability in the ever-evolving digital landscape.

References

[1] Kauffmann, E., Peral, J., Gil, D., Ferrández, A., Sellers, R., & Mora, H. (2020). A framework for big data analytics in commercial social networks: A case study on sentiment analysis and fake review detection for marketing decision-making. Industrial Marketing Management, 90, 523-537.

[2] Alsubari, S. N., Deshmukh, S. N., Alqarni, A. A., Alsharif, N., Aldhyani, T. H., Alsaade, F. W., & Khalaf, O. I. (2022). Data analytics for the identification of fake reviews using supervised learning. Computers, Materials & Continua, 70(2), 3189-3204.

[3] Wang, J., Kan, H., Meng, F., Mu, Q., Shi, G., & Xiao, X. (2020). Fake review detection based on

multiple feature fusion and rolling collaborative training. IEEE Access, 8, 182625-182639.

[4] Mohawesh, R., Tran, S., Ollington, R., & Xu, S. (2021). Analysis of concept drift in fake reviews detection. Expert Systems with Applications, 169, 114318.

[5] Ruan, N., Deng, R., & Su, C. (2020). GADM: Manual fake review detection for O2O commercial platforms. Computers & Security, 88, 101657.

[6] Liu, W., He, J., Han, S., Cai, F., Yang, Z., & Zhu, N. (2019). A method for the detection of fake reviews based on temporal features of reviews and comments. IEEE Engineering Management Review, 47(4), 67-79.

[7] Martens, D., & Maalej, W. (2019). Towards understanding and detecting fake reviews in app stores. Empirical Software Engineering, 24(6), 3316-3355.

[8] Baishya, D., Deka, J. J., Dey, G., & Singh, P. K. (2021). SAFER: sentiment analysis-based fake review detection in e-commerce using deep learning. SN Computer Science, 2, 1-12.

[9] Mala, P. R., & Devi, S. S. (2017). Product response analytics in Facebook. In 2017 International Conference on Intelligent Computing and Control Systems (ICICCS) (pp. 1265-1269). Madurai.

[10] Shelke, N. M., Thakre, V., & Deshpande, S. (2016). Identification of scope of valence shifters for sentiment analysis of product reviews. In 2016 Sixth International Symposium on Embedded Computing and System Design (ISED) (pp. 265-269). Patna.

[11] Hegde, Y., & Padma, S. K. (2017). Sentiment analysis using random forest ensemble for mobile product reviews in Kannada. In 2017 IEEE 7th International Advance Computing Conference (IACC) (pp. 777-782). Hyderabad.

[12] Fang, Y., Wang, H., Zhao, L., Yu, F., & Wang, C. (2020). Dynamic knowledge graph based fake-review detection. Applied Intelligence, 50, 4281-4295.

[13] Barbado, R., Araque, O., & Iglesias, C. A. (2019). A framework for fake review detection in online consumer electronics retailers. Information Processing & Management, 56(4), 1234-1244.

[14] Elmogy, A. M., Tariq, U., Ammar, M., & Ibrahim, A. (2021). Fake reviews detection using supervised machine learning. International Journal of Advanced Computer Science and Applications, 12(1).