Communications on Applied Nonlinear Analysis
ISSN: 1074-133X
Vol 32 No. 9s (2025)

AI Based Knowledge Representation on Spam Review Detection by Bottom -Up Learning Approach

Sripathi.S[1], Shanthi. P.M[2]

[1]Research Scholar, Department of Computer Science, JJ College of Arts and Science, Puthukottai, Bharathidasan University, Tiruchirapalli, Tamilnadu, India-620124
ORCID ID: 0009-0003-0188-4952
[2]Assistant Professor, Department of Computer Science, JJ College of Arts and Science, Puthukottai, Tamilnadu, India-622422
ORCID ID: 0000-0002-9318-6312

E-mail: [1] sripathimarch1982@gmail.com, [2] shanthisuman28@gmail.com

Article History:

Received: *07-01-2025*

Revised: *06-02-2025*

Accepted: *18-03-2025*

Abstract:

A bottom-up learning approach for AI-based knowledge representation in spam review detection utilizes machine learning and natural language processing to uncover fraudulent reviews. This method begins with raw data, identifying patterns and extracting features from text, metadata, and user behaviour. Unlike traditional rule-based top-down techniques, it gradually builds knowledge through unsupervised or semi-supervised learning, enhancing accuracy over time. Deep learning, sentiment analysis, and clustering techniques aid in distinguishing authentic from fake reviews. By continuously adapting to new data, this approach improves spam detection, making it more effective against evolving deceptive tactics on online platforms. Detecting spam reviews is crucial for maintaining the integrity of online platforms. AI-based knowledge representation using a bottom-up learning approach provides an effective solution for identifying fraudulent reviews. Unlike traditional rule-based methods that depend on predefined patterns, this approach begins with raw data and gradually builds knowledge by analysing text, metadata, and user behaviour leveraging unsupervised or semi-supervised learning, the system enhances its accuracy over time, allowing it to adapt to evolving spam tactics. Techniques such as deep learning, sentiment analysis, and clustering help differentiate authentic from fake reviews by examining writing styles, behavioural trends, and sentiment inconsistencies. This enables the model to detect deceptive reviews more effectively. A key advantage of this bottom-up approach is its ability to continuously learn from new data, making it more dynamic than static rule-based systems. It refines its detection mechanisms in real time, helping platforms stay ahead of increasingly sophisticated spam techniques. The adaptability of AI-driven models ensures that fraudulent reviews are identified with greater precision, fostering trust in online reviews across e-commerce, travel, and service-oriented platforms. In summary, AI-based knowledge representation through a bottom-up learning approach offers a scalable and intelligent method for spam detection. Its capacity to evolve with emerging threats makes it an essential tool for preserving the authenticity of online reviews, enhancing transparency, and boosting consumer confidence in digital marketplaces.

We present an innovative comparative analysis of six machine learning models for spam review detection: Long Short-Term Memory Networks (LSTM), Stochastic Gradient Descent (SGD), Classification Via Regression (CVR), PART, Random Tree (RT), and K-Star. Our study provides a thorough evaluation of these models on a large-scale Amazon product review dataset containing 26.7 million reviews. To improve model performance and efficiency, we implement a rigorous feature selection process. The analysis is conducted using multiple evaluation metrics, including accuracy, precision, recall, F-measure, ROC, and PRC, ensuring a comprehensive assessment of each model's effectiveness.

Keywords: Long Short-Term Memory Networks (LSTM), Stochastic Gradient Descent (SGD), Classification Via Regression (CVR), PART, Random Tree (RT), and K-Star

1. Introduction

In today's digital landscape, online reviews significantly impact consumer decisions, shaping purchasing behaviour across e-commerce platforms. Shoppers depend on these reviews to evaluate product quality, compare options, and make informed choices. However, the rise of spam reviews—fraudulent, biased, or manipulated feedback—poses a major challenge, eroding trust and transparency in online marketplaces. These deceptive reviews, often generated by bots or malicious users, aim to mislead buyers, artificially boost product ratings, or harm competitors. Combating this issue requires advanced and intelligent spam detection mechanisms capable of differentiating genuine reviews from fraudulent ones. Artificial Intelligence (AI) has emerged as a powerful solution to address spam detection challenges, offering automated, scalable, and adaptive techniques. One of the most effective strategies in this domain is AI-based knowledge representation, which structures and analyses data to uncover meaningful patterns. Within this framework, the **bottom-up learning approach** plays a vital role by incrementally building knowledge from raw data instead of relying on predefined rules or top-down heuristics. This approach enhances adaptability, enabling the system to dynamically evolve in response to emerging spam patterns and deceptive tactics.

1.1 Understanding Knowledge Representation in AI

In AI, knowledge representation involves structuring and organizing information so that machines can interpret, learn from, and make informed decisions. For spam review detection, this process entails encoding textual content, user behaviour, and metadata into structured formats that enable pattern recognition. Unlike conventional rule-based methods that depend on manually designed heuristics, AI-driven knowledge representation leverages machine learning (ML) and natural language processing (NLP) to automatically identify and classify spam reviews based on linguistic, statistical, and behavioural attributes.A major advantage of AI-based knowledge representation is its ability to generalize from data and recognize patterns beyond simple keyword matching. By utilizing deep learning models, sentiment analysis, and clustering techniques, AI systems can detect subtle irregularities in reviews, such as unnatural writing styles, excessive repetition, or abrupt increases in review activity. Additionally, these models incorporate contextual factors like user credibility, review timestamps, and sentiment polarity, further improving the accuracy of spam detection.

1.2 The Bottom-Up Learning Approach in Spam Detection

The bottom-up learning approach is a data-driven methodology that incrementally builds knowledge by starting with raw data and progressively refining its understanding. Unlike the top-down approach, which relies on predefined rules and assumptions, the bottom-up method enables systems to learn organically from real-world data. This flexibility is particularly advantageous in spam detection, where deceptive tactics continuously evolve to bypass traditional rule-based detection methods.

Within a bottom-up learning framework, machine learning models analyze large volumes of review data to uncover characteristics indicative of spam. This process typically involves:

- ✓ **Feature Extraction** – Identifying key attributes from review text, metadata, and user behaviour, such as word frequency, sentiment scores, reviewer credibility, and temporal patterns.
- ✓ **Unsupervised or Semi-Supervised Learning** – Applying clustering algorithms and deep learning techniques to discover hidden patterns in data, even in the absence of labelled examples.
- ✓ **Pattern Recognition and Classification** – Training classifiers to differentiate between genuine and spam reviews based on learned representations.
- ✓ **Continuous Learning and Adaptation** – Updating the model as new data becomes available, refining detection mechanisms over time.

By continuously learning from user interactions, emerging spam trends, and evolving linguistic patterns, the bottom-up approach enhances real-time spam detection. This adaptability strengthens spam detection systems, ensuring the accurate identification of fraudulent reviews with increasing precision.

1.3 Key AI Techniques for Spam Review Detection

Several AI-driven techniques contribute to effective spam detection within a bottom-up learning framework:

- ✓ **Natural Language Processing (NLP)** – NLP techniques extract semantic and syntactic features from review text, analyzing grammatical structures, sentiment scores, and keyword usage to differentiate authentic from fraudulent reviews.
- ✓ **Machine Learning Classifiers** – Both supervised and unsupervised ML algorithms, including Support Vector Machines (SVMs), Decision Trees, Random Forests, and Neural Networks, classify reviews based on identified patterns. Deep learning models such as Recurrent Neural Networks (RNNs) and Transformer-based models excel at capturing contextual dependencies in textual data.
- ✓ **Sentiment Analysis** – Evaluating sentiment distribution across reviews helps detect inconsistencies or unnatural sentiment spikes that may indicate spam.
- ✓ **Behavioral Analysis** – Examining reviewer behavior, including review frequency, length patterns, and IP addresses, provides additional insights into spam tendencies beyond textual content.

- ✓ **Graph-Based Detection** – Mapping relationships between reviewers, products, and ratings using graph-based techniques reveals suspicious review networks and bot-driven spam campaigns.

Advantages of AI-Based Bottom-Up Learning for Spam Detection

The AI-driven bottom-up learning approach offers several benefits over traditional rule-based or heuristic-driven spam detection methods:

- **Adaptability to Evolving Spam Tactics** – AI models continuously learn from new data, making them more resilient to emerging deceptive strategies.
- **Enhanced Detection Accuracy** – Machine learning and deep learning techniques improve classification precision by capturing complex patterns within review data.
- **Reduced Reliance on Manual Rule-Making** – Unlike top-down approaches that require predefined rules, the bottom-up approach automates pattern recognition, minimizing human intervention.
- **Scalability for Large Datasets** – AI-powered models efficiently process vast amounts of review data, making them ideal for large-scale e-commerce platforms.
- **Real-Time Spam Detection** – By utilizing real-time data processing, AI models can detect and flag fraudulent reviews as they appear, preventing the spread of misinformation.

This approach ensures a dynamic and efficient spam detection system, strengthening the credibility and transparency of online reviews.

2. Literature Review

Fake reviews are written without any real experience of the product or service they're about [3]. They can be written by people, or increasingly, generated by sophisticated AI. Because AI-generated text is so advanced, it's now very difficult to tell the difference between a real review and a fake one [4]. Online platforms make it easy for these fake reviews to spread widely, and since so many people read reviews, there's a big incentive to create them. This allows fake reviews to be produced at scale, making it much harder to combat the problem. Opinion spam is similar to fake reviews, as it involves writing false opinions to mislead people [1]. This can include creating seemingly authentic "fictitious opinions" [5]. Incentivized reviews are another related issue, where people are encouraged to write reviews through marketing campaigns like influencer endorsements [6] or by getting free products [7]. The key difference is that while fake reviews are often anonymous or AI-generated, incentivized reviews usually come from known influencers or real people. AI-generated spam reviews are much harder to spot than human-written ones for several reasons. AI excels at mimicking human writing thanks to advanced Natural Language Processing, making it tough to tell real from fake. The sheer volume of AI-generated reviews can overwhelm existing detection systems. Also, AI can create sophisticated, seemingly genuine reviews, making simple keyword or rule-based detection ineffective. The fact that AI can constantly adapt and improve in response to detection methods makes it even harder to catch. Finally, AI-generated content is often well-structured and contextually relevant, fooling both computer systems and human reviewers. Because AI can mimic how real people write, spotting spam becomes even harder. This highlights the need for strong detection systems that use Deep Learning and advanced Natural Language Processing to effectively combat this problem.

Researchers have explored several ways to spot fake reviews. One basic method is manual analysis, based on the idea that people can detect lies [8]. For example, some research has focused on rules to distinguish incentivized reviews from regular ones, looking at things like review length, sentiment, and helpfulness [7]. Other studies have identified general patterns in reviews, such as limited reviews per user and product, and low feedback [9], with deviations from these patterns suggesting fake reviews. Another approach examines how people judge review trustworthiness, considering factors like content, writing style, images, length, detail, and overly positive or negative language [10]. However, human judgment can be unreliable. Human review analysis isn't very accurate. For instance, one study found humans could only spot fake reviews 65% of the time, while a machine learning model did it 86% of the time [5]. Other studies show even worse human accuracy, around 52-57% [11,12], even when people are aware of the problem. The sheer volume of online reviews makes manual review impossible. TripAdvisor alone gets over 200 million reviews [2]. With millions of products and thousands of reviews per product, manual methods just don't scale. This has led researchers to focus on automated solutions. Online review spam, sometimes similar to opinion spam, involves deceptive tactics like creating fake reviews to trick consumers into thinking they're from real customers. Adding to the problem, some dishonest sellers even use people to spread fake reviews. This mix of factors makes it really hard to identify fraudulent reviews. On the other hand, algorithms can find patterns in data that humans might miss.
According to [13], the automated methods can be divided into
(a) content-based detection (focused on the textual content of reviews)
(b) behavior-based detection (focused on unusual and suspicious behaviours)
(c) information-based detection (focused on product characteristics) and
(d) spammer group detection (focused on identifying connections between reviewers).
Researchers have mainly concentrated on either text-based or non-text-based methods for detecting fake reviews. Those using Natural Language Processing (NLP) [11,14–16] analyze text features like keywords, phrases, n-grams, punctuation, semantic similarity, hidden topics, and writing style. Others [16] look at non-textual features such as user ID, location, and review count to distinguish real reviews from fake ones (whether human or AI-generated). Some research [2,10] combines both textual and non-textual features to improve detection accuracy.

The authors [2, 3] contributed significantly to spam review detection by defining three main types of problematic online reviews: First, "untruthful opinions" are reviews that intentionally mislead people or review analysis systems. This includes writing overly positive reviews for many products to boost them, and writing unfairly negative reviews for a few products to damage their reputation. Second, some reviews focus only on brand names, offering no real information about the product itself. Please specify the brands, manufacturers, or Suppliers of the aforementioned items. While brand-focused reviews might offer *some* perspective, they're considered spam because they're biased and lack specific product details. Beyond these, there are two other types of non-review spam:

(1) notices and

(2) irrelevant content (like random writing, questions, and answers) that don't offer any actual review information.

These three categories cover most online spam reviews. These deceptive reviews can significantly influence buying decisions and harm new businesses. Because of this, a lot of research has been done on how to detect these fraudulent product and ad reviews. Spammers often use online review platforms for malicious purposes, like promoting fake products or services, stealing social media contacts, or tricking customers into offline deals. These tactics confuse users and create a big challenge for anti-spam systems [5]. Much research focuses on understanding these deceptive techniques. This paper collects common misleading and incorrect viewpoints identified as "spam reviews" [2], specifically focusing on the two categories previously mentioned. Spam reviews aim to manipulate the perceived value of certain items and mislead users, unlike genuine product assessments. These reviews can come from various sources, including dishonest sellers, fraudulent groups, or individual marketers, and they often exhibit a wide range of patterns intentionally created by spammers [5,6].

3. Materials and Methods

This section outlines the materials and methods used in the research, specifically focusing on Amazon product reviews labeled as spam or not spam. The dataset comprises 26.7 million reviews, with 15.4 million labeled reviews. The classification labels are assigned as follows: spam reviews are marked as 0, while genuine (not spam) reviews are marked as 1. The architecture illustrates the workflow of this research, depicting the process from dataset collection to analysis. The collected dataset undergoes image filtering and feature selection using learning models within the WEKA 3.9.5 open-source tool. A 10:90 sampling technique is applied to ensure effective data distribution and model training.

This work considers following algorithms:

- Long Short-Term Memory Networks (LSTM)
- SGD
- Classification Via Regression (CVR)
- PART
- Random Tree (RT)
- K-Star (An instance-based classifier)

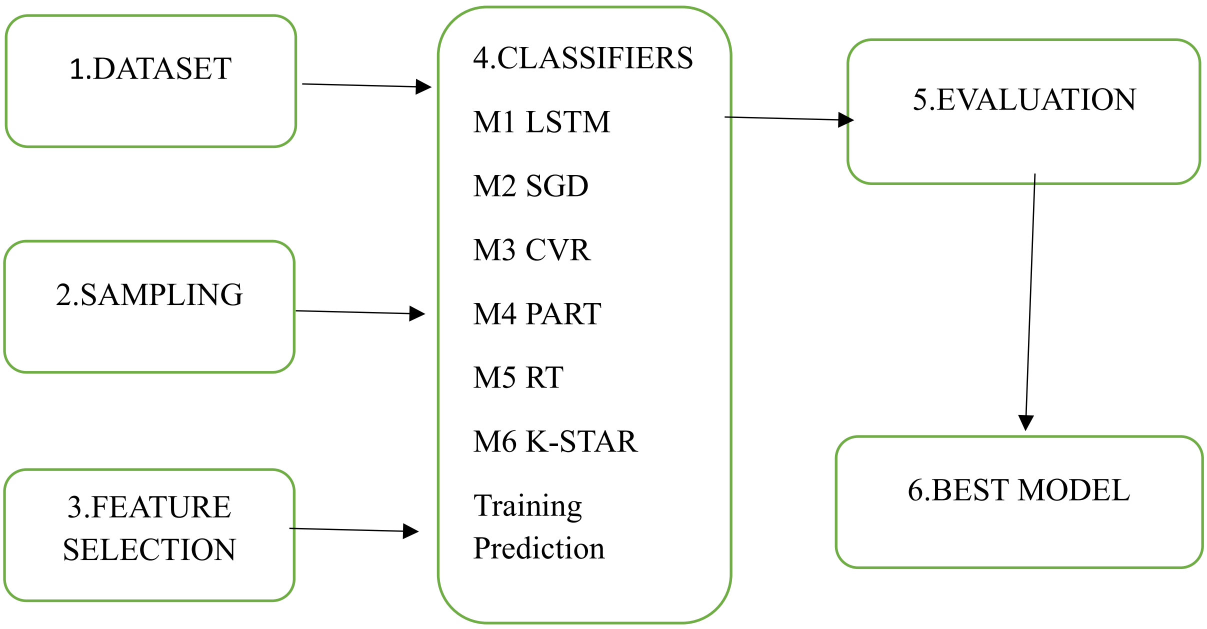

Figure 1: Proposed Architecture

Our proposed methodology for spam review detection employs a comparative analysis of six distinct machine learning models, aiming to tackle the intricate nature of spam reviews by examining various algorithmic approaches. The key innovation of this approach lies in the extensive evaluation of these models on a large-scale, real-world dataset, combined with a sophisticated feature selection process to enhance detection accuracy and efficiency.

3.1 Dataset and Preprocessing

We employ a large-scale Amazon product review dataset containing 26.7 million reviews, making it a valuable resource due to its real-world relevance and extensive size. This dataset serves as a strong foundation for training and evaluating our models. Our preprocessing steps include:

1. Data Cleaning: Eliminating irrelevant content and standardizing text formatting.
2. Feature Extraction: Leveraging both linguistic and behavioural attributes from the reviews.
3. Labeling: Assigning binary classifications - spam (0) and not spam (1) - To distinguish fraudulent from genuine reviews

3.2 Feature Selection

A key contribution of our research is the development of an enhanced feature selection process. We implement a hybrid approach that integrates filter and wrapper methods to identify the most relevant features for spam detection. This process consists of:

1. Information Gain Analysis: Ranking features based on their ability to distinguish between spam and genuine reviews.
2. Correlation-Based Feature Selection: Eliminating redundant features while retaining those most relevant to the classification task.
3. Recursive Feature Elimination: Iteratively refining the feature set through model training to optimize performance.

This strategic feature selection significantly improves model efficiency and interpretability, leading to more accurate spam detection.

3.3 Model Selection and Implementation

We carefully selected six diverse machine learning models to capture different aspects of the spam detection problem:

1. Long Short-Term Memory Networks (LSTM): Chosen for its efficiency in text classification tasks.
2. Stochastic Gradient Descent (SGD): Selected for its ability to handle large-scale learning problems.
3. Classification Via Regression (CVR): Employed for its potential to capture complex, non-linear relationships in the data.
4. PART: Utilized for its rule-based approach, offering interpretable decision-making.
5. Random Tree (RT): Chosen for its ability to handle high-dimensional data and capture feature interactions.
6. K-Star: Selected for its instance-based learning approach, potentially useful for detecting nuanced spam patterns.

3.4 Training and Evaluation Framework

We ensure robust and unbiased results through a rigorous training and evaluation process:

1. Data Splitting: This includes splitting our data into training and testing sets using a 90:10 ratio.
2. Cross-Validation: To guarantee model stability and generalizability, we employ 10-fold cross-validation.
3. Performance Metrics: Performance is evaluated using a comprehensive suite of metrics: accuracy, precision, recall, F-measure, ROC, PRC, Kappa, and MCC.
4. Error Analysis: Finally, we conduct detailed error analysis using MAE, RMSE, RAE, and RRSE to understand model behaviour and identify limitations.

3.5 Model Comparison and Selection

In the final stage, we rigorously compare model performance across all metrics. Statistical significance tests are used to confirm the superiority of the top-performing model.

3.6 Interpretability Analysis

To ensure practical application, we analyse the best-performing model for interpretability. This involves feature importance ranking and partial dependence plots to reveal the model's decision-making process. This algorithm describes the training, evaluation, and selection process for the optimal model to classify spam reviews.

4. Outcome and Interpretations

This section focuses the outcome of CVR, PART, LSTM, SGD, RT and K-Star Models

Table 1: Classifiers VS Outcomes

Classifier	Accuracy	Precision	Recall	ROC	PRC
RT	77.4%	0.82	0.76	0.87	0.84
CVR	82.70%	0.83	0.82	0.91	0.9
PART	79.70%	0.78	0.80	0.87	0.87
LSTM	89.12%	0.92	0.87	0.93	0.93
SGD	79.81%	0.8	0.78	0.86	0.86
K-Star	76.70%	0.75	0.77	0.78	0.72

The above table 1 shows the Accuracy, Precision, recall, receiver operating characteristic curve (ROC) and precision recall curve (PRC) value of RT, CVR, PART, LSTM, SGD and K-Star models.

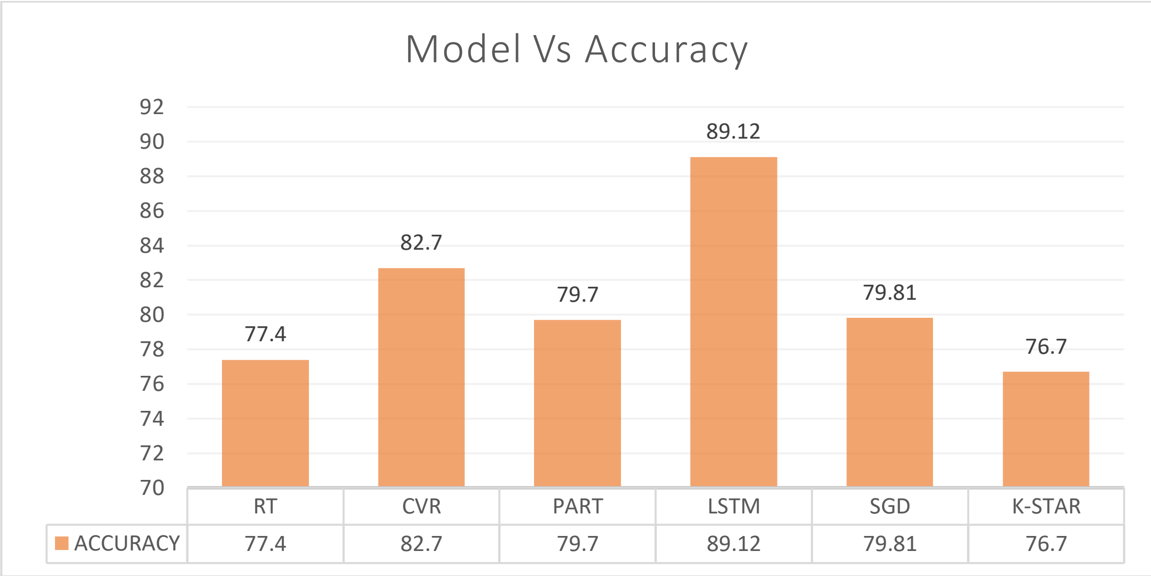

Figure 2: Model Vs Accuracy

Figure 2 shows the accuracy of several models: RT, CVR, PART, LSTM, SGD, and K-Star. LSTM performed best with 89.12% accuracy, while K-Star had the lowest at 76.70%. RT achieved 77.4% accuracy, and PART reached 79.70%.

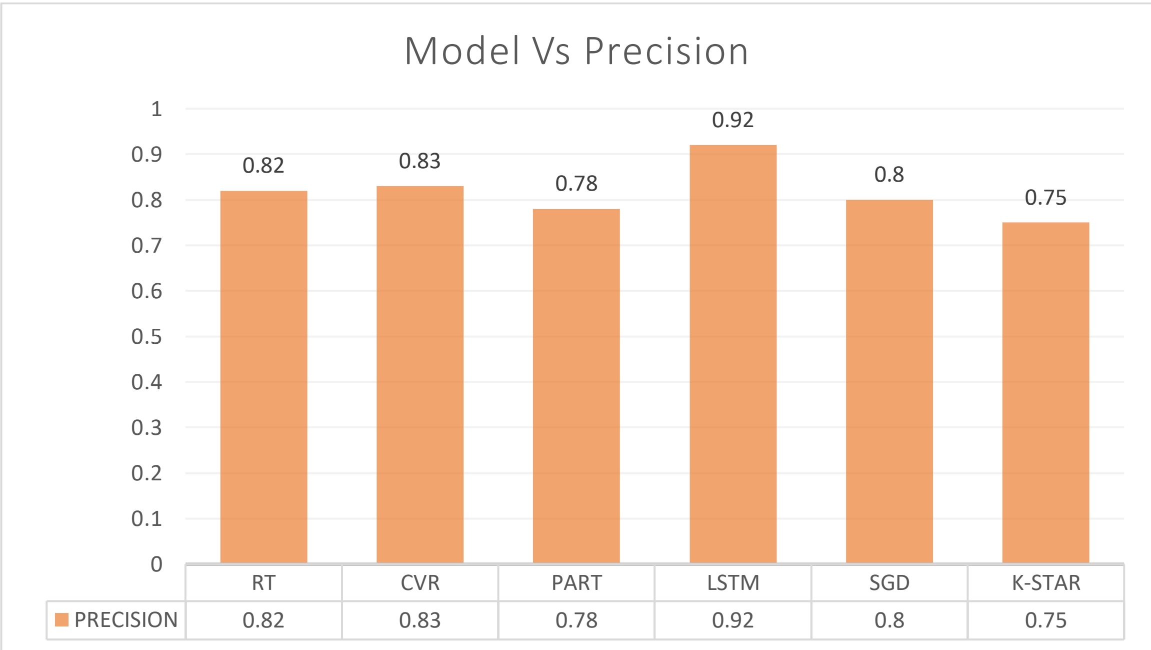

Figure 3: Model Vs Precision

Figure 3 shows the precision of the RT, CVR, PART, LSTM, SGD, and K-Star models. LSTM and CVR performed best, both achieving 0.92 precision. K-Star had the lowest precision at 0.75. PART reached 0.78 precision, and SGD had 0.8.

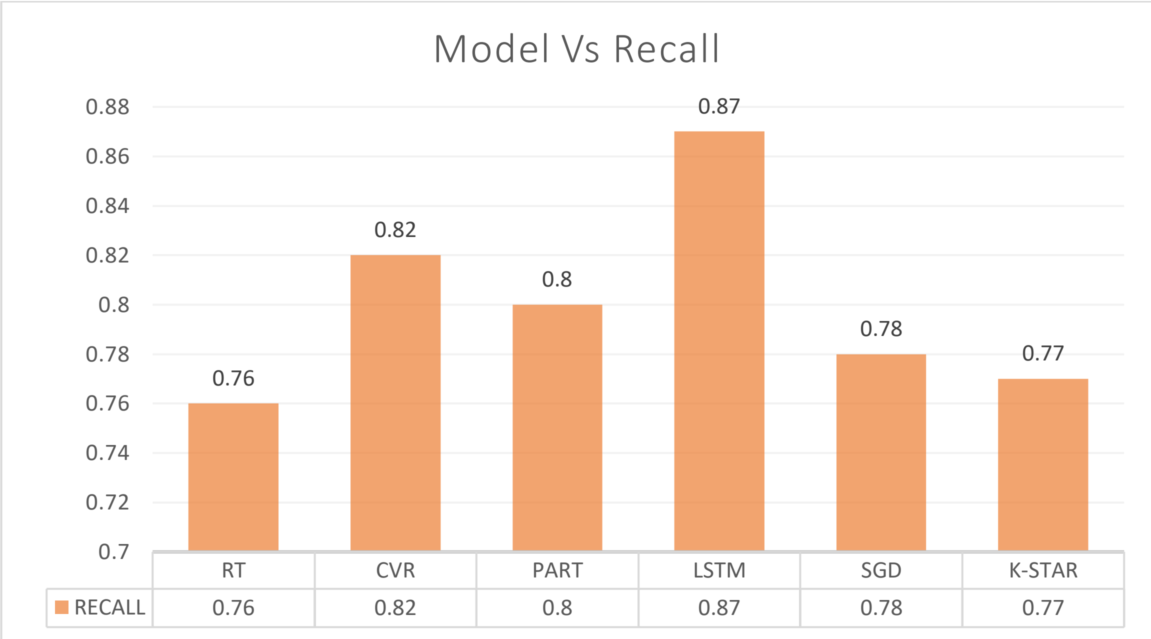

Figure 4: Model Vs Recall

Figure 4 displays the recall scores for the RT, CVR, PART, LSTM, SGD, and K-Star models. LSTM had the highest recall at 0.87, followed by CVR. K-Star had the lowest recall at 0.77. RT's recall was 0.76, PART was 0.8, and SGD was 0.78.

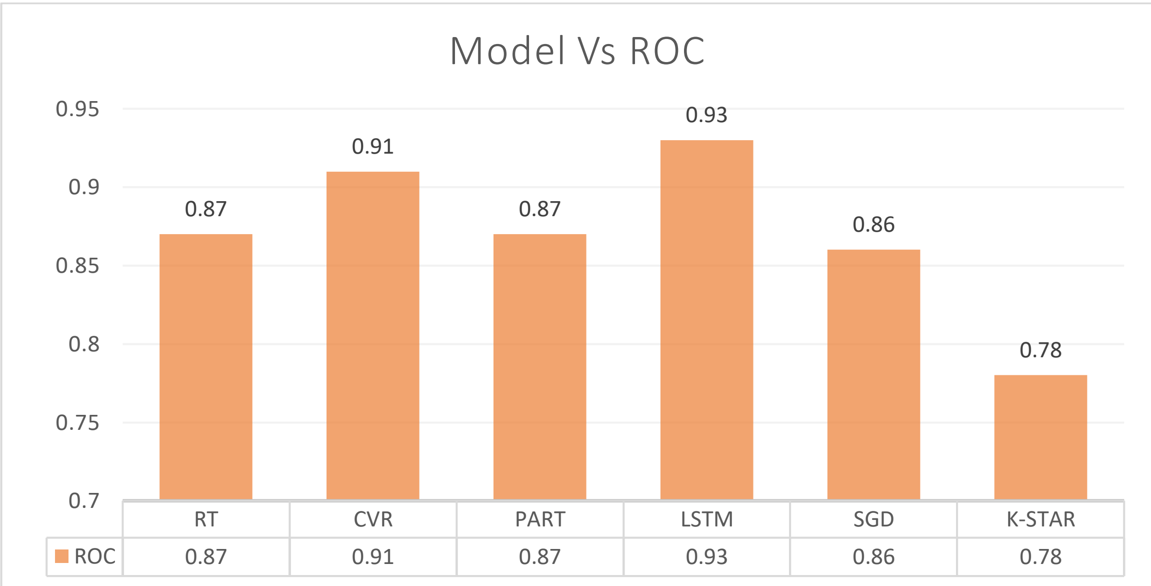

Figure 5: Model Vs ROC

Figure 5 shows the ROC (Receiver Operating Characteristic) values for the RT, CVR, PART, LSTM, SGD, and K-Star models. LSTM achieved the highest ROC score of 0.93, followed by CVR at 0.91. K-Star had the lowest ROC score of 0.78. RT and PART both had a ROC of 0.87, while SGD's was 0.86.

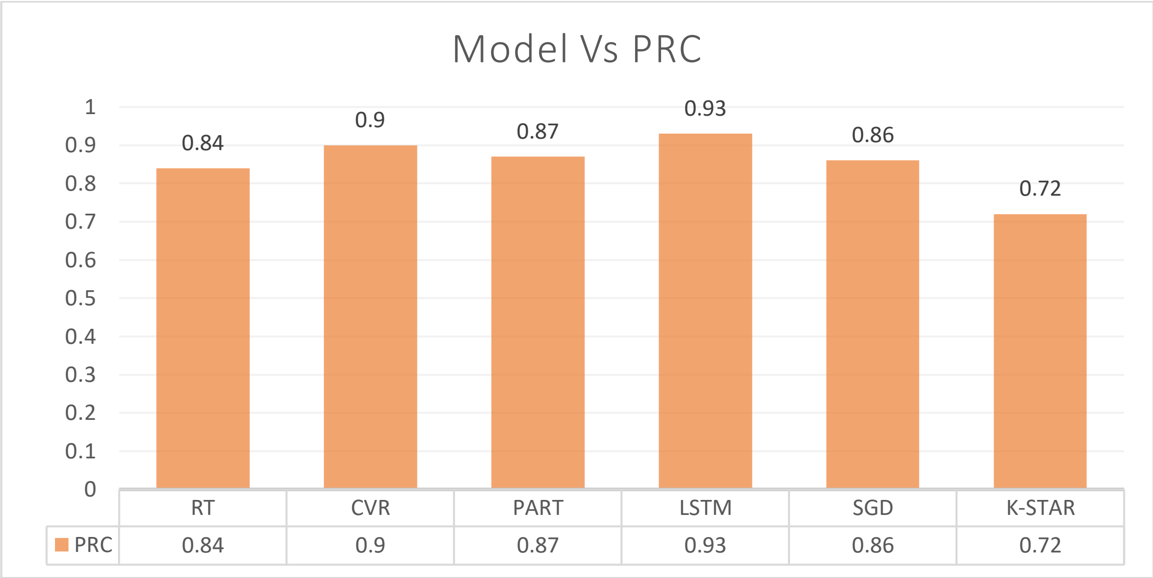

Figure 6: Model Vs PRC

Figure 6 presents the PRC (Precision-Recall Curve) values for the RT, CVR, PART, LSTM, SGD, and K-Star models. LSTM had the highest PRC at 0.93, followed by CVR at 0.90. K-Star had the lowest PRC at 0.72. The PRC for RT was 0.84, for PART it was 0.87, and for SGD it was 0.86.

Table 2: Classifiers Vs Outcomes

Classifier	Time	Kappa	F-Measure	MCC
RT	0.08	0.45	0.75	0.53
CVR	0.12	0.55	0.81	0.58
PART	3.1	0.48	0.27	0.52
LSTM	0.05	0.57	0.83	0.59
SGD	0.06	0.51	0.77	0.54
K-Star	1	0.39	0.76	0.46

The time consumption, Kappa, F-measure and Mathews Correlation Coefficient value (MCC) of the chosen models are presented in Table 2.

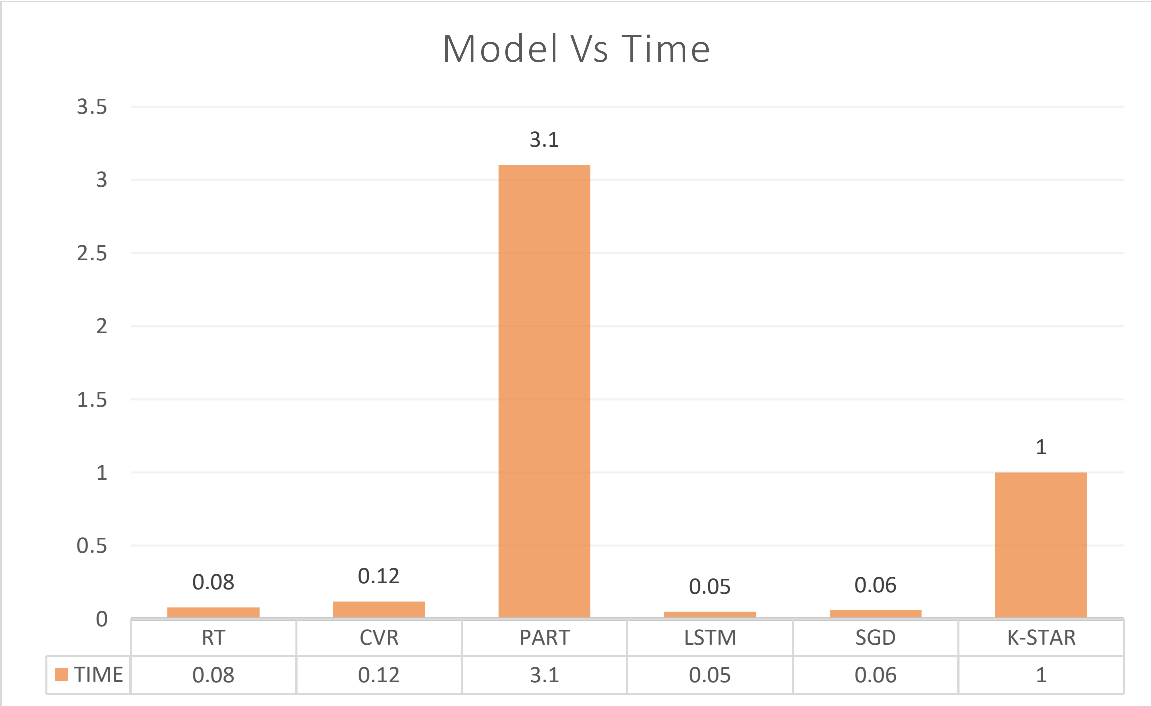

Figure 7: Model Vs Time

Figure 7 shows the model creation time (in seconds) for RT, CVR, PART, LSTM, SGD, and K-Star. PART took the longest to build, at 3.1 seconds. SGD and RT were the fastest, at 0.06 and 0.08 seconds, respectively. CVR took 0.12 seconds, LSTM took 0.05 seconds, and K-Star took 1 second.

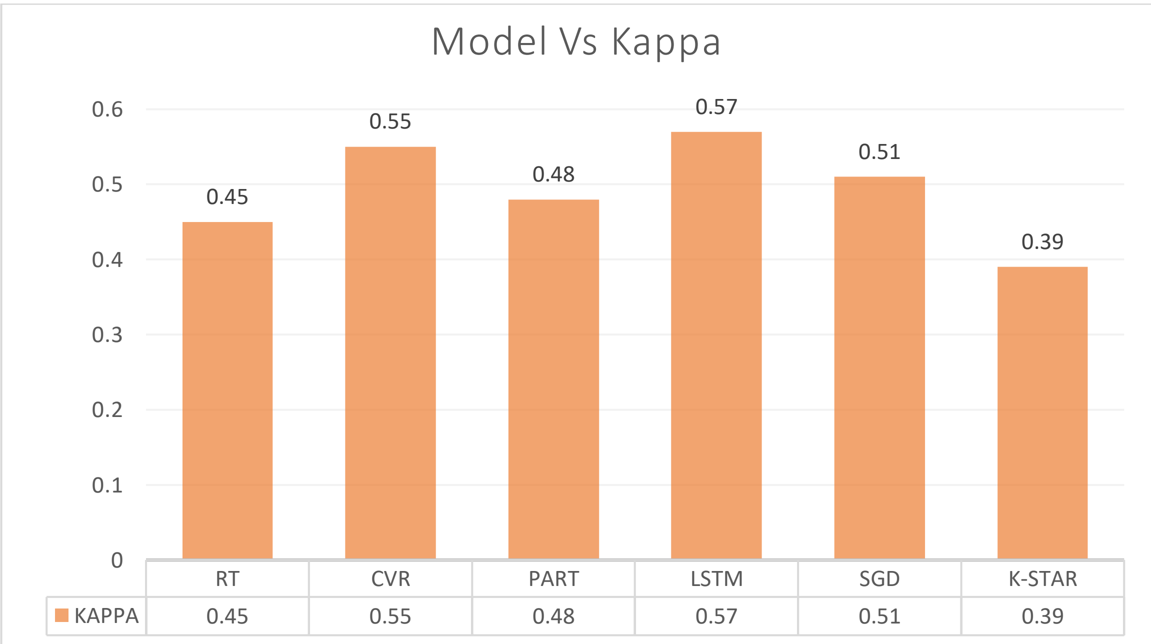

Figure 8: Model Vs Kappa

Figure 8 shows the Kappa values for the RT, CVR, PART, LSTM, SGD, and K-Star models. LSTM achieved the highest Kappa value of 0.57, followed by CVR. K-Star had the lowest Kappa value of 0.39. SGD had a Kappa of 0.51, and RT had a Kappa of 0.45.

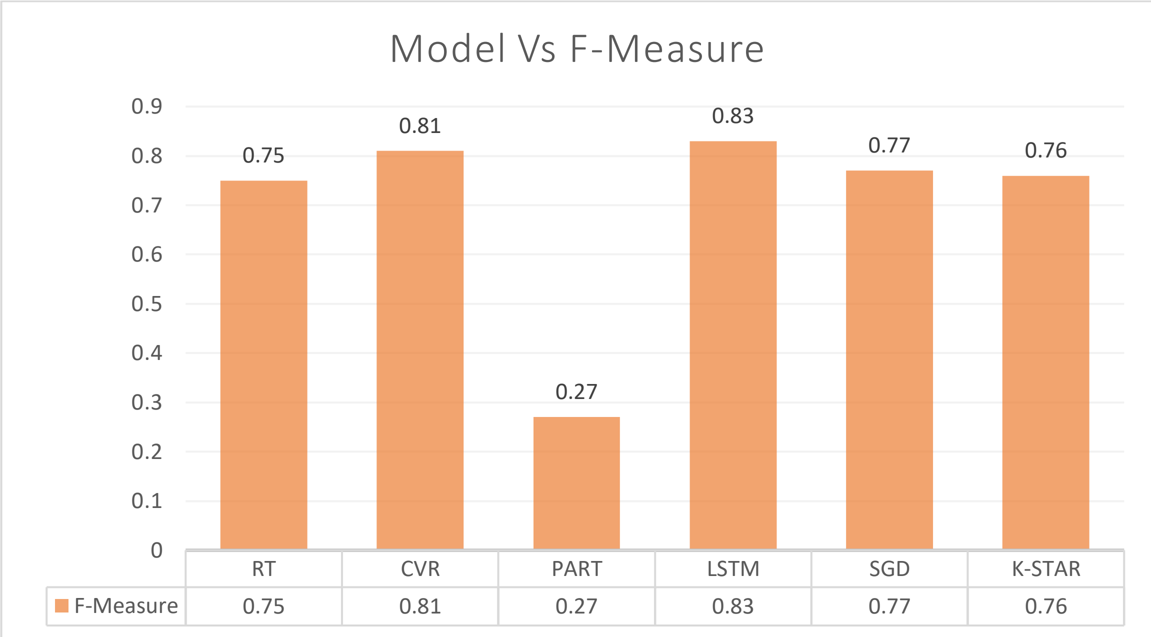

Figure 9: Model Vs F-Measure

Figure 9 shows the F-measure scores for the RT, CVR, PART, LSTM, SGD, and K-Star models. LSTM had the highest F-measure at 0.83, followed by CVR at 0.81. PART had the lowest F-measure at 0.27. RT scored 0.75, K-Star scored 0.76, and SGD scored 0.77.

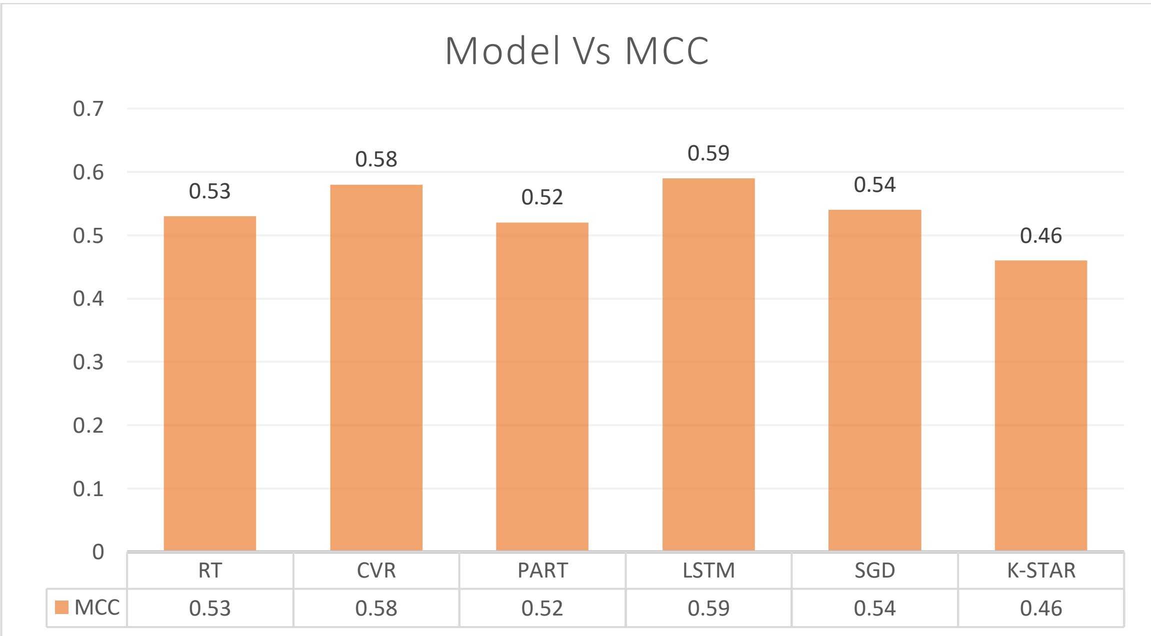

Figure 10: Model Vs MCC

Figure 10 displays the MCC (Matthews Correlation Coefficient) values for the RT, CVR, PART, LSTM, SGD, and K-Star models. LSTM achieved the highest MCC at 0.59, followed by CVR at 0.58. K-Star had the lowest MCC at 0.46. RT had an MCC of 0.53, PART had an MCC of 0.52, and SGD had an MCC of 0.54.

5. Conclusion

This study introduced a new AI-based system for detecting spam reviews, which is a growing problem on online platforms. Using advanced Deep Learning and Natural Language Processing, the system performed very well at identifying spam, especially AI-generated reviews. The research shows the need for sophisticated AI solutions to protect online platforms. However, understanding complex models, especially ensembles, can be difficult. The study addresses these challenges and explains how potential biases in expert validation were minimized. Future work will focus on making the models easier to understand and continuing to work with experts to improve them. Ethical considerations are crucial to our interpretable spam detection approach. We examine the ethical implications of revealing how the model makes decisions, including privacy concerns and the potential for misuse. We also discuss the broader impact of interpretable AI on e-commerce integrity and consumer trust. Our framework not only outperformed traditional machine learning models but also provides a scalable and accurate way to detect AI-generated spam, significantly improving the reliability of online reviews. This work sets the stage for future spam detection development, offering a powerful tool that can be easily integrated into various online environments to protect consumer trust and improve decision-making. Based on its performance, we recommend the LSTM model.

References

[1] Ahmed, H.; Traore, I.; Saad, S. Detecting opinion spams and fake news using text classification. Secur. Priv. 2018, 1, e9. [CrossRef]

[2] Crawford, M.; Khoshgoftaar, T.M.; Prusa, J.D.; Richter, A.N.; Al Najada, H. Survey of review spam detection using machine learning techniques. J. Big Data 2015, 2, 23. [CrossRef]

[3] Lee, K.D.; Han, K.; Myaeng, S.H. Capturing word choice patterns with LDA for fake review detection in sentiment analysis. In Proceedings of the International Conference on Web Intelligence, Mining and Semantics, Nimes, France, 13–15 June 2016; pp. 1–7.

[4] Floridi, L.; Chiriatti, M. GPT-3: Its nature, scope, limits, and consequences. Minds Mach. 2020, 30, 681–694. [CrossRef]

[5] Ott, M.; Choi, Y.; Cardie, C.; Hancock, J.T. Finding deceptive opinion spam by any stretch of the imagination. arXiv 2011,arXiv:1107.4557.

[6] Petrescu, M.; O'Leary, K.; Goldring, D.; Mrad, S.B. Incentivized reviews: Promising the moon for a few stars. J. Retail. Consum.Serv. 2018, 41, 288–295. [CrossRef]

[7] Costa, A.; Guerreiro, J.; Moro, S.; Henriques, R. Unfolding the characteristics of incentivized online reviews. J. Retail. Consum.Serv. 2019, 47, 272–281. [CrossRef]

[8] DePaulo, B.M.; Kashy, D.A.; Kirkendol, S.E.; Wyer, M.M.; Epstein, J.A. Lying in everyday life. J. Personal. Soc. Psychol. 1996,70, 979. [CrossRef]

[9] Jindal, N.; Liu, B. Opinion spam and analysis. In Proceedings of the 2008 International Conference on Web Search and DataMining, Palo Alto, CA, USA, 11–12 February 2008; pp. 219–230.

[10] Filieri, R. What makes an online consumer review trustworthy? Ann. Tour. Res. 2016, 58, 46–64. [CrossRef]
[11] Plotkina, D.; Munzel, A.; Pallud, J. Illusions of truth—Experimental insights into human and algorithmic detections of fake online reviews. J. Bus. Res. 2020, 109, 511–523. [CrossRef]
[12] Sun, H.; Morales, A.; Yan, X. Synthetic review spamming and defense. In Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Chicago, IL, USA, 11–14 August 2013; pp. 1088–1096.
[13] Cardoso, E.F.; Silva, R.M.; Almeida, T.A. Towards automatic filtering of fake reviews. Neurocomputing 2018, 309, 106–116.[CrossRef]
[14] Mihalcea, R.; Strapparava, C. The lie detector: Explorations in the automatic recognition of deceptive language. In Proceedings of the ACL-IJCNLP 2009 Conference Short Papers; Association for Computational Linguistics: Singapore, 2009; pp. 309–312.
[15] Mukherjee, A.; Venkataraman, V.; Liu, B.; Glance, N. What yelp fake review filter might be doing? Proc. Int. AAAI Conf. Web Soc.Media 2013, 7, 409–418. [CrossRef]
[16] Sandulescu, V.; Ester, M. Detecting singleton review spammers using semantic similarity. In Proceedings of the 24th International Conference on World WideWeb, Florence, Italy, 18–22 May 2015; pp. 971–976. Computers 2024, 13, 264 27 of 27
[17] Arvind Mewada, Rupesh Kumar Dewang, Research on false review detection Methods: A state-of-the-art review, Journal of King Saud University - Computer and Information Sciences, Volume 34, Issue 9
[18] https://www.kaggle.com/datasets/naveedhn/amazon-produc-review-spam-and-non-spam
[19] Yuanyuan Wu, Eric W.T. Ngai, Pengkun Wu, Chong Wu,Fake online reviews: Literature review, synthesis, and directions for future research, Decision Support Systems,Volume,132,2020,113280,ISSN0167-9236
[20] Ao Li, Zhou Qin, Runshi Liu, Yiqun Yang, and Dong Li. 2019. Spam Review Detection with Graph Convolutional Networks. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM '19). Association for Computing Machinery, New York, NY, USA, 2703–2711. https://doi.org/10.1145/3357384.3357820