Analysis and Detection of Deceptive Product Reviews for E-Commerce Platforms Using Machine learning and Text Mining: A Systematic Literature Review

[1]Akshay R Ghooli, [1]Abhinav Kulkarni, [1]Karna H Gowda, [1]Aishwarya T R, [1]Shruthi B S *Information and Science and Engineering, Malnad College of Engineering* Hassan-573202, India

Email id:akshghooli@gmail.com, abhinav.a.kulkarni123@gmail.com, karnahgowda@gmail.com, aishgowda2622@gmail.com

***Abstract*-Online reviews significantly influence modern business and commerce. Consumer decisionmaking processes for online purchases are significantly influenced by user feedback. However, this has led to the rise of fraudulent practices, where individuals or groups manipulate product reviews to serve their own interests. Posting fake reviews is a common deceptive practice used to unfairly promote products. This issue is especially prevalent in e-commerce platforms, where fake reviews can distort customer perceptions and impact purchasing decisions. To tackle this issue, we introduce a system that utilizes machine learning algorithms to identify and flag fraudulent reviews on e-commerce platforms. In our approach, we develop a novel e-commerce application where customers can register, browse products, make purchases, and share their feedback. The system utilizes data science techniques to analyze customer reviews and detect potential instances of fake reviews. This solution aims to enhance trust in online reviews and improve the overall shopping experience by ensuring the authenticity of customer feedback.**

Keywords— *Fake Review Detection, E-commerce Platforms, Machine Learning Algorithms, Naïve Bayes Classifier, Semi-Supervised Learning, Supervised Learning, Results.*

I. INTRODUCTION

Technological advancements are rapidly transforming the way tasks are accomplished, with older technologies being replaced by more efficient and sophisticated ones. One such innovation is the rise of online marketplaces, where people can shop and make reservations with ease. Today, online reviews play a pivotal role in shaping the reputation of businesses and influencing consumer decisions. Online reviews serve as a potent tool for advertising and promoting products and services. However, the proliferation of fabricated reviews presents a significant challenge. Such reviews are often used to promote products dishonestly or to tarnish the reputation of competitors, ultimately misleading consumers. In response to this problem, researchers have examined diverse methods for detecting fake reviews. These methods can broadly be categorized into content-based approaches, which analyze the textual content of reviews, and behaviour-based approaches, which examine factors such as location, IP address, and reviewer activity. Most existing systems rely on supervised classification techniques. Furthermore, some researchers have utilized semi-supervised models to overcome the limitations posed by the scarcity of labeled data. In our work, we adopt both supervised and semisupervised methods for fake review detection. algorithm for semisupervised
learning and a Statistical Naive Bayes classifier to enhance classification performance. Our approach focuses on features like word frequency, sentiment polarity, and review length to ensure accurate detection of fake reviews.

II. RELATED WORKS

This section summarizes and compares existing studies on fake review detection using ML models. The studies reviewed focus on the algorithms, datasets, and
methodologies used to classify and predict fake reviews

A. [1]Fake Reviews Detection

- **Authors:** Rami Mohawesh, Shuxiang Xu, Son N. Tran, Robert Ollington, Matthew Springer, Yaser Jararweh, and Sumbal Maqsood
- **Year:** 2021

Description:
This study explores the problem of fake reviews in ecommerce, highlighting their detrimental effects on consumer confidence and the ethical conduct of businesses. It emphasizes how manipulated reviews can skew product ratings and mislead purchasing decisions, leading to both economic and social consequences. The research underscores the critical role of trustworthy reviews in sustaining a fair and reliable online marketplace.

Methodology:
The authors conduct an analysis of various fake review detection techniques, including linguistic analysis of content, behavioral analysis of reviewer activity, and graph-based methods that leverage network structures. They explore machine learning approaches, including supervised and unsupervised models, as well as hybrid systems that combine multiple strategies. Additionally, the study reviews the use of publicly available datasets for training and evaluating detection models.

Limitations:
The research highlights several challenges, including the limited availability of high-quality labeled datasets and the difficulty in generalizing detection models across different platforms. Distinguishing between genuine strong opinions and fraudulent reviews proves to be a complex task. The paper also examines ethical considerations related to privacy and the potential for misuse of detection technologies.

Key Insights:

Fake reviews significantly influence consumer decisionmaking and business performance. The study reveals that hybrid detection methods provide enhanced accuracy in identifying deceptive reviews. It emphasizes the importance of collaboration among platforms, researchers, and policymakers to effectively combat this issue. Future research should prioritize the development of real-time detection systems and the formulation of strategies to counter adversarial techniques.

B. [2]Identifying Groups of Fake Reviewers Using a Semi-Supervised Approach

- **Authors:** Punit Rathore, Jayesh Soni, Nagarajan Prabakar, Marimuthu Palaniswami, and Paolo Santi
- **Year:** 2021

Description:

The paper explores the influence of coordinated fake reviewer groups on product sentiment in online marketplaces. These groups manipulate product ratings by creating multiple fake accounts and posting misleading reviews. The research underscores the need to detect such groups to maintain consumer confidence and ensure the integrity of digital marketplaces

Methodology:

The study proposes a graph-based approach where interactions between reviewers and products are modeled as a graph. Using the DeepWalk algorithm, embeddings are generated to represent these interactions. A semi-supervised clustering process is then applied to identify suspicious groups exhibiting behaviors like synchronized review timing and repetitive patterns.

Limitations:

The effectiveness of the framework depends on the availability and accuracy of graph data, which may not always be reliable. The approach also faces challenges in scaling to platforms with extensive datasets. Furthermore, the method may occasionally misclassify genuine reviewer groups with similar behavior as suspicious.

Key Insights:

Coordinated reviewer groups have a greater impact on product sentiment compared to individual fake reviewers. The combination of graph-based techniques and semisupervised learning has shown promising results. Future efforts should address scalability issues and aim to optimize computational efficiency.

C.[3]Machine Learning Approaches for Fake Reviews Detection

- **Authors:** Mohammed Ennaouri and Ahmed Zellour.
- **Year:** 2023 **Description:**

Fake reviews, or opinion spam, pose a challenge online. They mislead consumers and harm businesses. To address this challenge, researchers employ machine learning (ML) techniques. A common approach involves supervised learning using labeled datasets. Deep learning methods like CNNs and LSTMs also show promise. Labeled datasets from sources like Yelp support model evaluation. Effective detection frameworks can help mitigate the issue.

Methodology:

To prepare the dataset for analysis, the proposed system includes a preprocessing step that involves tokenization, removal of stop words, stemming, and part-of-speech tagging. Count vectorization is applied to identify term frequencies, followed by label encoding to convert text data into machine-readable formats. The dataset is partitioned into training and testing sets, maintaining an 80-20 ratio. Subsequently, SVM and Naive Bayes classifiers are trained on the processed data. A Flask-based web interface is developed for prediction, allowing users to input reviews and obtain classification results

Limitations:

The study's findings may have limited generalizability to other domains, as it relies solely on a dataset of Amazon reviews. The reliance on basic feature extraction methods like count vectorization limits the model's ability to capture complex textual patterns. SVM outperforms Naive Bayes in accuracy, but other advanced algorithms are not explored. Additional features, such as ratings and verified purchases, could enhance the robustness of the model.

Key Insights:

This paper reviews machine learning techniques for fake review detection, emphasizing supervised methods like SVM and Random Forests for high accuracy. Deep learning models, including Convolutional Neural Networks (CNNs) and Long Short-Term Memory networks (LSTMs), are emphasized for their capacity to effectively process intricate text patterns. Limited labeled datasets, reliance on Yelp and OTT data, and domain-specific focus are noted as major challenges. Feature extraction techniques like sentiment analysis and part-of-speech tagging are crucial for improving model performance. Future research is needed for real-time detection systems and platform-independent solutions.

D.[4]Detection of Fake Reviews on Products Using Machine Learning

- **Authors:** M. Narayana Royal, Rajula Pavan Kalyan Reddy, Gokina Sri Sangathya,B. Sai Madesh Pretam, Jayakumar Kaliappan, and C. Suganthan
- **Year:** 2023 **Description:**

The paper focuses on detecting fake product reviews using supervised machine learning algorithms. Reviews from Amazon are utilized, with an equal distribution of fake and genuine reviews to ensure balanced training. Preprocessing steps include tokenization, stemming, stopword removal, and part-of-speech tagging. Features are extracted using count vectorization, creating a structured vocabulary. SVM and Naive Bayes classifiers are implemented, with SVM demonstrating higher accuracy. The system includes a Python Flask-based web interface for real-time review classification. **Methodology:**

The proposed system incorporates a preprocessing stage that involves tokenization, removal of stop words, stemming, and part-of-speech tagging. Count vectorization is applied to identify term frequencies, followed by label encoding to convert text data into machine-readable formats. The dataset is partitioned into training and testing sets, maintaining an 8020 ratio. Subsequently, SVM and Naive Bayes classifiers are trained on the processed data. A Flask-based web interface is developed for prediction, allowing users to input reviews and obtain classification results.

Limitations:

The study's findings may have limited generalizability beyond Amazon reviews due to the use of a domain-specific dataset. Furthermore, the reliance on basic feature extraction techniques such as count vectorization may restrict the model's capacity to capture intricate textual patterns. SVM outperforms Naive Bayes in accuracy, but other advanced algorithms are not explored. Additional features, such as ratings and verified purchases, could enhance the robustness of the model. **Key Insights:**

The study utilizes SVM and Naive Bayes classifiers to identify fake reviews in Amazon data, achieving an accuracy of 80.1% with the SVM model. Preprocessing involves tokenization, stemming, and count vectorization for feature extraction. A Python Flask-based web interface enables real-time review classification. Challenges include limited dataset diversity and basic feature engineering. Future work should focus on multilingual datasets and advanced algorithms to improve scalability and robustness.

E. [5]Detecting Fake Reviews in E-Commerce Using Machine Learning: A Survey

- **Authors:** Maysam Jalal Abd and Mohsin Hasan Hussein**Year:** 2024 **Description:**

This paper examines the critical problem of fraudulent reviews in e-commerce and their influence on consumer trust and business outcomes. It provides an overview of machine learning strategies used to address this issue, focusing on supervised, unsupervised, and semi-supervised techniques. The research evaluates these methods based on their ability to identify fake reviews effectively, considering different features such as reviewer behavior and content attributes. By synthesizing insights from multiple studies, the paper offers valuable guidance for improving fraud detection in online marketplaces.

Methodology: The study explores three main strategies for identifying fake reviews: content-focused, reviewer-focused, and productfocused approaches. The study evaluates supervised learning algorithms, such as Support Vector Machines (SVM) and Decision Trees, for their effectiveness in classifying labeled datasets. Unsupervised techniques analyze unlabeled data using clustering and feature selection. Semi-supervised methods combine labeled and unlabeled data for enhanced detection. Emphasis is placed on factors such as writing style, user activity, and product-related metadata in improving model effectiveness.

Limitations:

Key challenges include a lack of diverse and standardized datasets, which limits the ability of models to generalize across different platforms and domains. The study highlights difficulties in processing the rapid influx of reviews, emphasizing the need for scalable solutions. Real-time detection remains a challenge, as most models are optimized for offline analysis. Furthermore, the reliance on feature quality and engineering significantly impacts detection accuracy, necessitating further refinement of input data.

Key Insights:

The research highlights that hybrid models, such as CNNRNN, can achieve near-perfect detection accuracy (up to 99%), with ensemble approaches and SVM performing robustly. Deep learning frameworks like Bi-LSTM excel in capturing complex patterns in review text. Unsupervised learning offers a practical solution for scenarios lacking labeled data but generally shows lower accuracy. The study underscores the necessity of real-time, adaptable systems to combat the growing challenge of fake reviews across diverse e-commerce platforms.

F. [6]Revisiting Semi-Supervised Learning for Online Deceptive Review Detection

- **Authors:** Jitendra Kumar Rout, Anmol Dalmia, Kim-Kwang Raymond Choo, Sambit Bakshi, and Sanjay Kumar Jena
- **Year:** 2017 **Description:**

This paper explores the use of semi-supervised learning techniques to detect deceptive reviews in online platforms, emphasizing the challenges posed by the dynamic and often ambiguous nature of such content. By focusing on methods that combine labeled and unlabeled data, the study enhances the ability to identify fake reviews efficiently. Using an enriched feature set including sentiment polarity, linguistic characteristics, and word count metrics, the research aims to improve detection accuracy compared to traditional supervised approaches. The findings contribute to understanding how semi-supervised techniques can address limitations in labeled data availability while maintaining robustness.

Methodology:

The authors evaluated four key semi-supervised learning algorithms: Co-Training, Expectation Maximization, Label Propagation, and Positive Unlabeled (PU) Learning. These methods utilize a mix of labeled and unlabeled datasets, incorporating features such as sentiment analysis, linguistic patterns, and part-of-speech tagging. Experiments were conducted using the "gold standard" Ott dataset of hotel reviews, partitioned into labeled and unlabeled subsets to simulate real-world scenarios. Performance metrics like accuracy, precision, and F-score were assessed to determine the effectiveness of each approach.

Limitations:

The study notes the inherent challenges of semi-supervised learning, including potential inaccuracies in the automated labeling process and dependency on feature quality. The dataset size and diversity may also limit generalization to other domains or languages. Furthermore, reliance on textual data excludes other contextual signals, such as multimedia content, that could improve detection reliability. Real-time applicability of these methods in live systems remains an area for future exploration.

Key Insights:

The research demonstrates that semi-supervised learning can significantly improve detection of deceptive reviews, with PU Learning achieving the highest F-score of 0.837. Co- Training and Expectation Maximization also performed robustly, offering effective solutions for scenarios with limited labeled data. The incorporation of enriched feature vectors, including linguistic and sentimental metrics, further enhances classification accuracy. The study highlights the importance of leveraging unlabeled data in deceptive review detection, suggesting future directions for incorporating diverse datasets and real-time application strategies.

G. [7] The Effect of Fake Reviews on e-Commerce During and After Covid-19 Pandemic: SKL-Based Fake Reviews Detection.Authors: Hina Tufail, M. Usman Ashraf, Khalid Alsubhi, and Hani Moaiteq Aljahdali.

- **Year:** 2022.

Description:

The research examines the prevalence and impact of fake reviews on e-commerce, particularly during the COVID-19 pandemic. It highlights how fraudulent reviews, driven by intense competition among businesses, influence consumer decisions and brand reputations. To combat this, the authors propose a machine learning model to detect fake reviews based on linguistic and behavioral features. The study achieves high accuracy using datasets from Yelp and

TripAdvisor, demonstrating the effectiveness of its approach. It sheds light on the growing importance of ensuring authenticity in online reviews.

Methodology:

The proposed methodology employs machine learning techniques, including Support Vector Machine (SVM), KNearest Neighbor (KNN), and Logistic Regression (LR). Features such as word frequency, pronouns, sentiment polarity, and bigram analysis were used to train the model. The preprocessing phase involved data filtering and feature selection, enhancing classification accuracy. The model was tested on datasets from Yelp and TripAdvisor, achieving accuracies of 95% and 89.03%, respectively. These results demonstrate the effectiveness of feature engineering combined with robust classification algorithms.

Limitations:

Despite its effectiveness, the research has several limitations. The datasets, primarily from Yelp and TripAdvisor, may not represent the full spectrum of fake reviews across different platforms and domains. Additionally, the evolving strategies used by fraudsters to create convincing fake reviews pose a challenge for long-term adaptability. The reliance on labeled datasets limits the model's generalizability to unlabeled or semi-supervised contexts. Moreover, while achieving high accuracy, the model's performance may vary with different datasets and user behavior patterns. Addressing these challenges requires future refinement and testing.

Key Insights :

The study reveals that fake reviews significantly impact consumer trust and business performance, especially in the competitive e-commerce landscape. The SKL-based model, combining SVM, KNN, and LR, proves highly effective in detecting fraudulent reviews using linguistic and sentimentbased features. SVM emerged as the best- performing classifier, demonstrating superior precision and recall. The research highlights the critical role of preprocessing and feature selection in improving detection accuracy. These insights emphasize the need for advanced detection mechanisms to safeguard the integrity of online marketplaces.

H.[8] Fake Review Detection Based on Multiple Feature Fusion and Rolling Collaborative Training

- **Authors:**Jingdong Wang, Haitao Kan, Fanqi Meng, Qizi Mu, Genhua Shi, and Xixi Xiao.
- **Year:** 2020.

Description:

This study investigates the pervasive issue of fake reviews in e-commerce, which undermine consumer trust and distort market dynamics. By integrating text features, sentiment analysis, and user behavior, the authors developed a multifeature fusion model. The research introduces rolling collaborative training to iteratively improve classifiers by leveraging labeled and unlabeled data. Experiments conducted on Yelp datasets revealed improved accuracy over baseline methods. The findings emphasize the need for adaptive and comprehensive systems in combating fraudulent reviews.

Methodology:

The methodology combines multiple feature fusion and rolling collaborative training to detect fake reviews with enhanced precision. Text features are extracted using Doc2Vec, and sentiment polarity is quantified to capture emotional consistency. Reviewer behaviors, such as abnormal posting patterns and rating deviations, are incorporated into the analysis. Classifiers like Support Vector Machine (SVM) and Random Forest (RF) are trained iteratively, dynamically updating with high-confidence labels. Statistical validation, including the Friedman test, supports the effectiveness of this adaptive model.

Limitations:

The study's limitations include restricted access to comprehensive e-commerce data, limiting insights into user behavior. The use of manually labeled datasets introduces potential bias, as they may not reflect the complexity of realworld reviews. While rolling collaborative training reduces manual labeling, the approach is computationally intensive and requires significant resources. Semi-supervised learning depends heavily on the quality of initial labeled data, constraining its potential for accuracy improvements. These factors highlight areas for further research and refinement **Key Insights :**

The research reveals that effective fake review detection requires combining linguistic, sentiment, and behavioral features. The rolling collaborative training approach enables classifiers to adapt dynamically to evolving review manipulation tactics. Support Vector Machine (SVM) and Random Forest (RF) emerged as the most effective classifiers for the task. The proposed model demonstrated significant accuracy improvements over traditional and deep learning methods. The study highlights the importance of integrating time-sensitive and multi-dimensional features to enhance detection reliability.

III. METHODOLOGY

The methodology focuses on detecting fake reviews in ecommerce platforms using machine learning techniques. Data comprising labeled genuine and fake reviews is collected, and key features such as word frequency counts, sentiment polarity, and review length are extracted to distinguish authentic reviews. Essential features are selected to reduce overfitting, avoiding derived features like bigrams or trigrams. The Naive Bayes algorithm is employed for supervised learning, achieving an accuracy of 86.32%, while the Expectation-Maximization algorithm is utilized for semisupervised learning with an accuracy of 85.21%. The system is implemented in a real-time e-commerce environment where users register, browse products, and post reviews, with the model predicting review authenticity. Performance is evaluated using metrics such as accuracy, precision, recall, and F1-score. Developed using Visual Studio for the frontend, SQL Server for the back-end, and languages like C#, the system demonstrates improved accuracy over existing approaches, ensuring a reliable platform for users and mitigating the impact of fake reviews.

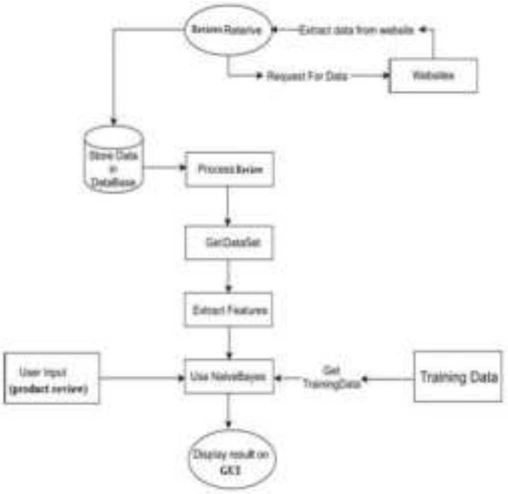

Fig 1. Methodology

IV. ALGORITHMS USED

4.1.1 Naive Bayes Classifier:

Naive Bayes is a probabilistic algorithm based on Bayes' theorem, assuming that features are conditionally independent given the class label. In this study, it is applied as a supervised learning method to classify reviews into fake or genuine. Key features like word frequency count, sentiment polarity, and review length are used for classification. The algorithm is computationally efficient and well-suited for text classification tasks, achieving an accuracy of 86.32%. Its simplicity, effectiveness in handling highdimensional data, and robustness to noise make it an excellent choice for detecting fake reviews in e-commerce platforms.

Likelihood Class Prior Probability

$$P(c \mid x) = \frac{P(x \mid c)P(c)}{P(x)}$$

Posterior Probability Predictor Prior Probability

$$P(c \mid X) = P(x_1 \mid c) \times P(x_2 \mid c) \times \cdots \times P(x_n \mid c) \times P(c)$$

Fig 2. Formula

4.1.2. Expectation-Maximization Algorithm:

The Expectation-Maximization (EM) algorithm is used for semi-supervised learning, especially useful in scenarios with limited labeled data. It alternates between the expectation step (assigning probabilities to unlabeled data based on the current model) and the maximization step (optimizing model parameters using both labeled and pseudo-labeled data). In this study, the EM algorithm enhances the dataset with highconfidence pseudo-labels, improving the accuracy of the classification model to 85.21%. This iterative process is particularly advantageous in real-world cases where obtaining labeled data is challenging, making it a valuable tool for improving fake review detection systems.

Expectation Maximization (EM) Algorithm

log of expectation of P(x|z)

Goal: $\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \log\left(\sum_{z} p(\mathbf{x}, \mathbf{z} \mid \theta)\right)$ $\quad f(\mathrm{E}[X]) \geq \mathrm{E}[f(X)]$

1. E-step: compute

expectation of log of P(x|z)

$$\mathrm{E}_{z \mid x, \theta^{(t)}}\left[\log(p(\mathbf{x}, \mathbf{z} \mid \theta))\right] = \sum_{z} \log(p(\mathbf{x}, \mathbf{z} \mid \theta)) p(\mathbf{z} \mid \mathbf{x}, \theta^{(t)})$$

2. M-step: solve

$$\theta^{(t+1)} = \underset{\theta}{\operatorname{argmax}} \sum_{z} \log(p(\mathbf{x}, \mathbf{z} \mid \theta)) p(\mathbf{z} \mid \mathbf{x}, \theta^{(t)})$$

Fig 3.EM Algorithm

V. RESULTS

A recent study demonstrated the effectiveness of a proposed system in detecting fake product reviews on e-commerce platforms using machine learning techniques, achieving high accuracy rates. Notably, the Naive Bayes classifier achieved an accuracy of 86.32% under supervised learning, while the Expectation-Maximization algorithm combined with Naïve Bayes reached an accuracy of 85.21% under semi-supervised learning, outperforming existing methods. The system showed improved accuracy compared to previous approaches, enabling faster and more reliable detection of fake reviews. The results highlighted the importance of carefully chosen features, including word frequency count, sentiment polarity, and review length, in enhancing model performance. Overall, the study validated the feasibility of implementing such models in real-time e-commerce applications, providing a more genuine and trustworthy platform for customers.

VI. CONCLUSION

Detecting fake reviews is an essential step toward building trust in e-commerce platforms. The proposed system uses machine learning, specifically the Naive Bayes algorithm, to identify fake reviews with high accuracy, achieving 86.32% in supervised classification. By focusing on key features like word frequency, sentiment polarity, and review length, the system offers reliable results while avoiding common pitfalls like overfitting. Although challenges such as limited labeled data remain, this approach provides a strong foundation for further improvements and contributes to a more transparent and trustworthy online shopping experience.

References

[1] Paper Title: Fake Reviews Detection Authors: Rami Mohawesh, Shuxiang Xu, Son N. Tran, Robert Ollington, Matthew Springer, Yaser Jararweh, and Sumbal Maqsood ,Year: 2021

[2] Paper Title: Stress Detection in College Students Using a Machine Learning Algorithm, Authors: Ms. Ancy Paul, Ms. Resija P R, Year: 2024

[3] Paper Title: Machine Learning Approaches for Fake Reviews Detection, Authors: Mohammed Ennaouri and Ahmed Zellour,Year: 2023

[4] Paper Title: Detection of Fake Reviews on Products Using Machine Learning, Authors: M. Narayana Royal, Rajula Pavan Kalyan Reddy, Gokina Sri Sangathya,B. Sai Madesh Pretam, Jayakumar Kaliappan, and C. Suganthan, Year: 2023

[5] Chengai Sun, Qiaolin Du and Gang Tian, "Exploiting Product Related Review Features for Fake Review Detection," Mathematical Problems in Engineering, 2016.

[6] A. Heydari, M. A. Tavakoli, N. Salim, and Z. Heydari, "Detection of review spam: a survey", Expert Systems with Applications, vol. 42, no. 7, pp. 3634–3642, 2015.

[7] M. Ott, Y. Choi, C. Cardie, and J. T. Hancock, "Finding deceptive opinion spam by any stretch of the imagination," in Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT), vol. 1, pp. 309–319, Association for Computational Linguistics, Portland, Ore, USA, June 2011.

[8] J. W. Pennebaker, M. E. Francis, and R. J. Booth, "Linguistic Inquiry and Word Count: Liwc," vol. 71, 2001.

[9] S. Feng, R. Banerjee, and Y. Choi, "Syntactic stylometry for deception detection," in Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers, Vol. 2, 2012.

[10] J. Li, M. Ott, C. Cardie, and E. Hovy, "Towards a general rule for identifying deceptive opinion spam," in Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL), 2014.

[11] E. P. Lim, V.-A. Nguyen, N. Jindal, B. Liu, and H. W. Lauw, "Detecting product review spammers using rating behaviors," in Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM), 2010.

[12] J. K. Rout, A. Dalmia, and K.-K. R. Choo, "Revisiting semisupervised learning for online deceptive review detection," IEEE Access, Vol. 5, pp. 1319–1327, 2017.