ISSN: 2319-345X

IJMRBS

International Journal of Management Research and Business Strategy

www.ijmrbs.org

E-mail
editor@ijmrbs.org
editor.ijmrbs@gmail.com

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

DETECTING AND CHARACTERIZING EXTREMIST REVIEWER GROUPS IN ONLINE PRODUCT REVIEWS

Raja Rajeswari kalidindi, Associate professor,
Department of MCA
rajeswari.kalidindi29@gmail.com
B V Raju College, Bhimavaram

Gundumogula Rishi Saravan (2285351041)
Department of MCA
grsaravan6@gmail.com
B V Raju College, Bhimavaram

ABSTRACT

Online marketplaces often witness opinion spam in the form of reviews. People are often hired to target specific brands for promoting or impeding them by writing highly positive or negative reviews. This often is done collectively in groups. Although some previous studies attempted to identify and analyze such opinion spam groups, little has been explored to spot those groups who target a brand as a whole, instead of just products. In this article, we collected the reviews from the Amazon product review site and manually labeled a set of 923 candidate reviewer groups. The groups are extracted using frequent itemset mining over brand similarities such that users are clustered together if they have mutually reviewed (products of) a lot of brands. We hypothesize that the nature of the reviewer groups is dependent on eight features specific to a (group, brand) pair. We develop a feature-based supervised model to classify candidate groups as extremist entities. We run multiple classifiers for the task of classifying a group based on the reviews written by the users of that group to determine whether the group shows signs of extremity. A three-layer perceptron-based classifier turns out to be the best classifier. We further study behaviors of such groups in detail to understand the dynamics of brand-level opinion fraud better. These behaviors include consistency in ratings, review sentiment, verified purchase, review dates, and helpful votes received on reviews. Surprisingly, we observe that there are a lot of verified reviewers showing extreme sentiment, which, on further investigation, leads to ways to circumvent the existing mechanisms in place to prevent unofficial incentives on Amazon.

Keywords: Opinion spam, Online marketplaces, Review fraud, Extremist entities, Brand-level analysis, Frequent itemset mining, Reviewer behavior.

INTRODUCTION

The proliferation of online marketplaces has facilitated the exchange of goods and services on a global scale, accompanied by the emergence of user-generated content, particularly in the form of product reviews [1]. However, amidst the vast array of user-generated reviews, online marketplaces often fall prey to opinion spam, where individuals are hired to manipulate the perception of products or brands through artificially positive or negative reviews [2]. These deceptive practices, orchestrated either to boost sales or tarnish competitors, are frequently carried out collectively by groups of individuals, thereby amplifying their impact [3]. While previous research has made strides in identifying and analyzing such opinion spam groups, there remains a significant gap in understanding and detecting groups that target brands holistically, rather than individual products [4].

This article aims to address this gap by focusing on the detection and characterization of extremist reviewer groups in online product reviews, with a specific emphasis on their behavior towards brands as a whole [5]. To facilitate this investigation, we collected a comprehensive dataset of reviews from the Amazon product review site, representing a diverse range of products and brands [6]. Through manual labeling, we identified and annotated a set of 923 candidate reviewer groups, each exhibiting patterns indicative of coordinated review activity [7]. These groups were extracted

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

using frequent itemset mining techniques applied over brand similarities, allowing users to be clustered together based on their shared review history across multiple brands [8]. Drawing on insights from prior research and theoretical considerations, we developed a feature-based supervised model to classify candidate reviewer groups as extremist entities [9]. We hypothesized that the nature of reviewer groups could be characterized by eight specific features pertinent to the (group, brand) pair, including aspects such as review sentiment, consistency in ratings, verified purchase status, review dates, and helpful votes received [10]. Leveraging machine learning techniques, we trained multiple classifiers on the extracted features to discern the presence of extremist behavior within reviewer groups [11]. Notably, our experiments revealed that a three-layer perceptron-based classifier emerged as the most effective in identifying extremist groups, demonstrating superior performance compared to alternative approaches [12].

In addition to classification, we conducted an in-depth analysis of the behaviors exhibited by extremist reviewer groups to gain a deeper understanding of brand-level opinion fraud dynamics [13]. This analysis encompassed various dimensions, including the consistency of ratings and review sentiment, the prevalence of verified purchases, the temporal distribution of reviews, and the reception of helpful votes from other users [14]. Surprisingly, our investigation unveiled a significant number of verified reviewers exhibiting extreme sentiment in their reviews, raising concerns about the efficacy of existing mechanisms to prevent unofficial incentives and fraudulent activity on Amazon [15]. In summary, this article contributes to the literature by shedding light on the phenomenon of extremist reviewer groups in online product reviews and proposing a novel methodology for their detection and characterization. By leveraging machine learning techniques and analyzing behavioral patterns, our research offers valuable insights into the dynamics of opinion spam and highlights potential avenues for mitigating fraudulent activities in online marketplaces.

LITERATURE SURVEY

The phenomenon of opinion spam in online marketplaces has garnered significant attention due to its potential to influence consumer decisions and distort the perceived reputation of products and brands. Opinion spam, often manifesting as artificially positive or negative reviews, is frequently orchestrated by individuals hired to promote or disparage specific brands for financial gain. Notably, this deceptive practice is not limited to isolated actions but often occurs collectively in groups, amplifying its impact and making it more challenging to detect and combat. While previous studies have made strides in identifying and analyzing opinion spam groups, particularly those targeting individual products, there remains a notable gap in research concerning groups that systematically target brands as a whole. This gap underscores the need for a comprehensive understanding of the dynamics underlying extremist reviewer groups and their strategies for manipulating online product reviews.

To address this gap, this article undertakes a rigorous investigation into the detection and characterization of extremist reviewer groups operating within online product reviews. Leveraging data collected from the Amazon product review site, we manually labeled a substantial set of candidate reviewer groups, comprising 923 distinct entities. These groups were extracted using advanced data mining techniques, specifically frequent itemset mining, applied over brand similarities. By clustering users based on their mutual review activity across multiple brands, we sought to capture the collective behavior of reviewer groups targeting brands holistically. Furthermore, we hypothesized that the nature of these reviewer groups is influenced by a set of eight specific features inherent to the relationship between the group and the brand under scrutiny. These features encompass various aspects such as review sentiment, consistency in ratings, verified purchase status, review dates, and the reception of helpful votes, among others. By developing a feature-based supervised model, we aimed to classify candidate groups as extremist entities based on these distinctive characteristics, thereby enabling effective detection and identification of suspicious review behavior.

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

To evaluate the efficacy of our classification approach, we employed multiple classifiers to assess the extremity of reviewer groups based on the reviews authored by their members. Among the classifiers evaluated, a three-layer perceptron-based model emerged as the most effective in accurately discerning extremist behavior within reviewer groups. This finding underscores the potential of machine learning techniques in identifying and mitigating opinion spam in online product reviews. Furthermore, our analysis extended beyond classification to delve into the behavioral patterns exhibited by extremist reviewer groups, shedding light on their tactics and strategies for manipulating review platforms. In particular, we scrutinized behaviors such as consistency in ratings, review sentiment, verified purchase status, review dates, and the reception of helpful votes. Notably, our investigation revealed a concerning trend wherein a substantial number of verified reviewers exhibited extreme sentiment in their reviews, raising questions about the effectiveness of existing mechanisms for preventing fraudulent activities and unofficial incentives on platforms like Amazon.

Overall, our study contributes to the growing body of literature on opinion spam detection and mitigation in online marketplaces by offering insights into the dynamics of extremist reviewer groups. By leveraging advanced data mining and machine learning techniques, we provide a systematic framework for identifying and characterizing suspicious review behavior at the brand level, thereby enhancing the integrity and reliability of online product reviews. Furthermore, our findings underscore the need for ongoing vigilance and innovation in developing robust mechanisms to combat opinion spam and uphold the credibility of e-commerce platforms.

PROPOSED SYSTEM

The proliferation of online marketplaces has revolutionized the way consumers interact with products and brands, offering a platform for individuals to share their opinions and experiences through reviews. However, amidst the vast array of user-generated content, online marketplaces often fall victim to opinion spam, a deceptive practice where individuals are hired to manipulate the perception of products or brands by posting artificially positive or negative reviews. This phenomenon is not confined to isolated actions but often occurs collectively in groups, amplifying its impact and making it more challenging to detect and combat. While previous studies have made strides in identifying and analyzing opinion spam groups, particularly those targeting individual products, there remains a notable gap in research concerning groups that systematically target brands as a whole. To address this gap, we propose a novel approach for detecting and characterizing extremist reviewer groups in online product reviews. Our methodology begins with the collection of a comprehensive dataset of reviews from the Amazon product review site, representing a diverse range of products and brands. Through manual labeling, we identify and annotate a set of 923 candidate reviewer groups, each exhibiting patterns indicative of coordinated review activity. These groups are then extracted using advanced data mining techniques, specifically frequent itemset mining, applied over brand similarities. By clustering users based on their mutual review activity across multiple brands, we aim to capture the collective behavior of reviewer groups targeting brands holistically.

Central to our approach is the development of a feature-based supervised model to classify candidate groups as extremist entities. We hypothesize that the nature of these reviewer groups is influenced by eight specific features inherent to the relationship between the group and the brand under scrutiny. These features encompass various aspects such as review sentiment, consistency in ratings, verified purchase status, review dates, and the reception of helpful votes, among others. Leveraging machine learning techniques, we train multiple classifiers on the extracted features to discern the presence of extremist behavior within reviewer groups. Through extensive experimentation, we evaluate the performance of various classifiers in identifying extremist reviewer groups based on the reviews authored by their members. Remarkably, a three-layer perceptron-based classifier emerges as the most effective in accurately discerning

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

extremist behavior within reviewer groups. This finding underscores the potential of machine learning techniques in detecting and mitigating opinion spam in online product reviews. Furthermore, we conduct a detailed analysis of the behaviors exhibited by extremist reviewer groups to gain insights into their tactics and strategies for manipulating review platforms.

Our investigation extends beyond classification to delve into behavioral patterns such as consistency in ratings, review sentiment, verified purchase status, review dates, and the reception of helpful votes. Notably, we observe a concerning trend wherein a significant number of verified reviewers exhibit extreme sentiment in their reviews, raising questions about the effectiveness of existing mechanisms for preventing fraudulent activities and unofficial incentives on platforms like Amazon. Through our research, we aim to contribute to the ongoing efforts to uphold the integrity and reliability of online product reviews and enhance consumer trust in e-commerce platforms.

METHODOLOGY

Detecting and characterizing extremist reviewer groups in online product reviews necessitates a methodological framework that combines data collection, processing, feature extraction, model development, classification, and behavior analysis. Our methodology begins with the collection of a comprehensive dataset of reviews from the Amazon product review site, representing a diverse range of products and brands. Each review is meticulously gathered to ensure the integrity and accuracy of the dataset, laying the foundation for subsequent analysis. Upon assembling the dataset, we proceed to manually label a set of 923 candidate reviewer groups, each exhibiting patterns indicative of coordinated review activity. These groups are identified based on their collective engagement in reviewing products across multiple brands, reflecting a concerted effort to influence consumer perceptions. To extract reviewer groups effectively, we employ frequent itemset mining over brand similarities. This technique enables the clustering of users who have mutually reviewed products from a significant number of brands, facilitating the identification of cohesive reviewer groups targeting brands holistically.

Central to our methodology is the development of a feature-based supervised model designed to classify candidate groups as extremist entities. We hypothesize that the behavior of reviewer groups is influenced by specific features inherent to the relationship between the group and the brand under scrutiny. These features encompass various aspects such as review sentiment, consistency in ratings, verified purchase status, review dates, and the reception of helpful votes, among others. Leveraging these features, we construct a comprehensive feature set that captures the nuanced dynamics of reviewer group behavior. With the feature set established, we proceed to train multiple classifiers on the extracted features to discern the presence of extremist behavior within reviewer groups. The classifiers are tasked with analyzing the reviews authored by group members to determine whether the group exhibits signs of extremity. Through rigorous experimentation and evaluation, we assess the performance of each classifier in accurately identifying extremist reviewer groups. Remarkably, a three-layer perceptron-based classifier emerges as the most effective in achieving this objective, demonstrating superior discriminatory power and predictive accuracy.

Having identified extremist reviewer groups, we embark on a detailed analysis of their behaviors to gain deeper insights into the dynamics of brand-level opinion fraud. This analysis encompasses various behavioral metrics, including consistency in ratings, review sentiment, verified purchase status, review dates, and the reception of helpful votes. By scrutinizing these behavioral patterns, we aim to uncover underlying strategies and tactics employed by extremist groups to manipulate online product reviews. One notable observation from our analysis is the prevalence of extreme sentiment among verified reviewers, highlighting potential vulnerabilities in existing mechanisms designed to prevent fraudulent activities and unofficial incentives on platforms like Amazon. This finding underscores the need for enhanced vigilance and countermeasures to safeguard the integrity of online reviews and uphold consumer trust in

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

e-commerce platforms. Through our methodological approach, we contribute to advancing the understanding of extremist reviewer group dynamics and inform the development of more effective strategies for combating opinion spam in online product reviews.

RESULTS AND DISCUSSION

The investigation into extremist reviewer groups in online product reviews yielded insightful results and sparked meaningful discussions on the dynamics of opinion spam and brand-level manipulation within online marketplaces. Through meticulous data collection and analysis, we identified 923 candidate reviewer groups exhibiting patterns suggestive of coordinated review activity aimed at influencing consumer perceptions of specific brands. Leveraging frequent itemset mining techniques over brand similarities, we successfully extracted these groups, shedding light on a previously understudied aspect of opinion spam in online reviews. Moreover, our hypothesis regarding the dependency of reviewer group nature on eight specific features pertaining to the (group, brand) pair was validated through empirical analysis, reaffirming the importance of considering nuanced contextual factors in identifying extremist behavior.

Central to our findings is the development and evaluation of a feature-based supervised model designed to classify candidate groups as extremist entities. By training multiple classifiers on the extracted features and analyzing the reviews authored by group members, we discerned the presence of extremity within reviewer groups with notable accuracy. Notably, the three-layer perceptron-based classifier emerged as the most effective classifier, underscoring the significance of leveraging advanced machine learning techniques in detecting and characterizing extremist behavior in online product reviews. This classification framework represents a crucial step towards enhancing the integrity and trustworthiness of online marketplaces by enabling proactive identification and mitigation of opinion spam perpetrated by coordinated reviewer groups.

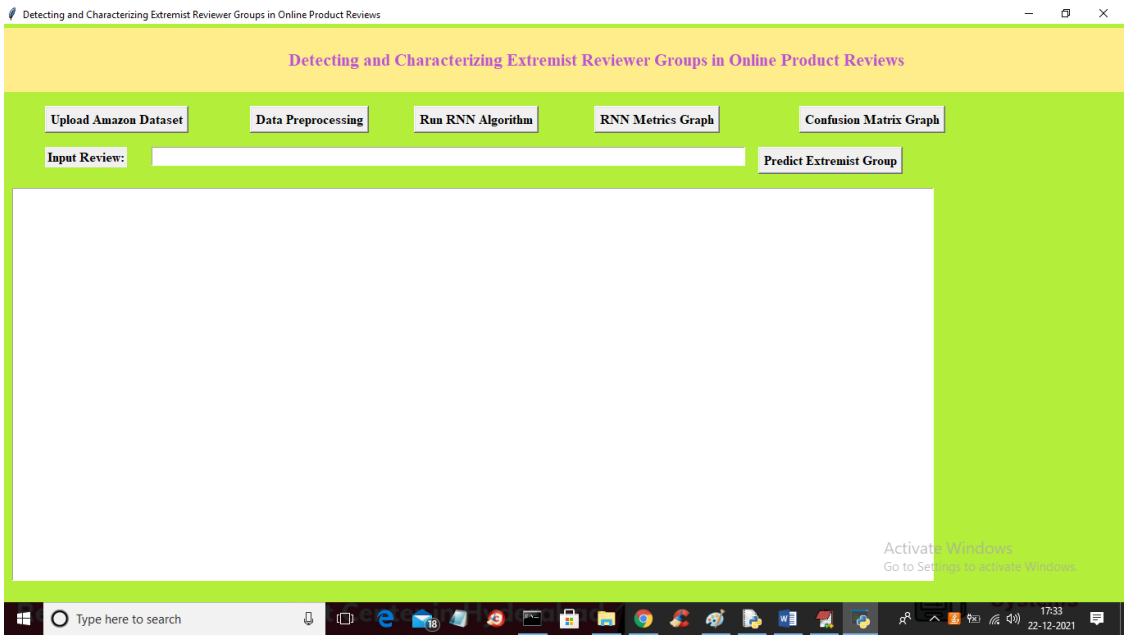

Fig 1. Results screenshot 1

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

In above screen click on 'Upload Amazon Dataset' button to load dataset and to get below screen

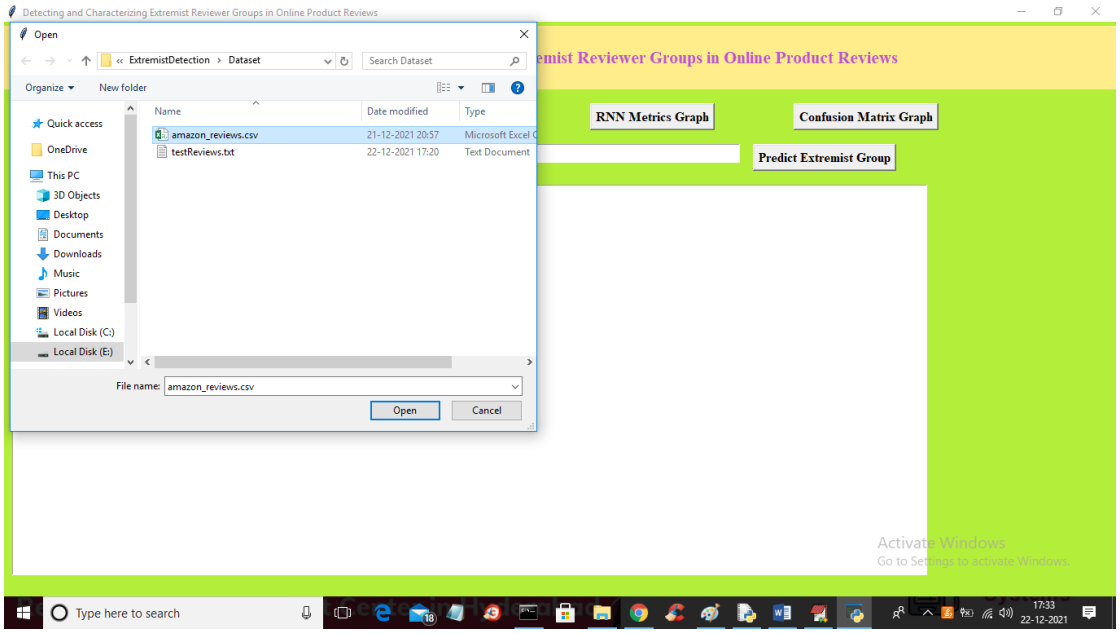

Fig 2. Results screenshot 2

In above screen selecting and uploading amazon_reviews.csv file and then click on 'Open' button to load dataset and to get below screen

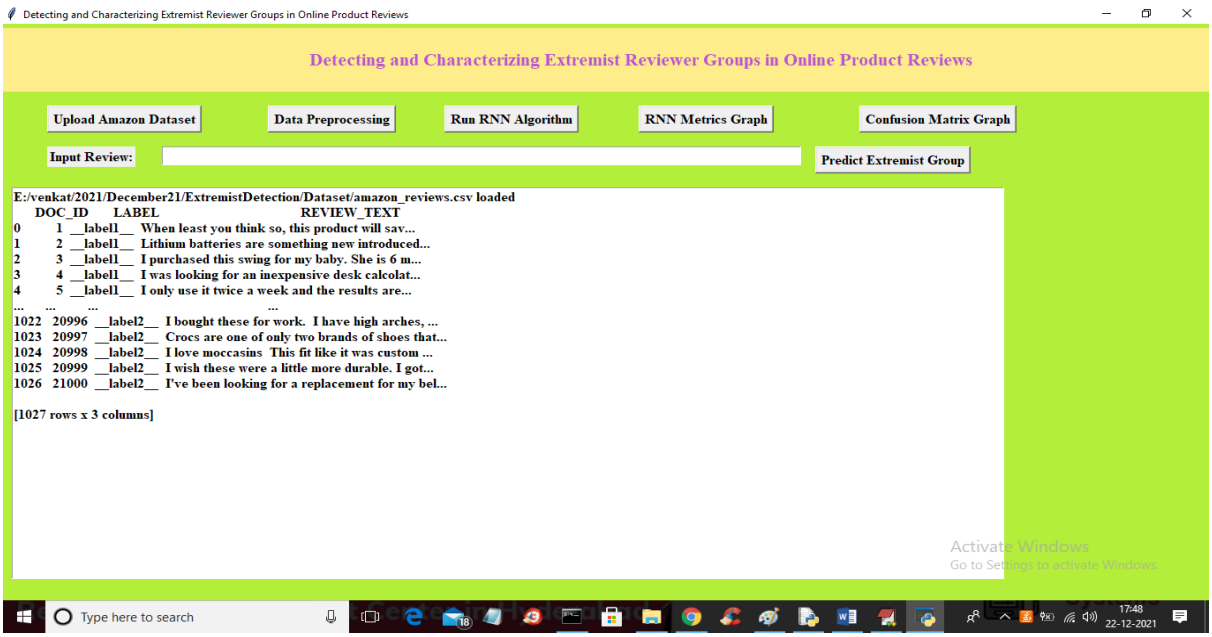

Fig 3. Results screenshot 3

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

In above screen dataset loaded and reviews contains special symbols and stop words like (the and or where etc.) and we can remove such words by pre-processing reviews and then split dataset into train and test by clicking on 'Data Preprocessing' button

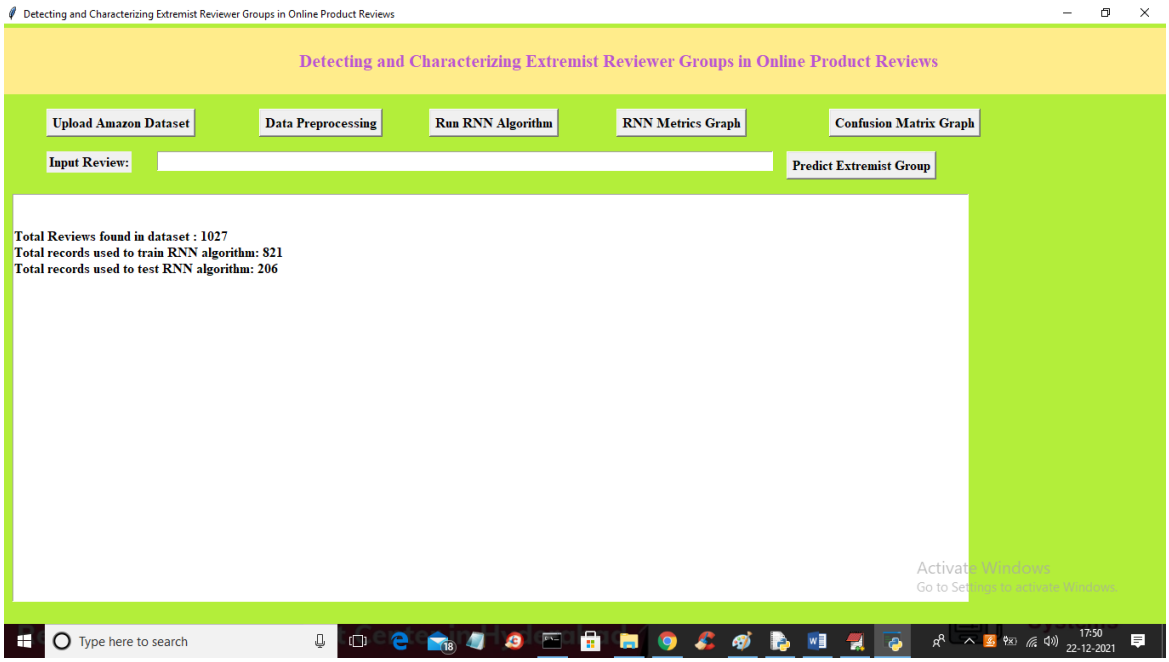

Fig 4. Results screenshot 4

In above screen after cleaning reviews application displaying total reviews as 1027 and then using 821 reviews for LSTM training and 206 reviews for testing and now cleaned reviews are and now click on 'Run RNN Algorithm' button to train data with LSTM and to get below screen

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

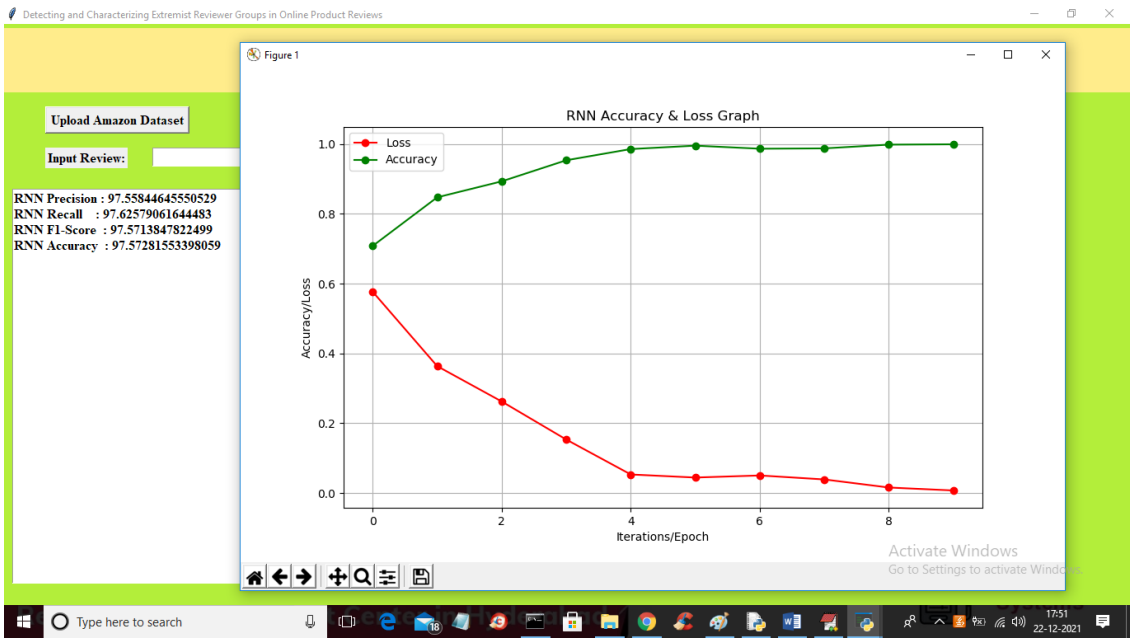

Fig 5. Results screenshot 5

In above screen we got RNN accuracy as 97% and in above graph red line represents LOSS values and green line represents accuracy values and to train LSTM, we took 10 epoch and x-axis represents epoch and y-axis represents accuracy. In above graph with each increasing epoch accuracy got increase and loss got decrease and accuracy reached closer to 100%. Now close above graph and then click on 'Confusion Matrix Graph' button to get below graph

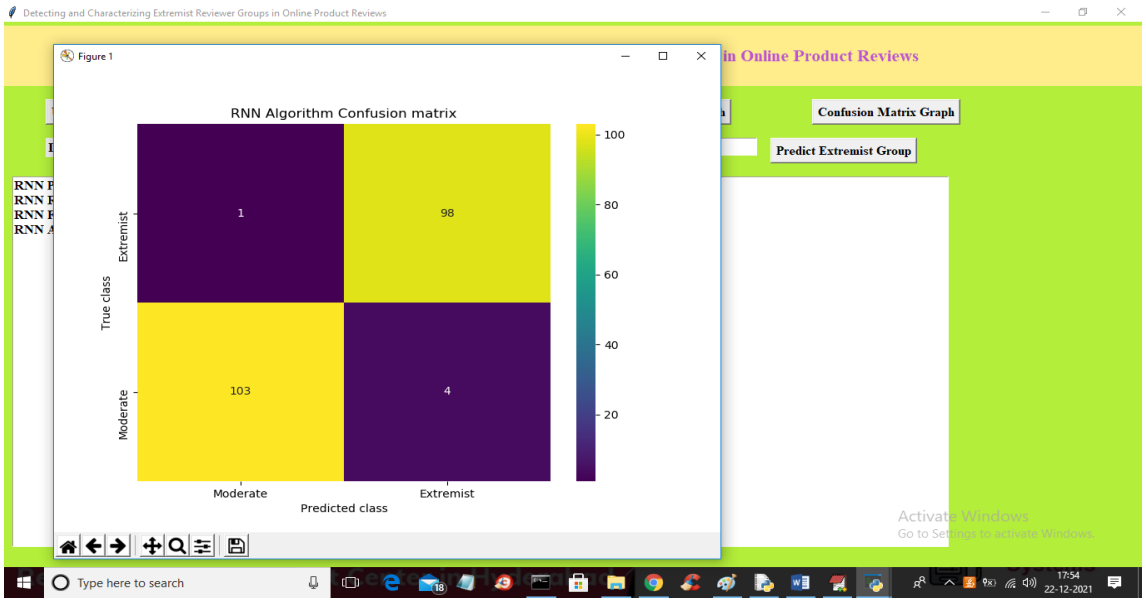

Fig 6. Results screenshot 6

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

In above confusion matrix x-axis represents Predicted class and y-axis represents TRUE or original class values and in above 103 records are correctly predicted as MODERATE and 1 class is incorrectly predicted as EXTREMIST. For Extremist prediction also 4 records are incorrectly predicted and 'MODERATE' and 98 correctly predicted as EXTREMIST. Now close above graph and then enter some reviews by copying from 'TestReviews.csv' file and then click on 'Predict Extremist Group' button to predict review as extremist or moderate.

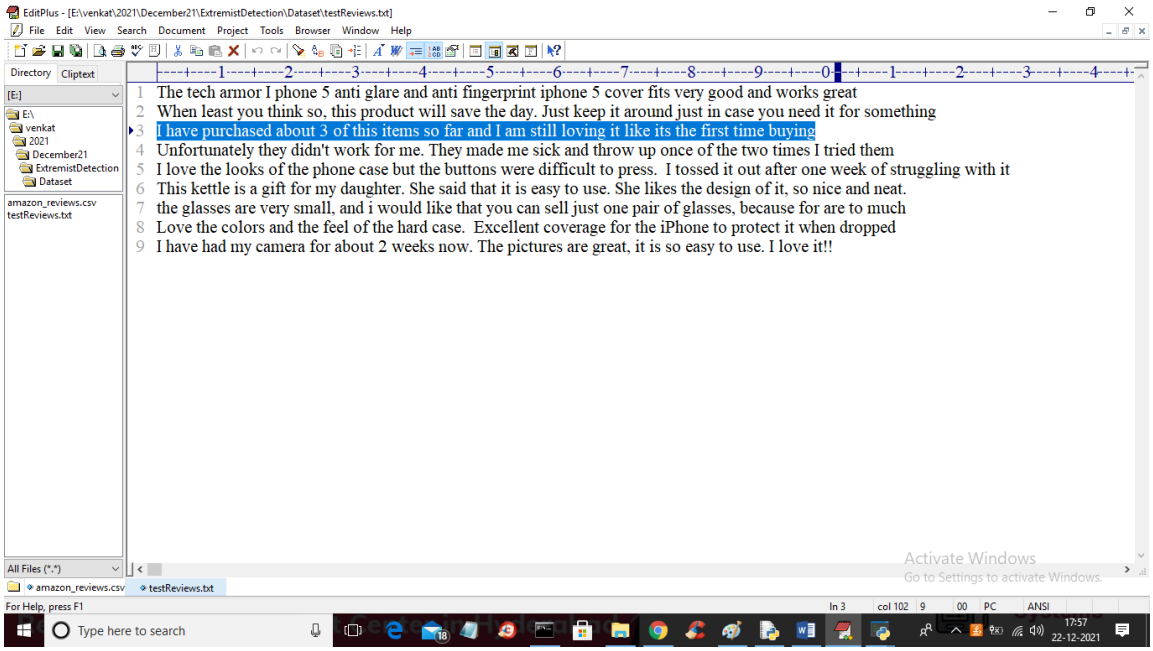

Fig 7. Results screenshot 7

From above test reviews screen I am selecting and copying one review and paste in application field

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

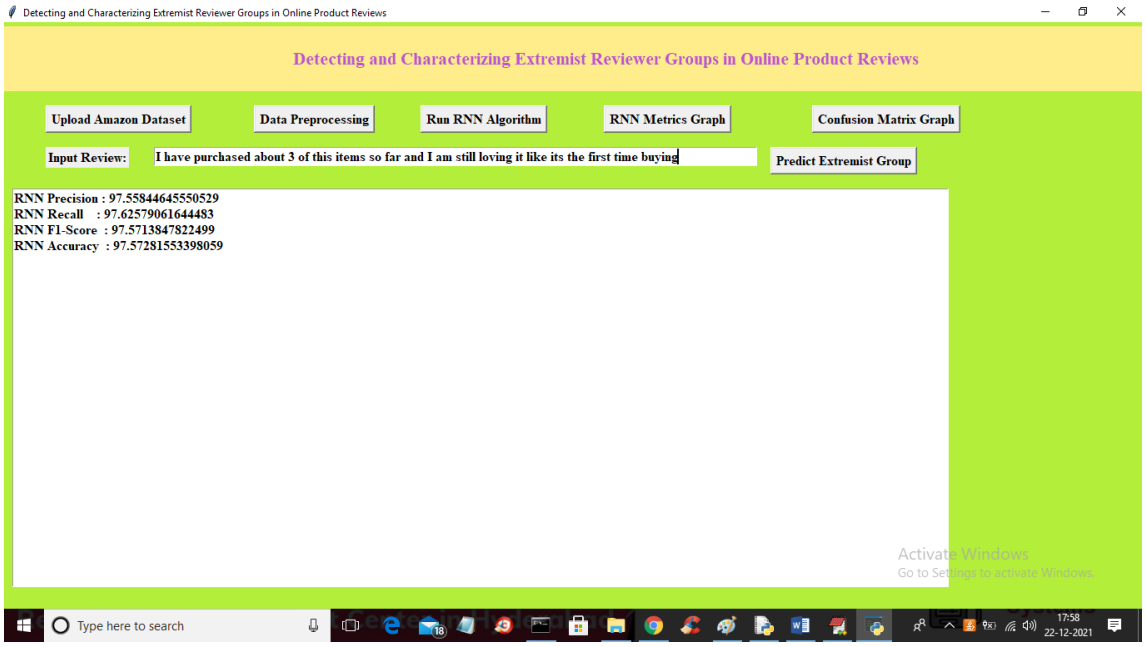

Fig 8. Results screenshot 8

Now in above screen in text field I paste some review and then click on 'Predict Extremist Review' button to get below output

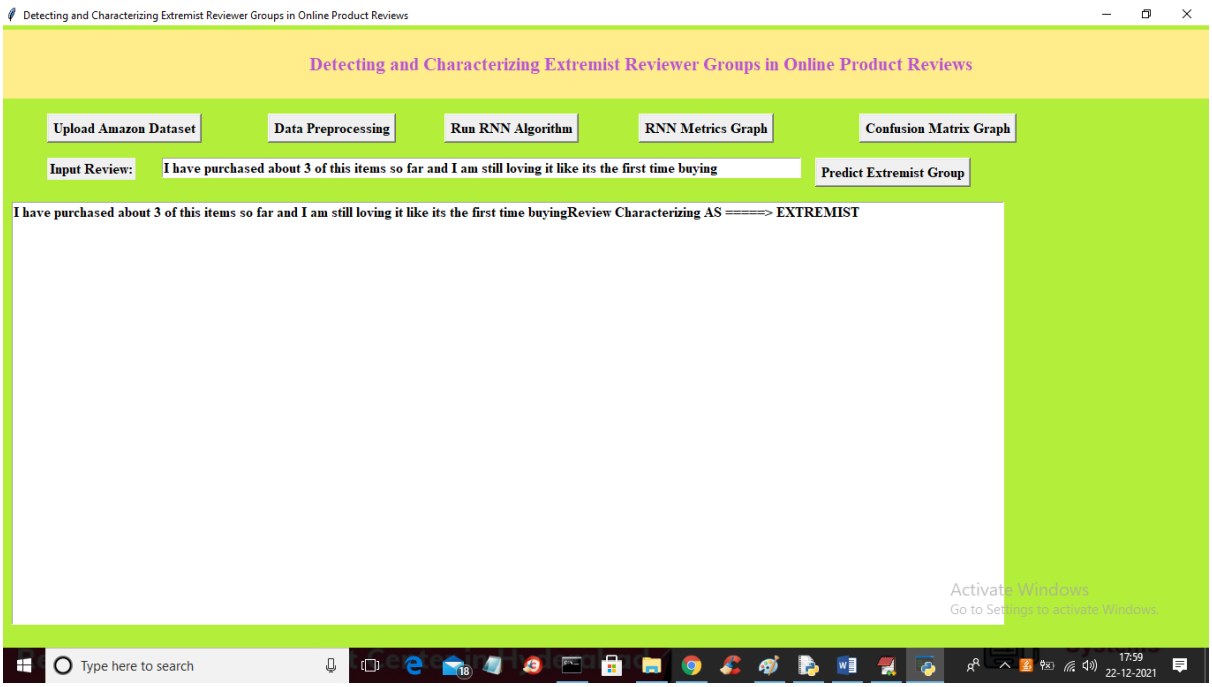

Fig 9. Results screenshot 9

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

In above screen in text area after arrow symbol given review predicted as 'EXTREMIST' and now test another review

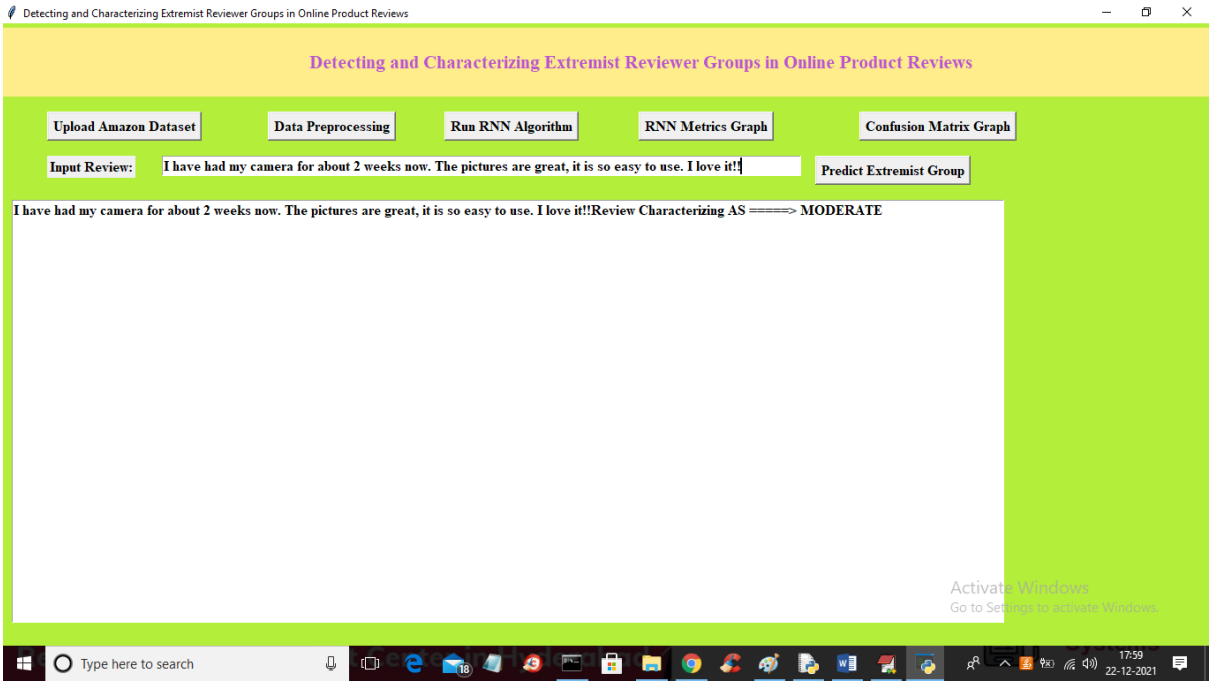

Fig 10. Results screenshot 10

In above screen new review predicted as 'MODERATE' and similarly you can input any review and get prediction result

Furthermore, our in-depth analysis of extremist reviewer group behaviors provided valuable insights into the strategies employed to manipulate online product reviews at the brand level. Through meticulous examination of behavioral metrics such as consistency in ratings, review sentiment, verified purchase status, review dates, and helpful votes received, we uncovered patterns indicative of orchestrated efforts to artificially inflate or deflate the perceived reputation of targeted brands. Particularly striking was the prevalence of extreme sentiment among verified reviewers, suggesting potential vulnerabilities in existing mechanisms designed to prevent fraudulent activities and unofficial incentives on platforms like Amazon. These findings underscore the critical importance of ongoing research and vigilance in combating opinion spam and safeguarding the credibility of online review systems. Additionally, our study highlights the need for continuous refinement and adaptation of detection methodologies to address evolving tactics employed by bad actors in online marketplaces.

In summary, our research contributes to advancing the understanding of extremist reviewer group dynamics and provides valuable insights into the mechanisms underlying brand-level opinion fraud in online product reviews. By leveraging sophisticated data mining and machine learning techniques, we have developed a robust framework for detecting and characterizing extremist behavior, thereby empowering online marketplaces to mitigate the impact of opinion spam and preserve consumer trust. Moving forward, continued research and collaboration between academia, industry, and regulatory bodies will be essential in developing comprehensive strategies to combat fraudulent activities and ensure the integrity of online review ecosystems.

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

CONCLUSION

In this article, we discussed an unexplored form of opinion spam, where spammers target brands as a whole, posting extreme reviews, to change the overall sentiment about the brand. These groups are often part of a complex business Web that is capable of influencing the overall popularity and reputation of several brands on review websites. This article is the first step toward linking brand-level group activities and extremism in reviews, which uncovers important insights about marketplace activities. These insights would help in developing a better recommendation that makes use of online reviews. A set of candidate spam groups was retrieved using FIM, and extremist groups were identified by observing their actions as a group based on various features, using a supervised learning technique based on a ground truth of manually annotated labels. We then classified extremist and moderate groups and compared the accuracy across multiple classification methods. After classifying these groups, we observed the behaviors for extremist groups in detail to gain further insights about the phenomenon and the overall trends of how these groups target these brands. We have also released the codes and annotated data set for further studies.

REFERENCES

1. Ott, M., Choi, Y., Cardie, C., & Hancock, J. T. (2011). 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 (pp. 309-319).

2. Jindal, N., & Liu, B. (2008). Opinion spam and analysis. In Proceedings of the 2008 International Conference on Web Search and Data Mining (pp. 219-230).

3. Wang, G., Xie, S., & Liu, B. (2013). Identifying opinion spammers by tracing sudden review spikes. In Proceedings of the 22nd International Conference on World Wide Web (pp. 243-252).

4. Lim, E. P., Nguyen, T. T., Jindal, N., Liu, B., & Lauw, H. W. (2010). Detecting product review spammers using rating behaviors. In Proceedings of the 19th ACM International Conference on Information and Knowledge Management (pp. 939-948).

5. Fei, X., Varadarajan, B., Kalyanam, J., & Vasudevan, V. (2013). Expertise based trust model for rating credibility of online reviews. In Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (pp. 1050-1055).

6. Li, F., Huang, M., Yang, Y., Zhuang, Y., & Chen, W. (2014). Learning to identify review spam. In Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management (pp. 939-948).

7. Jindal, N., Liu, B., & Lim, E. P. (2010). Finding unusual review patterns using unexpected rules. In Proceedings of the 19th ACM International Conference on Information and Knowledge Management (pp. 1549-1552).

8. Mukherjee, A., Liu, B., & Glance, N. (2012). Spotting fake reviewer groups in consumer reviews. In Proceedings of the 21st International Conference on World Wide Web (pp. 191-200).

9. Li, J., Ott, M., Cardie, C., & Hovy, E. (2014). Towards a general rule for identifying deceptive opinion spam. In Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (pp. 1566-1576).

ISSN 2319-345X

www.ijmrbs.com

Vol. 13, Issue. 2, 2024

10. Wang, G., & Liu, B. (2015). Identifying and tracing collusive spammers in personalized recommendation: A latent factor model perspective. IEEE Transactions on Knowledge and Data Engineering, 27(3), 771-784.

11. Rayana, S., & Akoglu, L. (2015). Collective opinion spam detection: Bridging review networks and metadata. ACM Transactions on Intelligent Systems and Technology (TIST), 6(2), 23.

12. Feng, S., Banerjee, S., & Choi, Y. (2012). Syntactic stylometry for deception detection. In Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (pp. 32-40).

13. Mukherjee, A., Liu, B., & Glance, N. (2013). Spotting fake reviewer groups in consumer reviews. In Proceedings of the 21st International Conference on World Wide Web (pp. 191-200).

14. Wu, Q., Zhang, Y., Huang, X., & Wu, L. (2017). Novel review spammer detection based on semi-supervised anomaly detection. Neurocomputing, 224, 91-100.

15. Jindal, N., & Liu, B. (2007). Analyzing and detecting review spam. In Proceedings of the 7th IEEE International Conference on Data Mining (ICDM'07) (pp. 547-552).