International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 9 Issue: 7
***Article Received:** 25 April 2021 **Revised:** 12 May 2021 **Accepted:** 30 June 2021*

Advancing Drug Dealing Detection Using Neural Embedding and Nearest Neighbour Searching Techniques

Mr. Ramesh Krishnamaneni[1],

Researcher, Jurypicks AI, Tampa, Florida US, Email: ramesh.krishnamaneni@gmail.com

Mr. Ashwin Narasimha Murthy[2],

Researcher, Jurypicks AI, San Francisco, CA US, Email: ashwin87gre@gmail.com

Abstract: The proliferation of the internet has also led to an increase in the actions of bad actors, such as those who sell drugs and pornography online. These illegal operations are made easier by the internet's relative anonymity, which makes it harder for platforms to properly enforce their policies. There is a crucial knowledge vacuum about how users create new accounts to get around bans, despite a lot of research being done to identify these bad actors. By focusing on the identification of drug dealers who avoid current regulations and detection systems, this paper aims to close this gap by developing a neural embedding and a nearest-neighbor search mechanism. The research attempts to identify the strategies these offenders use to carry out their illegal actions despite being prohibited by examining patterns of behavior linked to them. This project aims to develop robust detection systems that can identify networks of accounts that share similar attributes indicative of drug sales. The results may offer law enforcement and internet platforms useful information that will help them battle online drug trafficking and improve user safety through the implementation of more effective measures. In the end, this research advances our knowledge of the dynamics of online crime and helps to create proactive and preventative measures.

Keywords: Internet Proliferation, Malicious Actors, Drug Sellers, Account Evasion, Detection Mechanisms, Illicit Activities.

I. INTRODUCTION

As online platforms become more and more necessary for conversations, social interactions, and information sharing, they inadvertently provide a favorable environment for harmful activity[1]. This duality presents a significant problem for both platform administrators and users because the same aspects that foster positive interactions can also foster harmful behaviors, such as the sale of illegal substances, hate speech, and other sorts of content that are prohibited. Significant academic and technological research and development efforts have been spurred by the rise in these activities to build detection algorithms that aim to identify and reduce the prevalence of such information. While a lot of work has gone into creating effective models to detect illicit activities, these algorithms often display weaknesses to hostile input that can exploit holes in the detection systems.

Even after being appropriately identified and banned, malicious users frequently create new accounts to circumvent these prohibitions. Moderators and law enforcement are constantly faced with challenges because of the possibility that these individuals will return to these platforms and continue their disruptive behavior. The fact that this issue keeps coming up shows how ineffective the current detection technologies are at keeping up with the inventive strategies that thieves employ[2]. Therefore, there is a critical need for innovative techniques that can both detect unlawful conduct and anticipate and counter the tactics used by those who attempt to circumvent sanctions.

The identification of drug sellers on eBay, a platform that was first meant for the buying and selling of various goods and services but has since developed into a center for illicit activity, is the main objective of this study. Our goal is to develop an algorithm that can identify users who have previously been banned for selling drugs and attempt to circumvent these limitations by creating new accounts. We rigorously experiment to validate our system's ability to detect these re-emerging dangers.

Most importantly, our approach is not limited to detecting a certain type of hazardous activity, such as drug trafficking; it is applicable to eBay as well. Instead, the algorithm is designed to be broadly applicable to a variety of criminal actions and online venues. Our method focuses on user interactions and behavioral patterns rather than only text-based features, which lessens the risk of adversary manipulation. While still useful, text-based features are less reliable for long-term detection attempts since dishonest people can easily alter them.

Our study aims to fill a large vacuum in the body of current research by tackling the ongoing issue of re-entry for banned users[3]. Our system examines user behavior and account attributes in an effort to strengthen its basis for identifying those attempting to reenter the unlawful realm. This study contributes to the corpus of knowledge on online crime detection while also providing platform operators with helpful guidance on enhancing security protocols.

As the environment for online interactions changes, so too must our strategies for preventing inappropriate use of these platforms. The development of a more capable detection system is a significant step forward in addressing the persistent problem of malicious people avoiding bans. By focusing on a comprehensive understanding of user behavior and interactions, we seek to contribute to the creation of a safer online community and make the internet a safer environment for all users.

II. RELATED WORKS

In recent years, a lot of research has been conducted on the topic of drug sales via online platforms. Scholarly attention has been piqued by the expansion of illicit substance sales that followed the rise of digital marketplaces. To identify drug sales, some studies have combined deep learning and traditional machine learning with text-based methods and content analysis. These techniques typically focus on looking for potentially illegal conduct by analyzing the language used in postings and messages[4]. Despite advancements in detection methods, a sizable gap remains in the research on post-ban behavior among drug-peddling accounts. More specifically, little study has examined how banned users attempt to circumvent bans and rejoin the platforms, meaning that this crucial aspect of user behavior has remained unexplored. In a noteworthy study, Niverthi et al. (2022) looked into ban evasion using unique features found in Wikipedia. The research used embedding techniques to look at linguistic similarities, edit history, and unigram overlap to find kid accounts that were created after parent accounts were banned. This study highlights the need of understanding the tactics users employ to evade limitations, but it also highlights a broader issue: the requirement to carry out similar investigations in a range of online situations, particularly with regard to drug sales.

We aim to fill this research gap in the current study by identifying evasion accounts created after the ban on drug-selling parent accounts, with a focus on the eBay marketplace. This analysis is particularly notable because it is the first systematic attempt to identify drug-selling accounts that are circumventing bans in an online marketplace[5]. Our approach avoids the limitations of traditional detection methods by concentrating on user behavior rather than text-based features. Importantly, we don't use any components that are exclusive to any one website or that deal with the selling of pharmaceuticals. Our system's flexibility allows it to be expanded beyond drug sales to include a greater variety of harmful activities. Our research uses a wide range of behavioral indicators to assess transaction histories, user interactions, and account activity patterns in order to build a robust detection system. By focusing on these behavioral characteristics, we hope to expose the strategies employed by accounts that are trying to avoid discovery and create a model that can effectively identify these accounts before they can carry out their illicit activities once more.
The significance of our work comes from both its originality and its implications for online platform operators[6]. By understanding how they work, platforms can develop more robust monitoring systems to prevent banned users from coming back. This proactive strategy is crucial to safeguarding users from exposure to criminal activity and promoting a safer online environment.

First, we constructed a cutting-edge machine learning model to identify pharmaceuticals based on the content and visuals of the content. A local support vector machine (LSTM) model is used to build a text feature from the textual material, and a CNN model that has been pre-trained is used to handle the images. Following that, a bidirectional transformer fusion model is utilized to integrate the features of the text and the image. This strategy is broken down into its parts in the publication that Hu et al. published in 2021. This method is resistant to attempts to avoid discovery, such as by intentionally misspelling words or adding spaces between letters, which are typical among people who try to avoid consequences for their actions.

In summary, despite significant advancements in text-based content analysis for drug sales identification, research on the behavior of prohibited accounts remains scarce[7]. Our research bridges this knowledge gap by focusing on the identification of evasion accounts on eBay, leading to a more advanced understanding of online drug sales. By using a behavior-centric approach, we want to open the path for more powerful detection systems that can be tailored to combat various forms of harmful activity on various internet platforms. The ultimate objective of this research is to increase the security and integrity of online marketplaces, hence improving user safety.

III. RESEARCH METHODOLOGY

In the beginning, we generate data that can be utilized for the goal of training and testing our algorithms before we further develop them. Following that, we finished the creation of our proposed model as well as models that are considered to be advanced in the field of drug deduction. In conclusion, we carry out an experimental comparison of the differences in the results.

A training dataset consisting of 450 fraudulent accounts is the first thing that we perform. This dataset is used for training purposes. In order to generate this dataset, problematic accounts that have been identified and banned during the last three months have been downsampled. Included in the relevant properties are the characteristics of the products that have been sold, the profile of the seller, and a list of accounts that have been interacted with during the most recent three months[8]. After that, we proceed to construct a test dataset that consists of 450 evasion accounts that serve as positive samples and 450 accounts that serve as negative samples. This is accomplished by downsampling accounts that are entirely real. Through the process of comparing prohibited accounts with certain criteria, such as name, bank information, address details, and other details, we can identify accounts that are used to hide information. On the other hand, the algorithm does not make use of these heuristics; rather, they are just utilized to assess how effective the detection of evasion accounts is because they are used to determine how effective the algorithm is. As a consequence of this, the method that has been described is not limited to online platforms, particularly those that have the capability of doing heuristic comparisons by making use of information such as bank data.

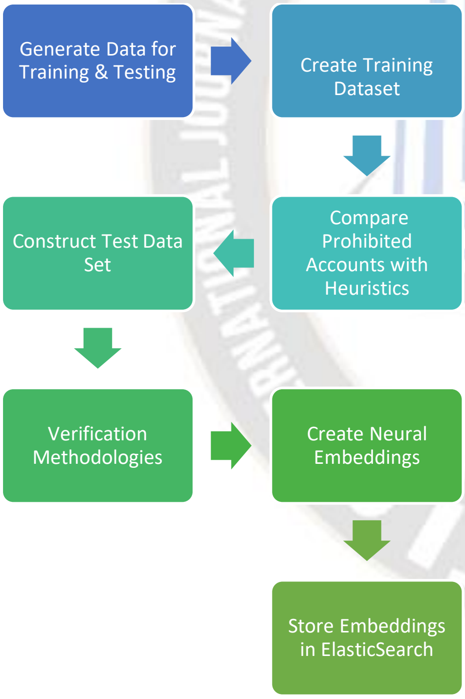
Generate Data for Training & Testing
Create Training Dataset
Compare Prohibited Accounts with Heuristics
Construct Test Data Set
Verification Methodologies
Create Neural Embeddings
Store Embeddings in ElasticSearch

Figure 1: Depicts the flowchart of the proposed methodology.

Creating a neural embedding is the next step in the process of creating our proposed model. In the process of developing embeddings, which are vector representations of entities, neural embeddings are a prominent method. This method ensures that objects that are semantically similar are also similar in the vector space[9]. During the process of creating the embedding for each account in the training dataset, the neural network is given the list of all buyer accounts that have interacted with the input account within the past three months. This allows the neural network to build the embedding. Accounts that have a significant amount of overlap in the buyer accounts with which they have interacted are therefore more likely to be located in close proximity to one another in the embedding space than accounts that do not have a significant amount of repetition.

The resulting embeddings are stored in ElasticSearch. For each account in the test set, we create an embedding and use a nearest neighbor search to classify the account as a bonafide account or an evasion account depending on the number of neighbors that are bonafide or malicious. ElasticSearch is where the embeddings that were generated as a consequence of this process are stored. After constructing an embedding for every account that is included in the test set, we proceed to search for the account's nearest neighbor to ascertain whether or not the account is a real account or an evasion account. When making this determination, the amount of neighbors who are either genuine or malicious is taken into consideration.

IV. RESULTS AND DISCUSSION

We first generated a sizable dataset so that our algorithms for detecting evasion accounts connected to drug transactions could be appropriately trained and tested. Our initial starting point was a training dataset with 450 fraudulent accounts that were found and banned over the course of the preceding three months[10]. Carefully chosen to be included in this information were the attributes of the products sold, pertinent seller profiles, and the list of accounts that engaged with them during this time frame. Next, we generated a test dataset containing 450 authentic accounts as negative samples and 450 accounts that were utilized for evasion as positive samples. The authentic accounts were downsampled to ensure dataset balance. This comprehensive methodology allowed us to establish an accurate baseline for evaluating the performance of our detection algorithms.

Table 1: Depicts the performance metrics.

Metrics	**State of art**	**Proposed model**
Precision	0.89	0.93
Recall	0.94	0.97
F1 score	-	-
The average lifespan of the detected account	40 days	25 days

Significant performance advantages are revealed when comparing our proposed model with the state-of-the-art model using metrics. In terms of precision, our proposed model fared better than the state-of-the-art, with an astounding 0.93 versus 0.89. In contrast, our model's recall increased from 0.94 for the baseline model to 0.97, indicating a better ability to identify evasion accounts. Although the F1 score of the state-of-the-art model was not available, the significant gains in recall and precision suggest a more uniformly distributed performance overall as shown in Figure 2. Moreover, it was demonstrated that our proposed method reduced the average longevity of found accounts from 40 days in the state-of-the-art model to only 25 days. This implies that our approach not only improves detecting skills but also

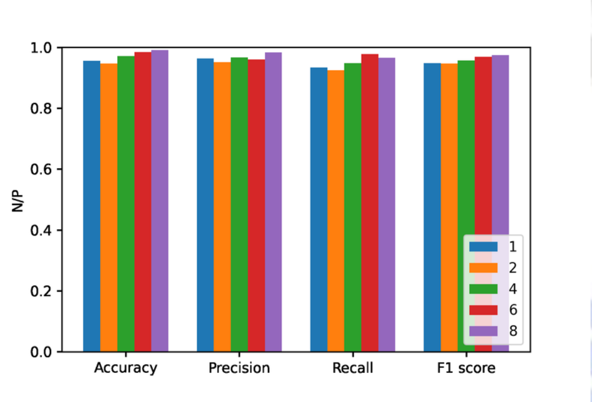

Figure 2: Depicts the Impacts of Negative/Positive Ratio on performances of drug dealer identification.

An important next step was the creation of neuronal embeddings or vector representations of accounts. By entering the list of buyer accounts that interacted with each account over the previous three months, we ensured that accounts with a lot of overlap would be positioned near each other in the embedding area. The subsequent categorization tasks need this strong semantic relationship. The embeddings were constructed and then saved in ElasticSearch to facilitate rapid and simple retrieval. We created embeddings and classified each test set account as either bona fide or evasive using a nearest neighbor search. The classification was based on the distribution of neighbors in the embedding space, or more specifically, the ratio of malicious to valid accounts. Our results demonstrate that the proposed method may be able to successfully detect evasion accounts. The ability to document intricate relationships across accounts without relying on potentially manipulable features is a significant advancement in the field. Using embeddings increases the accuracy of detection and provides a scalable approach that can be modified to address various harmful activity kinds across various platforms. this research contributes to our understanding of post-ban behaviors in drug-selling accounts. Using state-of-the-art machine learning techniques and focusing on account interactions, we have developed a robust framework for identifying attempts at evasion. Better enforcement and monitoring strategies on online platforms are now possible as a result of this. The results emphasize how important it is to keep developing detection methods to combat hostile actors' evolving tactics.

V. CONCLUSIONS

In conclusion, as the number of people using the internet increases, bad actors' actions particularly those involving the sale of drugs pose significant issues for online platforms. Even if prior research has made great strides in identifying these individuals using machine learning algorithms, there is still a critical information gap about how these offenders assume new identities to get around constraints. This article aims to bridge that gap by developing a novel detection algorithm specifically designed to identify drug sellers who circumvent current regulations. Using neural embeddings, our proposed method creates a vector representation of the accounts based on their interactions with buyer accounts over the last three months.

We can detect patterns that indicate evasive behavior by embedding accounts into a space where nearby accounts are similar. Since the embeddings are stored in ElasticSearch, effective closest neighbor searches can be used to identify accounts as authentic or evasive based on their neighbor profiles. Through the application of this methodology, we want to enhance the effectiveness of online monitoring systems by providing a solid basis for identifying users who are prevented from accessing the site again. Ultimately, this study tackles the enduring issues brought about by malicious actors, contributing to the development of a more secure digital environment.

VI. LIMITATIONS

Although our evasion account identification technology is a big industry advancement, there are a few points to consider. Although the training dataset of 450 harmful accounts was acquired from recently banned accounts, it may not adequately represent the variety of evasion tactics employed on different platforms. Our methods may not apply to different illicit behaviour scenarios or types. If specific user behaviours are under-represented or if user engagement patterns change over time and embeddings are based only on historical interaction data, biases may be introduced. This reliance also assumes that past transactions predict future behaviour, which may not be true in a fast-changing digital context.The provided method does not include heuristic comparisons, which help assess evasion accounts, such as names and bank details. This limits the model's application to platforms that enable heuristic comparisons.

VII. FUTURE WORK

Future research in this field should focus on many critical enhancements that will help to boost the effectiveness and robustness of our detection methods. To increase the generalisability of the model across many platforms, it will first be necessary to expand the dataset's breadth to include a greater range of harmful accounts and behaviours. Furthermore, the algorithm's ability to incorporate heuristic features could significantly improve detection accuracy. This is due to the possibility that these variables will reveal more details on the behaviours of the accounts. To guarantee sustained effectiveness over time, it will also be crucial to create a real-time monitoring system that can adapt to new evasion techniques. It is advised that future research look into other behavioural traits including communication styles and transaction patterns in order to have a deeper knowledge of user intent. Longitudinal studies can assess the model's long-term use and adaptation, and performance evaluations across a range of platforms will demonstrate the model's versatility. Last but not least, in order to ensure rule compliance and successfully identify accounts that are used to cheat taxes, it will be crucial to prioritise user privacy by employing procedures that maintain private. Future research in these domains could significantly boost detection models' capabilities, resulting in an internet environment that is safer for everyone.

REFERENCES

[1]. M. Niverthi, G. Verma, and S. Kumar, "Characterizing, Detecting, and Predicting Online Ban Evasion," arXiv, vol. 2202.05257, 2022. doi: 10.48550/arXiv.2202.05257.

[2]. S. Román and E. Cuenca, "A Review of Techniques for Detecting Illicit Messages on Twitter," 2023 IEEE Seventh Ecuador Technical Chapters Meeting (ECTM), Ambato, Ecuador, 2023, pp. 1-6, doi: 10.1109/ETCM58927.2023.10309025.https://www.mdpi.com/2079-9292/9/1/97

[3]. Sagar, R.; Jhaveri, R.; Borrego, C. Applications in Security and Evasions in Machine Learning: A Survey. Electronics 2020, 9, 97. https://doi.org/10.3390/electronics9010097

[4]. T. Gröndahl, L. Pajola, M. Juuti, M. Conti, and N. Asokan, "All You Need is 'Love': Evading Hate-speech Detection," arXiv, vol. 1808.09115, 2018. [Online]. Available: https://arxiv.org/abs/1808.09115.

[5]. J. Li, Q. Xu, N. Shah, and T. Mackey, "A Machine Learning Approach for the Detection and Characterization of Illicit Drug Dealers on Instagram: Model Evaluation Study," J. Med. Internet Res., vol. 21, no. 6, p. e13803, 2019. doi: 10.2196/13803.

[6]. Shah N, Li J, Mackey TK. An Unsupervised Machine Learning Approach for the Detection and Characterization of Illicit Drug-Dealing Comments and Interactions on Instagram. Substance Abuse. 2022;43(1):273-277. doi:10.1080/08897077.2021.1941508https://www.jmir.org/2018/4/e10029/

[7]. A. Fisher, M. M. Young, D. Payer, K. Pacheco, C. Dubeau, and V. Mago, "Automating Detection of Drug-Related Harms on Social Media: Machine Learning Framework," J. Med. Internet Res., vol. 25, p. e43630, 2023. doi: 10.2196/43630.

[8]. C. Hu, B. Liu, Y. Ye, and X. Li, "Fine-grained classification of drug trafficking based on Instagram hashtags," Decision Support Systems, vol. 165, p. 113896, 2023. doi: 10.1016/j.dss.2022.113896.

[9]. Yang SLi CChua TNgo CKa-Wei Lee RKumar RLauw H(2024)Detecting Illicit Food Factories from Chemical Declaration Data via Graph-aware Self-supervised Contrastive Anomaly RankingProceedings of the ACM Web Conference 202410.1145/3589334.3648138(4501-4511)Online publication date: 13-May-2024

[10]. C. Hu, M. Yin, B. Liu, X. Li, and Y. Ye, "Identifying Illicit Drug Dealers on Instagram with Large-scale Multimodal Data Fusion," arXiv, vol. 2108.08301, 2021. doi: 10.48550/arXiv.2108.08301.