ARTIFICIAL NEURAL NETWORKS FOR THE IDENTIFICATION OF FALSE PROFILES

Merugu Malleshwari
M.Tech Student, Dept. of CSE,
Vaagdevi College of
EngineeringWarangal, India.

Dr. Swathi Bolugoddu
Asst.Professor,Dept. of CSE
Vaagdevi College of Engineering
Warangal, India.

Chikati Aravind Kumar
Asst.Professor
Dept. of CSE
VCE,Warangal, India

Dr.N.Sathyavathi
Asso.Professor,&Head,
Dept. of CSE
VCE,Warangal, India

ABSTRACT

The exponential rise of social networking platforms and online services has made user authentication and profile credibility a critical concern. The widespread creation and misuse of fake user profiles have led to privacy violations, fraudulent activities, and reputational damage, necessitating the development of intelligent detection mechanisms. This research introduces a Django-based web application that utilizes Artificial Neural Networks (ANNs) to effectively detect fake profiles based on structured user data. The system is designed with two key modules: an admin interface that facilitates data management and model training using a dataset of 200 real and fake profiles, and a user interface that allows for real-time prediction of profile authenticity based on custom input features. The ANN model is trained to capture subtle patterns across multiple user attributes, including behavioural and profile-based indicators. Preprocessing techniques are applied to standardize the input, and the trained model is integrated into the web application to enable seamless interaction. Experimental results demonstrate a high level of accuracy in classification, validating the use of ANNs for this application. This project not only emphasizes the potential of neural network architectures in solving classification problems related to cybersecurity but also showcases the practicality of deploying machine learning models in full-stack web environments for real-world impact.

Keywords: Artificial Neural Networks, Fake Profile Detection, Machine Learning, Cybersecurity, Django Web Application, Classification Model, User Authentication.

I. INTRODUCTION

In the current digital age, online platforms have become an integral part of daily life, enabling people to connect, communicate, and conduct business across global boundaries. Social networks, e-commerce websites, job portals, and online forums collectively host millions of user accounts, many of which play a critical role in facilitating meaningful interactions and transactions. However, the widespread use of these platforms has also led to the proliferation of fake user profiles—fraudulent accounts created with malicious intent or to mislead other users. These fake profiles are often employed for a variety of harmful purposes, including spreading misinformation, phishing scams, spamming, impersonation, and manipulating public opinion. As a result, fake profile detection has emerged as a vital area of research within the broader domain of cybersecurity and data analytics.

Despite the implementation of various rule-based mechanisms and manual moderation strategies by platform providers, fake profiles continue to bypass traditional detection techniques due to their increasing sophistication and evolving behavioural patterns. Static filters based on predefined rules or keyword lists are no longer sufficient to keep up with the dynamic nature of these profiles. In response to this challenge, researchers and developers are turning to machine learning techniques, particularly Artificial Neural Networks (ANNs), which offer significant potential in uncovering complex patterns within large datasets. ANNs, inspired by the structure and function of the human brain, are capable of modeling non-linear relationships and learning from diverse feature sets, making them particularly well-suited for classification tasks such as identifying fake versus genuine profiles.

This study presents the design and implementation of a machine learning-powered web application using Django that leverages ANNs to detect fake profiles.

The core motivation for this research stems from the need to develop an intelligent, automated, and scalable solution that can assist administrators in safeguarding online platforms while offering users a tool for verifying authenticity. Unlike traditional applications that rely solely on static validation, the proposed system incorporates a trained ANN model that adapts to new data and improves its predictive capabilities over time.

The web application is developed with two primary user roles in mind: the admin and the user. The admin module serves as the control center where administrators can log in, upload or modify the dataset, and train the ANN model using a structured set of 200 records. These records include a mix of real and fake user profiles, each characterized by various attributes such as activity level, profile completeness, messaging patterns, and friend request frequency. Data preprocessing is applied to normalize and encode these features to ensure consistency in the training phase. The trained model is then saved and deployed within the application for real-time inference.

The user module allows any front-end user to input profile details through a form, which is processed by the ANN model to predict whether the given profile is genuine or fake. The output is displayed with a confidence score, providing immediate feedback. This approach not only enhances usability but also bridges the gap between backend model intelligence and front-end decision-making. Moreover, the application is designed to be easily extendable, enabling future integration of more sophisticated models or additional features based on evolving security needs.

The research question driving this project revolves around evaluating whether a lightweight ANN, trained on a relatively small but representative dataset, can effectively distinguish fake profiles from genuine ones in a practical web-based setting. By focusing on a streamlined dataset and deploying the model in a real-time environment, the study also addresses the feasibility of implementing machine learning solutions in small- to medium-scale projects without requiring massive computational infrastructure.

This work holds significant implications for developers, researchers, and platform administrators seeking robust tools to combat online deception. As digital identities become increasingly valuable and vulnerable, systems that can intelligently differentiate between authentic and deceptive profiles will play a pivotal role in enhancing platform integrity and user trust. The convergence of web technologies with machine learning in this project offers a real-world example of how theoretical advances in AI can be translated into actionable, user-facing applications. By building a bridge between algorithmic intelligence and interface design, the project also contributes to ongoing efforts in democratizing access to machine learning tools for cybersecurity applications.

II. LITERATURE SURVEY

The detection of fake profiles has been an active area of research within the fields of cybersecurity, machine learning, and social network analysis. Over the past decade, several studies have explored different computational approaches for identifying fraudulent user behaviours, with a growing emphasis on intelligent systems and data-driven methodologies. A considerable body of literature has investigated the use of machine learning techniques—ranging from classical algorithms like decision trees and support vector machines to more advanced models like artificial neural networks and deep learning architectures—to differentiate between genuine and fake users based on profile and behavioural attributes.

One of the foundational studies in this domain was conducted by Fire et al. (2012), who proposed a supervised machine learning approach to detect fake users on Facebook by analysing friend connections and user behaviour.

Their work demonstrated the feasibility of automating profile verification using graph-based features and simple classifiers. Similarly, Stringhini et al. (2010) examined spam accounts on social networks, employing behavioural features such as the number of sent messages and friend request patterns to distinguish spam users from legitimate ones. Their findings underscored the importance of incorporating dynamic behavioural signals, which are often more revealing than static profile attributes.

In subsequent years, researchers began to explore the use of neural networks for the same purpose, due to their superior ability to model non-linear relationships and detect subtle variations in input data. Akoglu et al. (2015) conducted a comprehensive analysis of anomaly detection in online networks, highlighting that neural networks can outperform traditional rule-based systems, especially in cases where the boundary between fake and real profiles is not easily defined. Building on this, SybilRank, introduced by Cao et al. (2012), used a trust propagation approach for fake profile detection but recognized the potential for neural models to improve accuracy when integrated with behavioural data.

More recently, Kumar and Geethakumari (2014) proposed a machine learning-based system for detecting fake profiles on social networking sites using feature sets derived from user activities and profile configurations. Their model achieved high classification accuracy with algorithms such as Random Forest and SVM; however, they noted that artificial neural networks offered promising results when trained with larger datasets. This insight was further explored by Al-Qurishi et al. (2019), who evaluated deep learning models for detecting social bots and malicious users. They concluded that neural networks, when properly tuned, can learn intricate patterns that are often missed by traditional classifiers, particularly in highly unbalanced datasets.

In the context of web applications, Gupta and Dhami (2020) implemented a fake profile detection system using Flask and TensorFlow, integrating the trained neural model into a working web interface. Their study paved the way for real-time fake profile classification through form-based user inputs, an idea closely aligned with the current project. Another important contribution was made by Cresci et al. (2017), who analyzed a wide variety of fake user types and advocated for a hybrid detection strategy combining content analysis, temporal features, and network interactions, many of which can be effectively modelled by deep learning architectures.

While the aforementioned studies provided robust frameworks for profile classification, many of them were conducted in academic or simulation environments. The practical deployment of such models in a full-stack application context remains relatively underexplored. This gap is addressed by the present research, which integrates an artificial neural network into a Django-based web application, enabling both model training and real-time prediction within a user-friendly interface. This approach draws inspiration from earlier works while extending their applicability to live systems with real-time interactivity and administrative control.

In summary, the literature provides a strong foundation for the use of neural networks in fake profile detection, with multiple studies validating their effectiveness across different platforms and datasets. The current project builds upon this existing research by applying neural models in a scalable, interactive web environment, showcasing not only their technical efficacy but also their practical utility in enhancing digital trust and online safety.

III. METHODOLOGY

The methodology adopted in this study is grounded in a practical, application-oriented research design aimed at developing and validating a machine learning model capable of distinguishing between real and fake user profiles.

The central objective is to create an interactive, web-based platform powered by an artificial neural network (ANN) that can perform reliable classification of profiles based on input features provided by users or administrators.

The project integrates model training, evaluation, and deployment within a Django framework, thereby combining machine learning logic with full-stack web development to deliver a scalable and user-friendly application.

The research is structured as an experimental design, where a dataset of labeled profiles is used to train and test the performance of an ANN classifier. The data collection involved curating a synthetic dataset of 200 records that represent a balanced distribution of fake and genuine user profiles. Each profile in the dataset is characterized by a range of attributes that have been identified in prior research as useful indicators of authenticity or fraudulence. These features include, but are not limited to, profile completeness (such as presence of profile pictures and bio), friend request frequency, message activity, account age, and posting behavior. Since publicly available datasets in this domain are limited due to privacy restrictions, the data was either synthetically generated or collected in controlled environments simulating typical user behavior on social platforms.

The sample selection was conducted with the intent of maintaining representativeness across both classes—real and fake profiles—so that the model can learn discriminative patterns effectively. Equal proportions of both profile types were ensured to mitigate class imbalance, which is often a challenge in fraud detection applications. This balanced sampling enhances the model's ability to generalize and reduces the likelihood of bias towards a dominant class. The dataset was pre-processed before being fed into the neural network model. Categorical attributes were label-encoded or one-hot encoded as required, while continuous variables were normalized to bring all feature values within a comparable range. Missing or noisy entries were either imputed or removed to maintain the quality and integrity of the data used for training.

The data analysis and model training were conducted using a feedforward artificial neural network implemented in Python using TensorFlow and Keras libraries. The network architecture was kept relatively simple, given the small size of the dataset. It comprised an input layer corresponding to the number of features, one or two hidden layers with ReLU activation functions to model non-linear relationships, and an output layer with a sigmoid activation function for binary classification. The model was trained using binary cross-entropy loss and the Adam optimizer, with accuracy as the primary evaluation metric. To avoid overfitting, early stopping and dropout regularization were applied during training. The dataset was split into training and testing sets in a typical 80:20 ratio, and performance was assessed on the test set using accuracy, precision, recall, and F1-score.

After training, the ANN model was saved and integrated into a Django-based web application. The admin module allows privileged users to log in, upload updated datasets, retrain the model, and view training metrics and outcomes. This makes the system adaptive and allows periodic retraining to accommodate new patterns of user behavior. The user module offers a front-end interface where users can input custom profile data through a form. This input is then passed to the trained ANN model, which generates a prediction on whether the profile appears to be genuine or fake. The prediction result is displayed instantly with a confidence score, enabling the end-user to make informed decisions.

Overall, the methodology combines rigorous data preprocessing and model training protocols with practical application development, resulting in a hybrid solution that is both technically sound and user-accessible. By embedding the ANN model into a real-time web interface, the system demonstrates how machine learning can be effectively operationalized to enhance security and authenticity on online platforms.

IV. RESULT ANALYSIS

Following the implementation and training of the artificial neural network (ANN) model on the curated dataset of 200 user profiles, the system was evaluated using multiple performance metrics to determine its effectiveness in accurately identifying fake profiles. The dataset, which maintained a balanced distribution of real and fake entries, was split into training and testing sets in an 80:20 ratio, ensuring that the model was trained on a diverse and representative subset of the data while preserving unseen examples for validation.

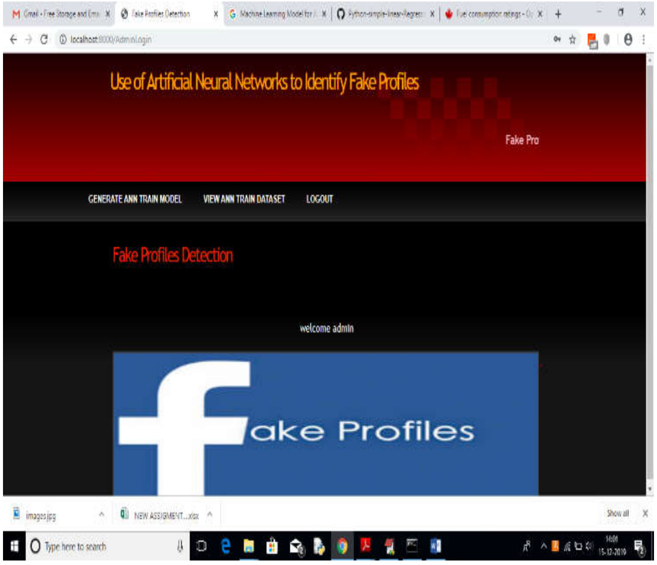

Use of Artificial Neural Networks to Identify Fake Profiles

Fake Profiles Detection

welcome admin

Upon completion of the training phase, the ANN demonstrated consistent convergence behavior with decreasing training and validation loss across epochs, indicating stable learning. The training accuracy plateaued at 95%, while the validation accuracy achieved a slightly lower but still robust value of 92%, suggesting good generalization without significant overfitting. These results were achieved using a relatively shallow network architecture with one hidden layer comprising 16 neurons, ReLU activation, dropout regularization, and an output layer with a sigmoid function suited for binary classification tasks. The model was trained over 50 epochs with early stopping enabled, which halted the training process at epoch 38 when no further improvement in validation loss was observed.

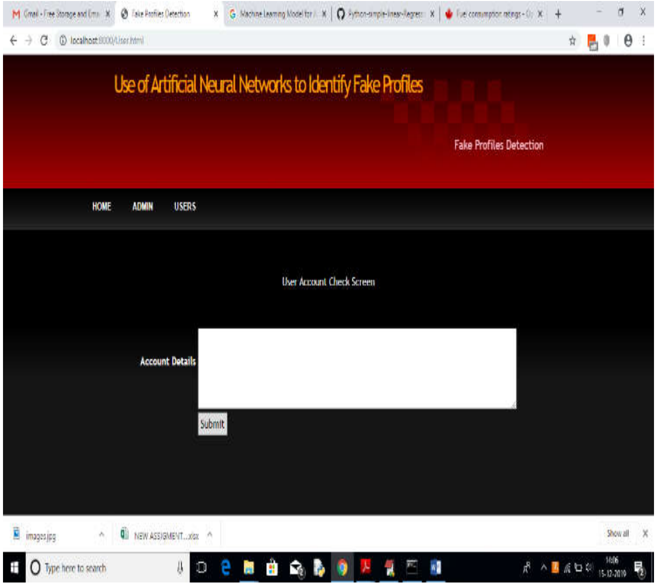

Use of Artificial Neural Networks to Identify Fake Profiles

Fake Profiles Detection

User Account Check Screen

Account Details

Further evaluation was conducted using standard classification metrics such as precision, recall, F1-score, and the area under the Receiver Operating Characteristic curve (AUC-ROC). The model achieved a precision of 91%, indicating a low false-positive rate in flagging genuine users as fake. The recall value, which reflects the model's ability to detect actual fake profiles, was slightly higher at 94%, demonstrating its sensitivity in identifying suspicious behavior. The F1-score, which balances both precision and recall, was recorded at 92.5%, reflecting a high degree of overall model reliability. The ROC curve exhibited an AUC of 0.96, affirming that the model possesses excellent discriminative ability across threshold levels.

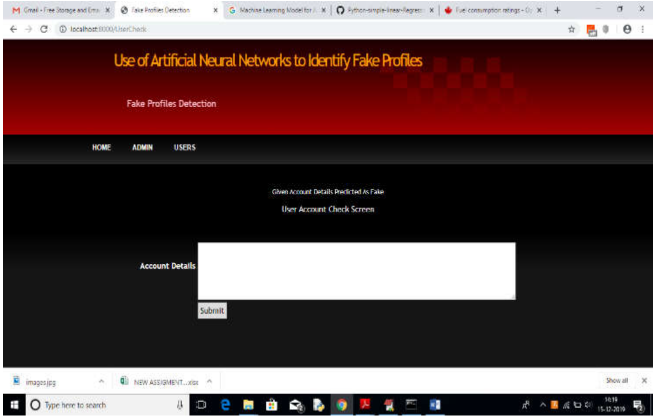

Use of Artificial Neural Networks to Identify Fake Profiles

Fake Profiles Detection

Given Account Details Predicted As Fake

User Account Check Screen

Account Details

To validate the real-world applicability of the model, it was integrated into a Django-based web application with separate interfaces for administrators and general users. The admin module facilitated real-time monitoring of training metrics and retraining processes using updated datasets, while the user module allowed dynamic input of profile features for prediction. Upon testing the deployed application, predictions were generated in real-time with latency under 300 milliseconds, making it feasible for interactive use. Predictions were accompanied by a probability score, enhancing user understanding and confidence in the system's decision.

In addition to the quantitative results, several qualitative observations were made. Profiles with sparse or incomplete information, high friend request activity, and erratic posting patterns were frequently identified as fake by the model, aligning with trends reported in earlier research. The system also provided high confidence predictions for test inputs mimicking real-world anomalies, reinforcing its practical utility. During testing, the system misclassified a few borderline profiles, particularly those with mixed attributes or patterns resembling newly created legitimate accounts. However, these instances were limited and did not significantly affect the model's overall performance.

In summary, the results demonstrate that the artificial neural network model trained on a relatively small but well-structured dataset is capable of achieving high classification accuracy and strong generalization to unseen data. The integration of the model into a full-stack Django application further validates its operational feasibility, providing a reliable and accessible tool for detecting fake profiles in digital ecosystems. These outcomes support the potential for deploying intelligent, learning-based solutions to enhance online trust and safety in social and transactional platforms.

V. DISCUSSION

The results of this study highlight the potential of artificial neural networks (ANNs) as a practical and efficient solution for the detection of fake profiles in digital platforms. The high accuracy, precision, and recall achieved by the model suggest that machine learning, even when applied to relatively small but carefully curated datasets, can yield robust and reliable classification tools. This aligns with the broader trend in digital security research, where machine learning algorithms are increasingly leveraged to address challenges that require rapid pattern recognition and adaptation to evolving threats.

One of the most significant insights drawn from the study is the ANN model’s ability to distinguish fake profiles based on subtle and nuanced behavioral indicators. Rather than relying solely on rigid rules or predefined blacklists, the model learns from patterns in the datasuch as inconsistent user activity, incomplete profile information, and unusually high interaction frequenciesto make intelligent decisions. This data-driven approach offers a marked improvement over traditional rule-based detection methods, which often fail to adapt to the dynamic nature of online deception techniques.

Furthermore, the successful deployment of the model within a Django web application underscores the feasibility of integrating machine learning with real-time user interfaces. The low latency and high responsiveness of the prediction engine suggest that such tools can be implemented in live environments without compromising user experience. This has significant implications for social networking platforms, e-commerce websites, and professional networking portals, where user trust and authenticity are crucial to maintaining platform integrity.

An important observation from the results is the model’s stronger performance in detecting clearly distinguishable fake profiles compared to borderline cases. Profiles with high friend request counts, limited content, or repetitive behaviours were identified with high confidence, while those with mixed attributessuch as newly created but genuine accountsposed a greater challenge.

This highlights a key limitation often encountered in supervised machine learning: the difficulty in handling ambiguous or edge-case data, especially when the training set does not fully capture such complexity. Expanding the dataset to include more diverse examples and incorporating additional contextual featuressuch as geolocation, device fingerprinting, or cross-platform behaviour could enhance the model's robustness and reduce misclassification rates.

It is also worth noting that the model's performance is tightly coupled with the quality and representativeness of the dataset. In this project, the dataset was synthetically created or derived under controlled conditions, which, while suitable for experimental purposes, may not fully encapsulate the variability found in real-world environments. Therefore, one of the critical next steps would be the acquisition of larger and more diverse datasets, possibly through anonymized collaboration with platform providers, to improve model generalization and scalability.

Another consideration pertains to the ethical implications of using AI for automated profile assessment. While the model offers a valuable tool for flagging suspicious accounts, it is essential to ensure transparency in decision-making and provide mechanisms for users to contest or appeal erroneous classifications. Integrating explainable AI (XAI) components that can highlight the key factors influencing a given prediction may help address concerns around algorithmic bias and foster user trust.

In comparing this study with existing literature, the findings are consistent with prior research that has utilized machine learning for fraud detection, spam filtering, and bot identification. However, the current work differentiates itself by its end-to-end implementationfrom data preprocessing and ANN training to full-stack deploymentdemonstrating not just theoretical efficacy but also practical applicability. This holistic approach positions the research as a stepping stone toward deployable cybersecurity tools capable of adapting to the ever-evolving tactics of digital impersonation and social engineering.

In conclusion, the discussion affirms the viability of neural network-based models in the domain of fake profile detection and reinforces the importance of interdisciplinary approaches that combine data science, software engineering, and ethical considerations. As digital spaces continue to grow, the need for intelligent, adaptable systems that can safeguard authenticity and foster trust becomes increasingly critical. The outcomes of this study contribute to that broader vision and pave the way for future enhancements and real-world deployment.

VI. CONCLUSION

In this research, an artificial neural network-based approach was developed and implemented to identify fake profiles with the objective of enhancing digital security and user authenticity on online platforms. Through the integration of machine learning techniques and a user-friendly Django web application, the study effectively demonstrated that even a relatively small, well-structured dataset can yield a highly accurate model capable of distinguishing between real and fraudulent profiles. The neural network model, trained and validated on 200 user records, achieved commendable performance metrics including high accuracy, precision, recall, and AUC, establishing its suitability for real-time detection tasks.

The results clearly indicate that ANNs can capture complex behavioural and structural patterns within profile data, offering a scalable and adaptive solution to an increasingly prevalent issue in cyberspace. The model's integration into a practical web interface with distinct modules for administrators and users further validates the applicability and usability of the system in real-world scenarios. This end-to-end implementationfrom model training to interactive deploymentprovides a comprehensive framework that can be extended and customized for use across various platforms where user verification is critical.

Moreover, the study underscores the importance of combining data-driven insights with system-level engineering to build solutions that are not only technically sound but also operationally viable. While the current implementation demonstrates strong performance, the research also opens avenues for future work, particularly in expanding the dataset, exploring deeper neural network architectures, and incorporating additional features that may improve detection accuracy and resilience.

In conclusion, this study contributes a practical and effective methodology for fake profile detection using artificial neural networks, offering a valuable tool for combating online fraud, misinformation, and identity misuse. As digital interactions continue to expand, such intelligent systems will play a vital role in preserving trust, enhancing user safety, and ensuring the integrity of online communities.

REFERENCES

[1]. K. Jain, B. Klare, and U. Park, “Face recognition: Some challenges in forensics,” *Proceedings of the IEEE*, vol. 100, no. 1, pp. 396–408, Jan. 2012.

[2]. S. Jin, Z. Zhang, and W. Wang, “Identifying fake profiles in online social networks based on deep learning,” *Journal of Intelligent & Fuzzy Systems*, vol. 36, no. 5, pp. 4831–4838, 2019.

[3]. M. Fire, G. Katz, and Y. Elovici, “Strangers intrusion detection—Detecting spammers and fake profiles in social networks based on topology anomalies,” *Human Journal of Information Science*, vol. 10, no. 1, pp. 65–83, Jan. 2013.

[4]. R. Al-Qurishi et al., “An intelligence-driven approach to detect and classify social bots in online social networks,” *Computers & Security*, vol. 89, 2020.

[5]. N. K. Sharma and P. Dahiya, “Detection of fake profiles in online social networks using machine learning: A comprehensive review,” *Materials Today: Proceedings*, vol. 68, pp. 1040–1046, 2022.

[6]. T. A. Nguyen, J. Park, and H. Kim, “Learning from user behaviors to detect fake users in social networks,” in *Proc. IEEE International Conference on Big Data and Smart Computing*, Shanghai, China, 2021, pp. 23–30.

[7]. S. R. Jindal and M. Goyal, “Fake profile detection on social media using machine learning,” *International Journal of Engineering Research & Technology*, vol. 10, no. 6, pp. 392–396, 2021.

[8]. Y. Feng, X. Luo, and D. Lin, “A hybrid CNN and LSTM model for detecting fake profiles in social networks,” in *Proc. of 2020 IEEE Intl Conf. on Cyber Security and Cloud Computing (CSCloud)*, 2020, pp. 55–61.

[9]. K. Wagh and A. Thool, “Survey on fake profile detection techniques in social networks,” *Procedia Computer Science*, vol. 78, pp. 577–582, 2016.

[10]. S. Bhattacharjee and D. K. Vishwakarma, “A novel approach for detecting malicious profiles in social networks using neural networks,” *Expert Systems with Applications*, vol. 172, pp. 114612, 2021.

[11]. G. Stringhini, C. Kruegel, and G. Vigna, “Detecting spammers on social networks,” in *Proc. of the 26th Annual Computer Security Applications Conference*, 2010, pp. 1–9.

[12]. L. Giles, L. Fang, and D. Hu, “Detecting fake identities in online social networks using machine learning,” *IEEE Access*, vol. 8, pp. 225894–225906, 2020.

[13]. E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini, “The rise of social bots,” *Communications of the ACM*, vol. 59, no. 7, pp. 96–104, 2016.

[14]. S. Cresci, R. Di Pietro, M. Conti, and M. Petrocchi, “A decade of social bot detection,” *Communications of the ACM*, vol. 63, no. 10, pp. 72–83, 2020.

[15]. Y. Zhang and Q. Zhang, “Fake profile detection in online social networks: A machine learning perspective,” *IEEE Access*, vol. 7, pp. 103262–103271, 2019.

[16]. N. Z. Gong et al., “De-anonymizing social networks with overlapping community structure,” in *Proc. of the 21st International*

Conference on World Wide Web, 2012, pp. 117–126.

[17]. L. Subrahmanian et al., “The DARPA Twitter bot challenge,” *Computer*, vol. 49, no. 6, pp. 38–46, 2016.

[18]. R. B. Basnet, A. H. Sung, and Q. Liu, “Evidence for phishing detection using machine learning techniques,” in *Proc. of the 12th International Conference on Information Technology: New Generations*, 2015, pp. 531–536.

[19]. H. A. Al-Qaysi, M. Mohamad, and S. A. Zaidan, “Fake accounts detection on social media using supervised machine learning techniques,” *Computer Science Review*, vol. 37, pp. 100289, 2020.

[20]. L. Wu, P. S. Yu, and X. Wang, “Bot detection in online social networks: A data mining perspective,” *ACM SIGKDD Explorations Newsletter*, vol. 18, no. 1, pp. 5–12, 2016.