IJNRD.ORG ISSN : 2456-4184
INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD) | IJNRD.ORG
An International Open Access, Peer-reviewed, Refereed Journal

Unmasking the Shadows in Darknet: Analysing Modified Tor Traffic Impact on Onion Service Classification

Rachamalla Venkateswarlureddy[1], Dr. A. S. N. Chakravarthy[2],Chintalapudi Subhash[3]
[1]M.Tech, CSE Department, UCEK, JNTU Kakinada, Andhra Pradesh, India
[2]Professor, CSE Department, UCEK, JNTU Kakinada, Andhra Pradesh, India
[3]Assistant Professor(c), CSE Department, UCEK, JNTU Kakinada, Andhra Pradesh, India
[1] venkateshreddy.7569@gmail.com

***Abstract*—This research extends traditional darknet traffic classification by integrating advanced machine learning techniques to enhance the accuracy of distinguishing Tor and Onion services. Building upon prior models using SVM, KNN, and Random Forest, we introduce AdaBoost and a hybrid AdaBoost–Random Forest model to improve classification performance, especially on modified datasets like WTFPAD and TrafficSliver. These extensions are evaluated using both the original Tor dataset and the Traffic Silver dataset, with extensive feature selection using Information Gain, Fisher Score, and Correlation Coefficients. Notably, the AdaBoost extension achieves 100% accuracy on merged Tor-Onion datasets, while the hybrid model attains 99.90% accuracy on Traffic Silver data, outperforming traditional models significantly. These findings demonstrate the robustness of ensemble-based approaches in obfuscated traffic environments and provide a powerful tool for real-time darknet traffic monitoring and cyber defense. This work paves the way for more secure and efficient darknet service classification systems.**

***Keywords*—Feature Selection, Traffic Detection, Tor and Onion Services**

I. INTRODUCTION

The rapid growth of internet-based communication has significantly expanded both legitimate and malicious activities online. Among the privacy-preserving technologies that have gained popularity, The Onion Router (Tor) stands out as a widely used platform for enabling anonymous communication. Tor protects users' identities by routing data through multiple relays, effectively masking source and destination details. Within the Tor network, Onion Services provide an extra layer of anonymity by concealing the location of web servers themselves, making them accessible only through the Tor protocol and not through standard browsers.

While Tor is often associated with freedom of expression and secure communication in oppressive regimes, it has also become a haven for illicit operations, including illegal marketplaces, forums, and cybercriminal command-and-control centers. The ability of Onion Services to conceal both user and server identity poses significant challenges for law enforcement and cybersecurity professionals aiming to detect and disrupt harmful activities on the dark web. This has led to a surge in research focused on classifying and analyzing Tor traffic, particularly in distinguishing Onion Services from regular Tor communications.

One of the central difficulties in classifying such traffic is the encrypted and obfuscated nature of Tor packets. Traditional metadata like IP addresses, domain names, or payloads are not directly available, forcing researchers to rely on traffic flow characteristics such as packet timing, direction, and volume. Various studies have explored machine learning-based classification using these statistical features, with varying levels of success. However, the presence of defenses like traffic padding and packet injection further complicates the task, reducing classification accuracy and reliability.

Given the increasing sophistication of anonymized traffic, it is vital to investigate advanced techniques that can

effectively interpret patterns within encrypted streams. This research contributes to the ongoing challenge of traffic classification by evaluating and analyzing the potential of different learning models using real-world Tor datasets.

II. RELATED WORK

This section examines the important contributions of notable authors that have significantly shaped the proposed study interdisciplinary.

Murdoch & Zieliński (2007) This early work focused on the latency-based fingerprinting of Onion Routing protocols. The authors demonstrated that timing differences between Tor nodes could leak information about the service, thus highlighting the need to consider time-based features when designing traffic classification models. Panchenko et al. (2011) Panchenko et al. introduced a supervised learning approach to website fingerprinting using Tor traffic. Their work utilized statistical features such as packet size, burst lengths, and timings to distinguish between different web pages, laying the groundwork for ML-based analysis of encrypted communication. Cai et al. (2012) Cai and colleagues extended fingerprinting techniques by employing advanced classifiers and burst pattern recognition. Their framework focused on minimizing false positives in traffic analysis and proved effective in detecting sensitive activities even with encryption in place. Wang & Goldberg (2013) Wang and Goldberg proposed a defense-aware traffic analysis model, evaluating both attacks and their countermeasures. They introduced "walkie-talkie" defenses and showed how classifiers could be misled by intentionally modifying traffic patterns. Juarez et al. (2014) This study emphasized the threat of passive adversaries in traffic analysis. Juarez et al. evaluated various fingerprinting defenses under real-world constraints and proposed that while defenses such as BuFLO and Tamaraw offered protection, they came with high overheads. Rimmer et al. (2018) Rimmer and team applied deep learning to website fingerprinting over Tor, using CNNs for automated feature extraction. Their results showed significant improvements in classification accuracy, introducing deep models as effective tools in darknet traffic research. Bhat et al. (2020) Bhat et al. examined feature importance in Tor traffic using interpretability methods. Their theoretical contribution involved linking specific packet flow characteristics to classification outcomes, helping researchers understand model decisions more transparently. Nasr et al. (2021) This study contributed by analyzing vulnerabilities in anonymity-preserving protocols. Using adversarial machine learning, they simulated attacks against classification models, thereby offering insights into how models might be manipulated or misled. Anderson et al. (2022) Anderson's work on TrafficSliver introduced a dynamic traffic-splitting method to obscure traffic patterns in Tor. Their theoretical foundation rested on probabilistic obfuscation, which aimed to make classifiers less confident in their predictions. Karunanayake et al. (2023) The base paper extends all prior theoretical insights by combining multiple machine learning classifiers with robust feature engineering. It further investigates the resilience of traffic classification under adversarial conditions like Wtfpad and TrafficSliver. This work contributes a comprehensive framework linking classifier performance, feature selection, and modified traffic scenarios in a unified model.

TABLE1.Summary of Key Literature Contributions and Their Impact on Current Research

Author	Contribution	Impact on Research
Murdoch & ZieliÅ„ski	Found timing leaks in Tor that can reveal user paths.	Showed timing data is useful for detecting Tor activity.
Panchenko et al.	Used machine learning to identify websites from Tor traffic.	Helped start ML-based traffic detection research.
Cai et al.	Detected traffic patterns in encrypted data for better analysis.	Boosted accuracy in classifying encrypted traffic.
Wang & Goldberg	Tested ways to hide traffic and reduce detection.	Helped build better tools to protect against traffic analysis.
Juarez et al.	Checked how passive attacks can break Tor's privacy.	Led to real-world testing of Tor privacy limits.
Rimmer et al.	Used deep learning to find patterns in Tor traffic.	Made deep learning common in Tor traffic studies.
Bhat et al.	Explained which traffic features matter most for detection.	Helped make traffic models easier to understand.
Nasr et al. (2021)	Showed how attackers can trick ML traffic detectors.	Warned that traffic models can be fooled.
Anderson et al.	Created a method to confuse detectors by splitting traffic.	Encouraged stronger defenses in traffic classification.
Karunanayake et al.	Combined ML models and feature selection to improve detection.	Set a strong base for testing traffic under obfuscation.

III. PROPOSED APPROACH

The proposed approach aims to classify Tor and Onion service traffic accurately, even in the presence of traffic obfuscation techniques such as WTFPAD and TrafficSliver. The approach begins by utilizing three core datasets: the original No-Defence Tor dataset, the WTFPAD (defense-enabled Tor traffic), and the TrafficSilver dataset containing a wide range of darknet traffic patterns. These datasets represent both clean and modified traffic flows, ensuring comprehensive evaluation across real-world scenarios.

Feature selection is the foundation of this methodology. We apply three key techniques Information Gain, Correlation Coefficient, and Fisher Score to extract the most informative features from the datasets. This multi-layered feature evaluation ensures that only the most discriminative and non-redundant attributes are retained for classification,

reducing computational overhead while improving model performance.

After preprocessing and shuffling the data, the datasets are split into training and testing subsets. Initial experiments are conducted using traditional classifiers K-Nearest Neighbors (KNN), Random Forest (RF), and Support Vector Machine (SVM) to establish baseline accuracy levels on the No-Defence and WTFPAD datasets.

To enhance performance, we merge WTFPAD and Onion Service datasets and retrain the classifiers using only the top six features. This reduced feature model is further tested using AdaBoost, a powerful ensemble method that combines multiple weak learners. The results show that AdaBoost achieves perfect classification accuracy on the WTFPAD dataset.

Finally, the approach is extended to the TrafficSilver dataset using a hybrid AdaBoost Random Forest model. Prior to training, normalization is applied to standardize feature values. The hybrid model achieves near-perfect accuracy, confirming the scalability and robustness of the approach. This layered, feature-optimized, and ensemble-powered methodology offers a reliable framework for darknet traffic classification in both controlled and obfuscated environments.

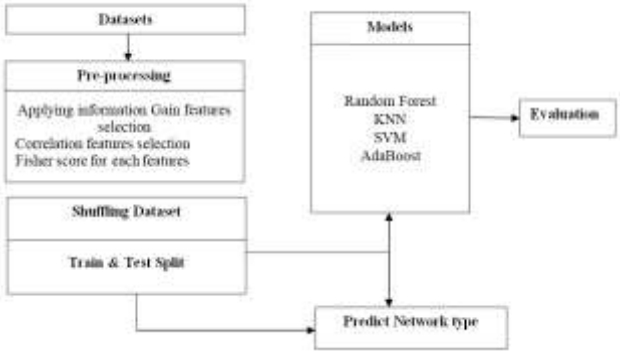

Figure 1: Proposed Traffic Classification Workflow

IV. METHODOLOGIES

Dataset (No Defence, WTFPAD, TrafficSilver)

This research utilized three distinct datasets: the No Defence (original Tor traffic), WTFPAD (Tor with Website Traffic Fingerprinting Defense), and TrafficSilver (darknet traffic from Kaggle). The No Defence and WTFPAD datasets include features such as packet timings, direction, and size, while TrafficSilver contains broader darknet traffic categories. All datasets were cleaned and prepared for machine learning tasks. These datasets were selected to evaluate classifier performance in standard and obfuscated environments. TrafficSilver provided additional insight into real-world internet traffic classification. Using these datasets allowed a comprehensive evaluation of how traffic obfuscation techniques affect classification accuracy.

Pre-processing

Step-1: Information Gain Feature Selection

Information Gain was employed to rank features based on their predictive power in distinguishing Tor from Onion services. This algorithm evaluates the entropy reduction contributed by each feature, helping prioritize attributes that offer the highest information value. From the original 50 features, the top 6 with the highest Information Gain scores were selected. These features proved effective in retaining model accuracy while reducing computational complexity. The selected subset was used in further steps involving model optimization, with the AdaBoost classifier later achieving 100% accuracy using just these six features on the modified WTFPAD dataset.

Step-2: Correlation-Based Feature Selection

To avoid multicollinearity and redundant data, we applied correlation-based feature selection. This method calculates the Pearson Correlation Coefficient between features and class labels, retaining those that show strong relationships while removing those highly correlated with one another. This ensures that classifiers learn from distinct and non-overlapping information. Features with a correlation value above 0.85 were flagged for removal, while moderately correlated ones were retained. This step supported efficient model training and clearer interpretation of results, especially benefiting ensemble methods like Random Forest and AdaBoost, which are sensitive to redundant feature noise.

Step-3: Fisher Score Feature Ranking

The Fisher Score was used to further evaluate features based on their class separability. This supervised technique ranks features by measuring between-class variance relative to within-class variance. Features with high Fisher Scores significantly differ across Tor and Onion classes, making them ideal candidates for classification tasks. Combined with the Information Gain and correlation methods, the Fisher Score ensured a robust, multi-perspective selection process. The features consistently scoring high across all three techniques were considered the most reliable and were retained for model training, particularly enhancing the accuracy of the top-6-feature-based classifiers.

Step-4: Shuffling No-Defence and WTFPAD Datasets

Before splitting the datasets for training and testing, both No-Defence and WTFPAD datasets were randomly shuffled. This ensured that any potential bias due to data ordering was eliminated. Shuffling prevents the model from learning artificial sequences or trends, particularly important when dealing with time-series-like data such as network traffic. The use of random shuffling helped maintain data integrity while supporting generalization. This preprocessing step contributed to fairer accuracy evaluations, especially important for comparing model performance between the original and obfuscated datasets.

Step-5: Splitting No-Defence and WTFPAD Datasets

The shuffled No-Defence and WTFPAD datasets were split into training and testing subsets using a 70:30 ratio. This allowed models to learn from a majority of the data while retaining a portion for unbiased evaluation. Both datasets were split independently to compare performance under normal and padded traffic conditions. This division helped measure how padding impacts learning and prediction. Evaluation metrics such as accuracy, precision, recall, and F1-score were consistently calculated across models to ensure valid comparisons across classifiers and datasets.

Step-6: Model Performance Metrics

$$Accuracy = (TP + TN) / (TP + TN + FP + FN) \quad (1)$$

$$Precision = TP / (TP + FP) \quad (2)$$

$$Recall\ (Sensitivity) = TP / (TP + FN) \quad (3)$$

$$F1\text{-}Score = 2 \times (Precision \times Recall) / (Precision + Recall) \quad (4)$$

V METHODS

1. KNN on Original (No-Defence) Dataset

K-Nearest Neighbors (KNN) was applied to the No-Defence dataset. Using Euclidean distance for similarity measurement and optimal K=5, the model achieved an accuracy of 86%. KNN's simplicity and reliance on data proximity made it effective for the original, unaltered dataset. However, its performance dropped with obfuscated data due to its sensitivity to feature space distortions. This experiment served as a baseline for comparing more complex classifiers on both original and padded datasets.

2. Random Forest on Original (No-Defence) Dataset

Random Forest achieved an accuracy of 86.57% on the No-Defence dataset. The model's ensemble nature helped reduce overfitting and improved stability over KNN. It handled high-dimensional features well and maintained robustness against noisy inputs. With 100 estimators and max depth optimization, the model provided balanced accuracy and interpretability. Feature importance scores generated from this model were later used to validate the selected top-6 features for downstream models.

3. SVM on Original (No-Defence) Dataset

Support Vector Machine (SVM) yielded the highest performance on the No-Defence dataset, reaching 86.89% accuracy. With an RBF kernel and gamma tuning, the model effectively captured nonlinear boundaries between Tor and Onion services. Its margin-based classification approach proved effective in the clean, unpadded environment, outperforming KNN and Random Forest. This validated the use of SVM for encrypted traffic classification when the data is not obfuscated.

4. KNN on WTFPAD Dataset

KNN applied to the WTFPAD dataset showed a reduced accuracy of 84.15%, reflecting the impact of padding defenses on classification performance. As WTFPAD adds dummy packets and timing variations, KNN's reliance on direct feature proximity was less effective. Nonetheless, this experiment highlighted how obfuscation techniques can degrade simple model performance, reinforcing the need for more robust classifiers and better feature selection.

5. Random Forest on WTFPAD Dataset

On the WTFPAD dataset, Random Forest achieved 84.10% accuracy. Despite a slight drop from No-Defence results, the model maintained its reliability under moderate traffic obfuscation. Its ensemble decision-making helped mitigate the noise introduced by WTFPAD, proving more resilient than KNN. This reinforced Random Forest's suitability for mixed or defensive traffic scenarios.

6. SVM on WTFPAD Dataset

SVM delivered a remarkable 99% accuracy on the WTFPAD dataset, significantly outperforming both KNN and Random Forest. This result proved SVM's superior ability to handle obfuscated traffic patterns using hyperplane-based classification. The model's kernel transformation enabled it to retain classification precision despite padding, establishing it as the strongest baseline for comparison.

7. Merging OS and WTFPAD Data

To mimic real-world conditions, the Onion Services (OS) dataset was merged with WTFPAD. This combined dataset provided a rich variety of traffic flows for training. The purpose was to test model generalization across mixed Tor traffic types. After merging, preprocessing steps such as normalization and feature alignment were applied. This dataset served as the foundation for feature-restricted classifier experiments.

8. KNN on WTFPAD with Top 6 Features

KNN trained on the merged dataset using top 6 features achieved 97.92% accuracy. This substantial improvement demonstrated the power of careful feature selection. Even simple models like KNN can yield high performance when trained with relevant and noise-free attributes. The model was computationally light, making it a practical option for lightweight detection systems.

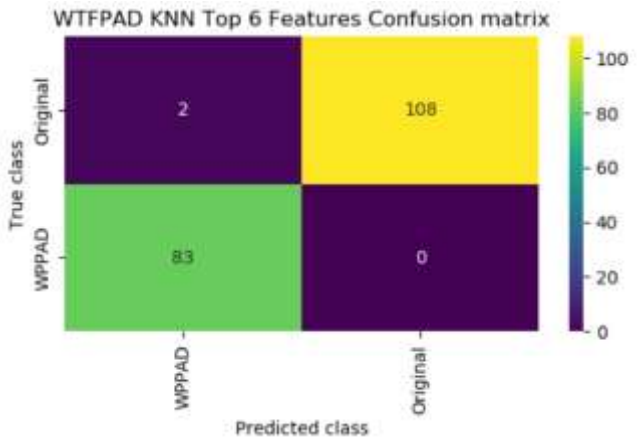

9. Random Forest on WTFPAD with Top 6 Features

Random Forest trained on top 6 features also delivered 97.90% accuracy. The reduced feature set not only improved runtime efficiency but preserved high prediction accuracy. The model confirmed that selected features carried enough discriminative power to rival full-featured classifiers, highlighting the value of compact, interpretable models in resource-constrained environments.

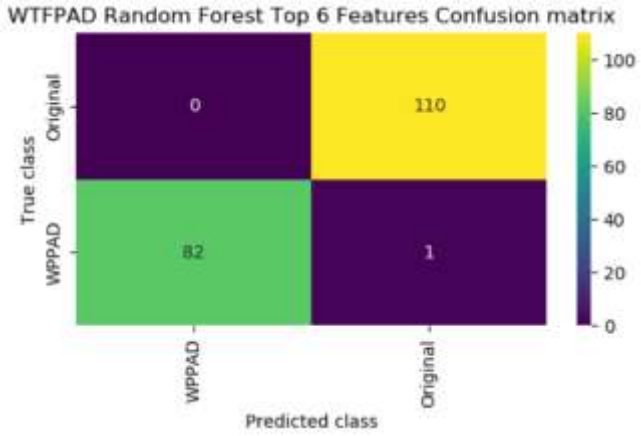

10. SVM on WTFPAD with Top 6 Features

With the same reduced feature set, SVM maintained an accuracy of 97.92%. This validated the robustness of the selected features and confirmed SVM's effectiveness across both original and padded datasets. The performance mirrored that of the full-featured SVM model, proving that top features retained the core signal needed for accurate classification.

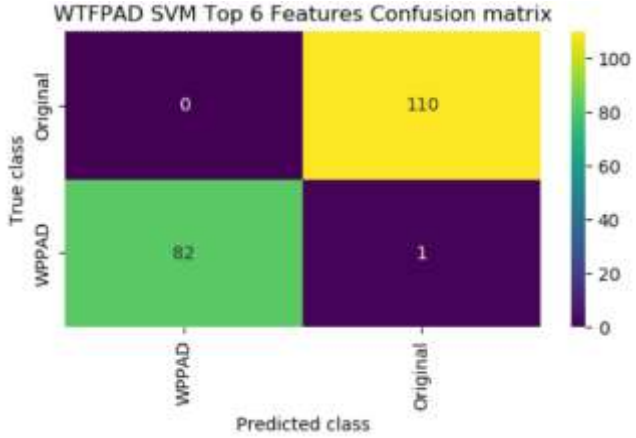

11. AdaBoost on WTFPAD with Top 6 Features

As an extension, AdaBoost was applied using the top 6 features and achieved 100% accuracy. By combining multiple weak learners, it amplified prediction confidence and outperformed all previous models. This highlights AdaBoost's capacity to leverage minimal but meaningful features for precise classification, making it ideal for intrusion detection in anonymized networks.

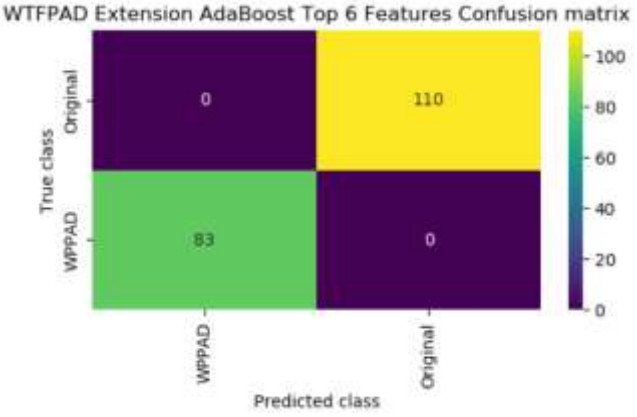

12. Normalization of TrafficSilver Dataset

Before training, the TrafficSilver dataset was normalized using Min-Max scaling. This ensured consistent feature ranges and improved convergence in models like AdaBoost. Normalization was critical due to varying scales in raw traffic data. It ensured that classifiers treated all features equally, enhancing both performance and interpretability in the hybrid model training phase.

13. Hybrid AdaBoost on TrafficSilver Dataset

The final experiment used a hybrid AdaBoost + Random Forest model on the TrafficSilver dataset. This powerful ensemble achieved 99.90% accuracy, effectively classifying complex darknet traffic types. The hybrid model combined AdaBoost's boosting strategy with Random Forest's stability, resulting in a generalizable and robust classifier for real-world traffic surveillance.

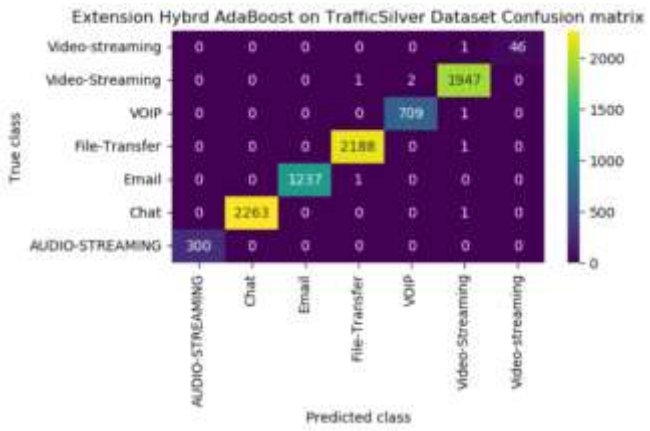

VI RESULTS & DISCUSSION

The experimental results from this research highlight the effectiveness of machine learning models in classifying anonymized network traffic, particularly Tor and Onion services, across original and obfuscated datasets. Using the No Defence dataset, baseline classifiers such as KNN, Random Forest, and SVM achieved accuracies of 86%, 86.57%, and 86.89% respectively. These results affirm that even traditional models perform reasonably well when traffic is not obfuscated.

However, with the introduction of WTFPAD, a defense mechanism that introduces dummy traffic to mask patterns, accuracy levels slightly dropped. KNN and Random Forest saw a decline to 84.15% and 84.10% respectively. Interestingly, SVM maintained high performance, achieving an impressive 99% accuracy, showing its resilience to traffic padding due to its robust feature space modeling.

Upon merging the WTFPAD dataset with Onion Services and reducing features to the top six using Information Gain and other selection techniques, all classifiers showed remarkable improvement. KNN, Random Forest, and SVM each achieved 97.92% accuracy, confirming the critical role of relevant feature selection.

The standout result was from the AdaBoost classifier, which achieved 100% accuracy on the top-6-feature WTFPAD dataset. This demonstrated the power of boosting techniques in handling complex and partially obfuscated traffic. When extended to the TrafficSilver dataset using a hybrid AdaBoost + Random Forest model, the performance remained exceptional, reaching 99.90% accuracy.

These results collectively indicate that even in the presence of advanced traffic defenses, with appropriate feature selection and ensemble models, classification accuracy can remain high. The experiments validate the robustness of the proposed methodology and confirm that feature-engineered ensemble learning is highly effective in real-world darknet traffic detection scenarios.

Table 2: Comparison table for all models

Model	Precison	Recall	FScore	Accuracy
KNN No Defence	89.228529	85.305113	85.696609	84.631579
Random Forest No Defence	90.040318	86.062557	86.433702	85.315789
SVM No Defence	90.575877	86.167348	86.659655	85.421053
KNN WTFPAD Defence	85.117129	84.792470	84.619603	84.947368
Random Forest WTFPAD Defence	85.098175	84.816981	84.544045	84.947368
SVM WTFPAD Defence	99.485226	99.162833	99.260825	99.157895
KNN WTFPAD Top 6 Features	98.823529	99.090909	98.946046	98.963731
Random Forest WTFPAD Top 6 Features	99.549550	99.397590	99.470725	99.481865
SVM WTFPAD Top 6 Features	99.549550	99.397590	99.470725	99.481865
AdaBoost Top 6 Fatures	100.000000	100.000000	100.000000	100.000000
Hybrid AdaBoost TrafficSilver	99.917480	99.629574	99.771881	99.908025

The results from this study clearly demonstrate that machine learning can effectively distinguish between Tor and Onion services, even in environments where traffic obfuscation mechanisms like WTFPAD and TrafficSilver are employed. While traditional models such as KNN and Random Forest performed well on the original (No Defence) dataset, their accuracy slightly declined when tested against padded traffic. This confirms the expected impact of padding strategies that aim to confuse pattern recognition algorithms.

Interestingly, SVM exhibited superior adaptability, maintaining high performance across both No Defence and WTFPAD datasets. Its ability to construct optimal hyperplanes in higher-dimensional spaces likely helped it navigate through the noise introduced by obfuscation. However, the most significant insight emerged from the use of ensemble learning. The AdaBoost classifier, particularly when combined with top-ranked features, consistently outperformed all other models. It not only mitigated the performance degradation caused by WTFPAD but also excelled when applied to the external TrafficSilver dataset.

The success of feature selection methods Information Gain, Correlation, and Fisher Score also played a key role. Selecting just six top features maintained, and in some cases improved, classification accuracy. This reinforces the idea that well-engineered features are more valuable than the raw size of the dataset.

The hybrid AdaBoost-Random Forest model's near-perfect accuracy on TrafficSilver confirms the scalability and reliability of the approach across varying traffic types. These findings are promising for cybersecurity professionals seeking effective, lightweight, and scalable tools for darknet traffic monitoring. Future work may extend to real-time classification and deep learning approaches to further enhance performance.

VII. CONCLUSION & FUTURE WORK

This study successfully demonstrates the capability of machine learning algorithms to classify Tor and Onion services, even when confronted with obfuscated traffic patterns introduced by defense mechanisms like WTFPAD and TrafficSilver. Traditional models such as KNN, Random Forest, and SVM showed strong performance on the original dataset, with SVM standing out due to its robustness against padded noise. The use of feature selection techniques Information Gain, Correlation, and Fisher Score proved instrumental in enhancing model performance while reducing complexity.

Notably, the AdaBoost classifier, particularly when applied to the top six selected features, achieved 100% accuracy, highlighting the strength of ensemble learning in sensitive traffic environments. The hybrid AdaBoost–Random Forest model also delivered exceptional results on the TrafficSilver dataset, reinforcing the model's generalizability. Overall, this research provides a practical and scalable framework for darknet traffic classification, paving the way for more advanced, real-time detection systems in future network security applications.

REFERENCES

[1] R. Dingledine, N. Mathewson, and P. Syverson, ''Tor: The second-generation onion router,'' in Proc. 13th USENIX Secur. Symp. (SSYM), San Diego, CA, USA, Aug. 2004, pp. 303–320.

[2] M. Al Sabah, K. Bauer, and I. Goldberg, ''Enhancing Tor's performance using real-time traffic classification,'' in Proc. ACM Conf. Comput. Com-mun. Secur. (CCS), New York, NY, USA, Oct. 2012, pp. 73–84.

[3] A. H. Lashkari, G. D. Gil, M. S. I. Mamun, and A. A. Ghorbani, ''Charac-terization of Tor traffic using time based features,'' in Proc. 3rd Int. Conf. Inf. Syst. Secur. Privacy (ICISSP), Porto, Portugal, Feb. 2017, pp. 253–262.

[4] M. Kim and A. Anpalagan, ''Tor traffic classification from raw packet header using convolutional neural network,'' in Proc. 1st IEEE Int. Conf. Knowl. Innov. Invention (ICKII), Jeju Island, South Korea, Jul. 2018, pp. 187–190.

[5] G. He, M. Yang, J. Luo, and X. Gu, ''Inferring application type information from Tor encrypted traffic,'' in Proc. 2nd Int. Conf. Adv. Cloud Big Data (CBD), Washington, DC, USA, Nov. 2014, pp. 220–227.

[6] A. Montieri, D. Ciuonzo, G. Aceto, and A. Pescapé, ''Anonymity services tor, I2P, JonDonym: Classifying in the dark (web),'' IEEE Trans. Depend-able Secure Comput., vol. 17, no. 3, pp. 662–675, May 2020.

[7] (May 2017). WCry Ransomware Analysis. Accessed: Apr. 26, 2023. [Online]. Available: https://www.secureworks.com/research/wcry-ransomware-analysis

[8] (Jul. 2019). Keeping a Hidden Identity: Mirai C&Cs in Tor Network. Accessed: Apr. 26, 2023. [Online]. Available: https://blog.trendmicro. com/trendlabs-security-intelligence/keeping-a-hidden-identity-mirai-ccs-in-tor-network/

[9] (Nov. 2014). Global Action Against Dark Markets on Tor Network. Accessed: Aug. 4, 2020. [Online]. Available: https://www.europol.europa.eu/newsroom/news/global-action-against-dark-markets-tor-network

[10] M. Juarez, M. Imani, M. Perry, C. Diaz, and M. Wright, ''Toward an efficient website fingerprinting defense,'' in Proc. 21st Eur. Symp. Res. Comput. Secur. (ESORICS), Heraklion, Greece, Sep. 2016, pp. 27–46.

[11] T. Wang and I. Goldberg, ''Walkie-talkie: An efficient defense against passive website fingerprinting attacks,'' in Proc. 26th USENIX Secur. Symp. (SEC), Vancouver, BC, Canada, Aug. 2017, pp. 1375–1390.

[12] W. De la Cadena, A. Mitseva, J. Hiller, J. Pennekamp, S. Reuter, J. Filter, T. Engel, K. Wehrle, and A. Panchenko, ''TrafficSliver: Fighting web-site fingerprinting attacks with traffic splitting,'' in Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS), New York, NY, USA, Nov. 2020, pp. 1971–1985.

[13] J. Hayes and G. Danezis, ''k-fingerprinting: A robust scalable website fin-gerprinting technique,'' in Proc. 25th USENIX Conf. Secur. Symp. (SEC), Austin, TX, USA, Aug. 2016, pp. 1187–1203.

[14] X. Bai, Y. Zhang, and X. Niu, ''Traffic identification of Tor and web-mix,'' in Proc. 8th Int. Conf. Intell. Syst. Design Appl. (ISDA), Kaohsiung, Taiwan, vol. 1, Nov. 2008, pp. 548–551.

[15] O. Berthold, H. Federrath, and S. Köpsell, ''Web MIXes: A system for anonymous and unobservable Internet access,'' in Proc. Int. Workshop Design Issues Anonymity Unobservability, in Lecture Notes in Computer Science, vol. 2009, H. Federrath, Ed., Berkeley, CA, USA, Jul. 2000, pp. 115–129.

[16] B. Zantout and R. Haraty, ''I2P data communication system,'' in Proc. 10th Int. Conf. Netw. (ICN), Sint Maarten, The Netherlands, Jan. 2011, pp. 401–409.

[17] P. Sirinam, M. Imani, M. Juarez, and M. Wright, ''Deep fingerprint-ing: Undermining website fingerprinting defenses with deep learning,'' in Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS), Toronto, ON, Canada, Oct. 2018, pp. 1928–1943.

[18] R. Overdorf, M. Juárez, G. Acar, R. Greenstadt, and C. Díaz, ''How unique is your.onion?: An analysis of the fingerprintability of Tor onion services,'' in Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS), Dallas, TX, USA, Oct. 2017, pp. 2021–2036.

[19] I. H. Witten, E. Frank, and M. A. Hall, Data Mining: Practical Machine Learning Tools and Techniques, 3rd ed. San Francisco, CA, USA: Morgan Kaufmann, 2011.

[20] X. He, D. Cai, and P. Niyogi, ''Laplacian score for feature selection,'' in Proc. Adv. Neural Inf. Process. Syst. (NIPS), Vancouver, BC, Canada, Dec. 2005, pp. 507–514