Unmasking the Phishermen: Phishing Domain Detection with Machine Learning and Multi-Source Intelligence

Radek Hranický
Faculty of Information Technology
Brno University of Technology
Brno, Czechia
hranicky@fit.vut.cz

Adam Horák
Faculty of Information Technology
Brno University of Technology
Brno, Czechia
ihorak@fit.vut.cz

Jan Polišenský
Faculty of Information Technology
Brno University of Technology
Brno, Czechia
xpolis04@stud.fit.vut.cz

Kamil Jeřábek
Faculty of Information Technology
Brno University of Technology
Brno, Czechia
ijerabek@fit.vut.cz

Ondřej Ryšavý
Faculty of Information Technology
Brno University of Technology
Brno, Czechia
rysavy@fit.vut.cz

***Abstract*—In the digital landscape, phishing attacks have rapidly evolved into a major cybersecurity challenge, posing significant risks to individuals and organizations. This short paper presents our preliminary research on detecting phishing domains. Our approach amalgamates intelligence from multiple sources: DNS servers, WHOIS/RDAP, TLS certificates, and GeoIP data. We created a rich 15.8 GB dataset of information about benign and phishing domains, from which we derived a comprehensive 80-feature vector for training and testing machine learning classifiers. We propose preliminary results with a fine-tuned XGBoost model, achieving 0.9716 precision rate, 0.9540 F-1 score, and false positive rate of 0.23%.**

***Index Terms*—Phishing, Domain, Detection, Machine learning, XGBoost, Features, DNS, RDAP, TLS, GeoIP**

I. Introduction

In today's digital era, phishing has become a significant cybersecurity concern. The widespread use of e-mail and social media has created a vast arena for phishers, allowing them to target masses of unsuspected users easily. Many individuals and organizations underestimate the sophistication of modern phishing campaigns, leading to catastrophic consequences. Those include data breaches, identity theft, substantial financial loss, or data damage due to frequent ransomware attacks.

Various countermeasures have been studied and deployed over the years. Educating users helps, but it is not a panacea [1]–[3]. Even a knowledgeable user might be caught off guard and tricked by a sophisticated campaign. Hence, additional systematic techniques based on allow/deny lists [4]–[7] and heuristics [8], [9] are deployed. Specialized crafted blocklists or heuristics might effectively cover known already reported suspicious websites but are weak against zero-hour attacks [10]. Therefore, researchers utilize machine learning to propose solutions capable of uncovering new potentially suspicious assets based on the known data patterns from the past [11]–[13].

Works proposing solutions to fight phishing analyze different data sources, including e-mail, instant messaging, social networks, advertisements, websites, or domain names. Domain names are an indispensable part of user interaction on the Internet, and they usually represent the first action triggered by the user willing to visit any website. A wide range of publicly available information is obtainable from the DNS system, WHOIS/RDAP, IP geolocation databases, or information about deployed TLS certificates. Therefore, our research focuses on utilizing the information from all listed data sources to train a machine-learning classifier capable of distinguishing benign and suspicious phishing domains. This paper presents the preliminary results of the proposed solution.

Our contribution is threefold. At first, we created and published a dataset containing 15.8 GB of domain-related data. Secondly, we designed a feature vector for accurate classification. And thirdly, we propose a classifier and present preliminary achieved results. The novelty of this study lies in combining five information sources, that as far as we know were not used together in any previous study, unique crafted feature vector and proposing highly accurate phishing domain detection method.

II. Related Work

Throughout time, network administrators and cybersecurity experts have sought various methods to combat phishing. Sheng et al. [1] demonstrated that anti-phishing training reduced phishing susceptibility by 40%, yet participants still succumbed to 28% of phishing attempts. Görling [2] noted users often bypass security models that interfere with their tasks, and the Anonymous attack on HBGary [3] underscores that even experts are not immune. Unsurprisingly, there is a high demand for network forensics and incident response professionals to address such incidents [14]. Countermeasures like allow/deny lists [10] or DNSBL [6] are limited to exact match scenarios with the domain name or URL. Even predictive solutions like PhishNet [7] cannot withstand zero-hour attacks. Heuristic-based solutions like the AntiPhish [9], Phishwish [8], and IceShield [15] browser extensions face challenges of high false positive rates and require periodic updates [10]. Alternative visual similarity-based anti-phishing schemes, like those proposed by Chen et al. [16] and Hara et al. [17], provide satisfactory results but are computationally demanding and unusable on a larger scale [10].

To combat zero-hour attacks, machine learning models can be trained to detect phishing e-mails [11]–[13], domains [18]–[25], URLs [26]–[28], or webpage contents [12], [29]–[31].

Chandrasekaran et al. [11] detected phishing e-mails with SVM using 25 features, including keyword search for terms like ”account” or ”security”, resulting in 95% detection rate. Fette et al. [12] introduced the PILFER algorithm for detecting both phishing e-mails and websites with a 0.96 accuracy using Random Forest. Abu-Nimeh et al. [13] compared multiple classification methods, achieving precision up to to 0.9511.

While HTML/DOM analysis can identify phishing web pages [29]–[31], this method requires computationally intensive full-page scraping and processing.

To detect malicious domains, lexical features can be derived from the domain name itself. Drichel et al. [18] validated this method by examining 136 unique lexical features to identify C&C domains. Counting characters of the given type, analyzing length and subdomain level is also usable for detecting malicious URLs, as shown by Do et al. [27] and Patgiri et al. [28] who also utilized IP reputation data.

DNS is vital for domain credibility assessment. Bilge et al. [19] identified phishing and botnet domains through passive DNS traffic analysis. While we do not analyze network traffic directly, we adopted several features like the count of IP addresses, countries, and numerical characters. Perdisci et al. [32] similarly used passive DNS analysis with unique IP-based features. Antonakakis et al. [20] used statistical features from BGP prefixes, AS numbers, and more, incorporating requester IP distributions and malware-related reputation scores.

Potential indicators of domain maliciousness are also evident in TLS certificates. Research by Hageman et al. [21] highlighted that in Q4 2020, 84% of detected phishing attacks occurred over HTTPS, relying on a relatively small group of certificate issuers. Torroledo et al. [22] achieved 0.8963 precision rate in detecting phishing and malware domains through 30 TLS-based features.

Lexical, DNS, and TLS-based features have proven effective, and we believe combining these sources might yield even better results. Kuyama et al. [23] detected malicious domains using 9 WHOIS and 8 DNS-based features. Shi et al. [24] extended this by adding three lexical and two IP-based features. Nevertheless, the two approaches focused mainly on malware/C&C domains. Hason et al. [33] identified phishing and C&C domains using 9 features categorized as non-robust, semi-robust, and robust. Consecutive characters were seen as a robust feature, while attributes like domain entropy, lifetime, activity, and the number of distinct IP addresses were deemed semi-robust. Chatterjee et al. [34] achieved 0.867 precision using 14 features, such as the HTTPS presence and domain age. Sadique et al. [25] blended lexical, host-based, WHOIS, and GeoIP features, reaching 0.87 accuracy, with lexical features dominating the top 10 list according to importance. However, they did not utilize DNS or TLS data.

With the exception of Sadique et al. [25], the majority of existing ML-based phishing detection techniques used data from one or two sources (e.g., WHOIS and DNS). To see how detection works across a broader spectrum, we focused our research on combining lexical attributes and five different external information sources. Additionally, many ML-driven malicious domain detection efforts have barely achieved precision over 0.90, indicating a significant room for improvements. We also noted most phishing-detection attempts aim at malicious web contents, URLs, or e-mails. In contrast, domain-centric analyses were primarily used for malware and botnet domains. Inspired by this observation, we directed our efforts into exclusively domain-based phishing detection.

III. Phishing Domain Detection Methodology

Figure 1 outlines the whole methodology for creating our phishing domain classifier. The process is split into three core steps: A) creating a suitable dataset with domain-related data, B) extracting convenient features for phishing detection, and C) training the classifier.

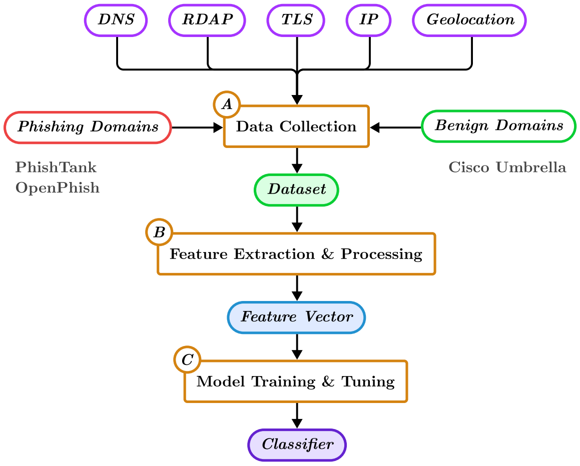

Fig. 1. A holistic overview of the classifier creation.

A. Data Collection

Training effective classifiers require a sufficiently large dataset of high quality and granularity. We started by creating lists of benign (trusted) domains and phishing domains.

To obtain benign domains, we chose the Cisco Umbrella [35] platform that protects organizations on the DNS level and provides vast allow/deny lists of Internet domains. To collect only the most likely benign domains and thereby ensure the dataset's quality, we performed a filtering process on the available lists similar to Rahbarinia et al. [36]. First, we took monthly data for the top one million domains across the past year. We then collected only the domains that consistently appeared in these top results. We gathered a list of 432,572 benign domains from Cisco Umbrella using this technique.

The phishing domains were collected from OpenPhish [37], a fully automated platform for phishing intelligence that keeps track of live malicious hostnames and URLs, and PhishTank [38], operated by Cisco Talos Intelligence Group. Both services verify the reported domains, minimizing the ocurrence of false positives, which is essential for a quality dataset. From these two sources, accessed via the MISP platform [39], we obtained a list of 36,993 verified phishing domains.

For each domain, we collected DNS records, as they provide valuable information for malicious domain detection, such as the domain's name servers, mail servers [23], and related IP addresses [19], [24]. For each IP address, we also measured the Route-trip time (RTT) with ICMP Echo. Our next source, RDAP, a WHOIS replacement, provides domain registration information that may be linked to known phishing activity [24], [33]. Further, we downloaded the domains' TLS certificate chains to search for discrepancies and suspicious information [22]. Additionally, we collected domain-related GeoIP data to identify whether a domain's IP addresses are located in high-risk regions or linked to recognized phishing activities. Finally, we published the complete dataset on Zenodo [40].

B. Feature Extraction and Processing

In related work, we identified 183 unique domain-related features. We then excluded those unsuitable for our scope (e.g., not extractable from our data) and those deemed irrelevant for phishing detection in prior research. Conversely, we crafted manually a number of novel, potentially helpful features, for instance, related to TLS certificate extensions and policies. This approach resulted in 106 features for manual examination.

Through result analysis and examining feature importance and correlation, we further refined our list. We then eliminated features that did not contribute meaningfully to identifying phishing domains or those duplicating information to prevent computational strain or result distortion. Our final feature vector comprised 80 features divided into several categories. The entire feature vector is also described on Zenodo [40].

The first group comprises 22 lexical features (prefixed with *lex_*) extracted from the sole domain name. Those included the length and normalized entropy of the domain name, ratios of digits, vowels, and consonants. Such features have effectively detected C&C DGA domains in previous works [18], [24]. Observing possible algorithmically generated examples in our phishing list, like "help-center267.crabdance.com", we assessed similar lexical features for phishing detection. We also counted the presence of 47 common clickbait words like "account" or "free" which we discovered in phishing domain names. Finally, we created statistics of the most frequent {2-5}-grams amongst phishing domain names and calculated their occurrence within the domain name.

Another group of 25 DNS-related (prefixed with *dns_*) features included counts for specific record types, as Kuyama et al. proved their usability for malicious domain detection [23]. As lexical approaches showed success [18], [24], [25], we applied these to DNS-derived strings. We thus measure the length, subdomain level, digit count, and mean entropy of the primary nameserver and the administrator's e-mail address. We also assessed the entropy and average length of MX names and TXT record values. Furthermore, we flagged SPF and DKIM presence as potential domain credibility indicators.

From DNS records, we also extracted information about related IP addresses and crafted five additional IP-based features (prefixed with *ip_*) like the total counts of IP, IPv4, and IPv6 addresses, usable in previous detection attempts [19], [20], [24]. We also included the average entropy of IP prefixes, as Perdisci et al. [32] suggested low IP diversity is often an indicator of high-flux malicious domains. Lastly, we included the average RTT for all domain-related IP addresses, as we suppose credible services might exhibit reduced latencies, particularly when regionally aligned.

Examining TLS certificate chains resulted in 17 TLS-related features (prefixed with *tls_*). As the classifiers of Torolledo et al. [22] showed a high precision rate, we adopted some of the features, for instance, extension count or the validity length of certificates. Additional features like root certificate validity or numbers of certificates fulfilling ISO/Joint ISO-ITU policies were identified through manual examination.

We derived seven features from WHOIS/RDAP data for the domain and related IP addresses. (prefixed with *rdap_*), including the domain's age, underpinned by our hypothesis that longer-standing domains are generally more trustworthy. We also incorporated the time since the last change, positing that the nature and timing of updates might signal a domain's legitimacy. Additionally, we included a flag indicating the support for DNSSEC as a potential marker for a domain's credibility. The domain registration period and the domain active time were also included, as previous studies have underscored their usability [24], [33], [41]. Other features cover information about the registrar and the average length of the admin contact name for IP addresses.

Finally, we added three features associated with geolocation information for domain-related IP addresses (prefixed with *geo_*). We included the total number of distinct countries, a metric previously highlighted for its utility in malicious domain detection [19], [20], [24], [33]. Recognizing that specific phishing campaigns might be orchestrated by blackhat groups operating within distinct sets of countries or continents, we further enriched our feature vector with: a unique hash for each combination of countries, and a unique hash for each combination of continents where the countries were situated.

C. Model Training and Tuning

For preliminary verification of the feature vector's usability, we employed the XGBoost classifier [42] due to its fast training capabilities, easy interpretability of results, general well performance on imbalanced data and remarkable performance on tabular data, further confirmed in several Kaggle competitions [43].

We conducted the Train-Test split with 70% of data reserved for training and 30% for model validation. On the training part, we applied additional stratified 5-fold cross-validation for grid search-based model tuning. We ended with 998 decision trees with a maximum depth of 11, training rate of 0.10 to reach the "binary:logistic" objective with GPU histogram algorithm and gradient-based sampling method. We also set *min_child_weight* to 1.1, *subsample* to 0.9 and *max_bin* to 2048. The *scale_pos_weight* was 4.0 to compensate the imbalance of the benign and phishing classes. Other hyperparameters remained on default values.

IV. Preliminary Results

Following the proposed methodology, we achieved 0.9716 precision rate, 0.9370 recall, and 0.9540 F-1 score. The latter we consider the most important due to the imbalance between the benign and phishing classes. The false positive rate was 0.23%. Table I shows the confusion matrix, illustrating the classification results in detail.

Figure 2 shows the SHAP [44] summary plot for the TOP 10 most important features. The plot indicates the distribution of domains assigned to benign (to the left of the vertical line) and phishing (to the right) classes. The chart's height corresponds to the number of samples, while the color indicates the feature value. The classification of samples situated around the middle is highly dependent on other attributes. The detailed description of all 80 features is provided on Zenodo [40].

The top feature in the list is domain age, which confirms our assumptions that long-running services are statistically more credible. Big Internet players distribute their services across numerous nodes, often spanning multiple countries, which likely explains the importance of the IPv4 count. The results confirm that lexical properties like the number of subdomains, TLD, or malicious domain n-grams, usable for detecting C&C DGA domains, are also instrumental in identifying phishing. The importance of DNS TTL further confirms the findings of Bilge et al. [19]. TLS and geolocation features were also contributive but did not reach the top 10 list. From the TLS-based features, the most important were the length of the certificate chain and the number of TLS extensions. The most crucial geolocation feature was the combination of countries where the servers of domain-related IPs were located.

Overall, the results validate the usability of the examined data sources for phishing domain detection. While there is still space for further improvements, our current feature vector performed well in the given classification task and successfully identified most of the phishing domains, keeping the false positive rate on a low level.

TABLE I
Confusion matrix for the Phishing domain classifier.

		Predicted	
		Phishing	**Benign**
Actual	**Phishing**	10 399	699
	Benign	304	129 468

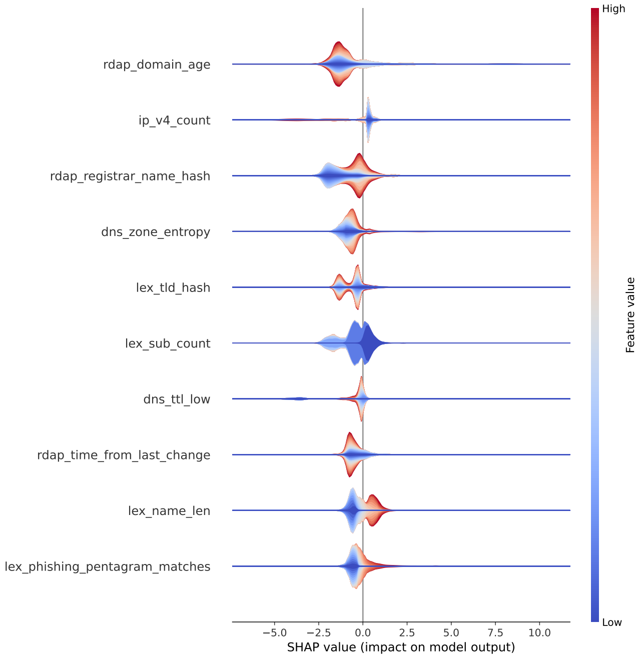

Fig. 2. SHAP summary plot with TOP 10 most important features

V. Conclusion

In the paper, we presented our approach to phishing domain detection, including early experiments on a large dataset of domain-related information. Preliminary results indicate that our methodology, which leverages intelligence from multiple data sources, is effective in identifying phishing domains. The achieved scores of standard metrics are promising and comparable to related studies. While we believe we are on the right track towards effective phishing detection, additional effort is required to deeper analyze the effects individual features and draw more detailed conclusions.

In the upcoming research, we aim to compare the distinct data sources according to their information value and contributions to the decision-making process. Additionally, we plan to expand our feature vector with more RDAP-based features that we have not analyzed yet, and explore the usability of our approach with other clasification methods, such as AdaBoost, LightGBM, SVM, and Deep Neural Networks.

Acknowledgment

The research presented in this paper has been supported by the "Flow-based Encrypted Traffic Analysis" project, no. VJ02010024, granted by the Ministry of the Interior of the Czech Republic and the "Smart Information Technology for a Resilient Society" project, no. FIT-S-23-8209, granted by Brno University of Technology.

References

[1] S. Sheng, M. Holbrook, P. Kumaraguru, L. F. Cranor, and J. Downs, "Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2010, pp. 373–382.

[2] S. Gorling, "The Myth of User Education," in *Proceedings of the 16th Virus Bulletin International Conference*, Montreal, CN, October 2006.

[3] B. A. Gyunka and A. O. Christiana, "Analysis of human factors in cyber security: A case study of anonymous attack on HBGary." *Computing & Information Systems*, vol. 21, no. 2, 2017.

[4] S. Bell and P. Komisarczuk, "An analysis of phishing blacklists: Google Safe Browsing, OpenPhish, and Phishtank," in *Proceedings of the Australasian Computer Science Week Multiconference*, 2020, pp. 1–11.

[5] S. Sheng, B. Wardman, G. Warner, L. Cranor, J. Hong, and C. Zhang, "An empirical analysis of phishing blacklists," 2009.

[6] C. Lewis and M. Sergeant, "Overview of Best Email DNS-Based List (DNSBL) Operational Practices," RFC 6471 (Informational), Internet Engineering Task Force (IETF), Request for Comments (RFC) 6471, Jan. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6471.txt

[7] P. Prakash, M. Kumar, R. R. Kompella, and M. Gupta, "Phishnet: predictive blacklisting to detect phishing attacks," in *2010 Proceedings of IEEE INFOCOM*. IEEE, 2010, pp. 1–5.

[8] D. L. Cook, V. K. Gurbani, and M. Daniluk, "Phishwish: a stateless phishing filter using minimal rules," in *Financial Cryptography and Data Security: 12th International Conference, Cozumel, Mexico, January 28-31. Revised Selected Papers 12*. Springer, 2008, pp. 182–186.

[9] E. Kirda and C. Kruegel, "Protecting users against phishing attacks with AntiPhish," in *29th Annual International Computer Software and Applications Conference*, vol. 1, 2005, pp. 517–524 Vol. 2.

[10] M. Khonji, Y. Iraqi, and A. Jones, "Phishing detection: a literature survey," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 2091–2121, 2013.

[11] M. Chandrasekaran, K. Narayanan, and S. Upadhyaya, "Phishing email detection based on structural properties," in *NYS cyber security conference*, vol. 3. Albany, New York, 2006, pp. 2–8.

[12] I. Fette, N. Sadeh, and A. Tomasic, "Learning to detect phishing emails," in *Proceedings of the 16th international conference on World Wide Web*, 2007, pp. 649–656.

[13] S. Abu-Nimeh, D. Nappa, X. Wang, and S. Nair, "A comparison of machine learning techniques for phishing detection," in *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*, 2007, pp. 60–69.

[14] R. Hranický, F. Breitinger, O. Ryšavý, J. Sheppard, F. Schaedler, H. Morgenstern, and S. Malik, "What do incident response practitioners need to know? A skillmap for the years ahead," *Forensic Science International: Digital Investigation*, vol. 37, p. 301184, 2021.

[15] M. Heiderich, T. Frosch, and T. Holz, "Iceshield: Detection and mitigation of malicious websites with a frozen dom," in *Recent Advances in Intrusion Detection: 14th International Symposium, RAID 2011, Menlo Park, CA, USA, September 20-21, 2011. Proceedings 14*. Springer, 2011, pp. 281–300.

[16] K.-T. Chen, J.-Y. Chen, C.-R. Huang, and C.-S. Chen, "Fighting phishing with discriminative keypoint features," *IEEE Internet Computing*, vol. 13, no. 3, pp. 56–63, 2009.

[17] M. Hara, A. Yamada, and Y. Miyake, "Visual similarity-based phishing detection without victim site information," in *2009 IEEE Symposium on Computational Intelligence in Cyber Security*. IEEE, 2009, pp. 30–36.

[18] A. Drichel, N. Faerber, and U. Meyer, "First step towards explainable dga multiclass classification," in *Proceedings of the 16th International Conference on Availability, Reliability and Security*, 2021, pp. 1–13.

[19] L. Bilge, E. Kirda, C. Kruegel, and M. Balduzzi, "EXPOSURE: Finding malicious domains using passive DNS analysis." in *NDSS*, 2011, pp. 1–17.

[20] M. Antonakakis, R. Perdisci, W. Lee, N. Vasiloglou II, and D. Dagon, "Detecting malware domains at the upper {DNS} hierarchy," in *20th USENIX Security Symposium (USENIX Security 11)*, 2011.

[21] K. Hageman, E. Kidmose, R. R. Hansen, and J. M. Pedersen, "Can a TLS certificate be phishy?" in *18th International Conference on Security and Cryptography*. SCITEPRESS Digital Library, 2021, pp. 38–49.

[22] I. Torroledo, L. D. Camacho, and A. C. Bahnsen, "Hunting malicious TLS certificates with deep neural networks," in *Proceedings of the 11th ACM workshop on Artificial Intelligence and Security*, 2018, pp. 64–73.

[23] M. Kuyama, Y. Kakizaki, and R. Sasaki, "Method for detecting a malicious domain by using whois and dns features," in *3rd international conference on digital security and forensics*, vol. 74, 2016.

[24] Y. Shi, G. Chen, and J. Li, "Malicious domain name detection based on extreme machine learning," *Neural Processing Letters*, vol. 48, pp. 1347–1357, 2018.

[25] F. Sadique, R. Kaul, S. Badsha, and S. Sengupta, "An automated framework for real-time phishing URL detection," in *2020 10th Annual Computing and Communication Workshop and Conference (CCWC)*. IEEE, 2020, pp. 0335–0341.

[26] D. Sahoo, C. Liu, and S. C. Hoi, "Malicious URL detection using machine learning: A survey," *arXiv preprint arXiv:1701.07179*, 2017.

[27] C. Do Xuan, H. D. Nguyen, and V. N. Tisenko, "Malicious URL detection based on machine learning," *International Journal of Advanced Computer Science and Applications*, vol. 11, no. 1, 2020.

[28] R. Patgiri, H. Katari, R. Kumar, and D. Sharma, "Empirical study on malicious URL detection using machine learning," in *Distributed Computing and Internet Technology: 15th International Conference, ICDCIT 2019, Bhubaneswar, India, January 10–13, 2019, Proceedings 15*. Springer, 2019, pp. 380–388.

[29] A. Niakanlahiji, B.-T. Chu, and E. Al-Shaer, "Phishmon: A machine learning framework for detecting phishing webpages," in *2018 IEEE International Conference on Intelligence and Security Informatics (ISI)*. IEEE, 2018, pp. 220–225.

[30] A. Singh and N. Goyal, "A comparison of machine learning attributes for detecting malicious websites," in *11th International Conference on Communication Systems & Networks*. IEEE, 2019, pp. 352–358.

[31] J. McGahagan, D. Bhansali, C. Pinto-Coelho, and M. Cukier, "A comprehensive evaluation of webpage content features for detecting malicious websites," in *2019 9th Latin-American Symposium on Dependable Computing (LADC)*. IEEE, 2019, pp. 1–10.

[32] R. Perdisci, I. Corona, and G. Giacinto, "Early detection of malicious flux networks via large-scale passive DNS traffic analysis," *IEEE Transactions on Dependable and Secure Computing*, vol. 9, no. 5, pp. 714–726, 2012.

[33] N. Hason, A. Dvir, and C. Hajaj, "Robust malicious domain detection," in *Cyber Security Cryptography and Machine Learning: Fourth International Symposium, CSCML 2020, Be'er Sheva, Israel, July 2–3, 2020, Proceedings 4*. Springer, 2020, pp. 45–61.

[34] M. Chatterjee and A.-S. Namin, "Detecting phishing websites through deep reinforcement learning," in *IEEE 43rd Annual Computer Software and Applications Conference*, vol. 2. IEEE, 2019, pp. 227–232.

[35] Cisco Systems, Inc. Cisco umbrella. Accessed: January 25, 2023. [Online]. Available: https://umbrella.cisco.com/

[36] B. Rahbarinia, R. Perdisci, and M. Antonakakis, "Segugio: Efficient behavior-based tracking of malware-control domains in large ISP networks," in *2015 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. IEEE, 2015, pp. 403–414.

[37] Openphish. Accessed: February 12, 2023. [Online]. Available: https://openphish.com/

[38] OpenDNS. Phishtank. Accessed: February 24, 2023. [Online]. Available: https://phishtank.org/

[39] C. Wagner, A. Dulaunoy, G. Wagener, and A. Iklody, "MISP: The design and implementation of a collaborative threat intelligence sharing platform," in *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*, 2016, pp. 49–56.

[40] R. Hranický, A. Horák, J. Polišenský, P. Pouč, and O. Ondryáš, "Phishing and Benign Domain Dataset (DNS, IP, WHOIS/RDAP, TLS, GeoIP), Version 1.0," Sep. 2023, Brno University of Technology. [Online]. Available: https://zenodo.org/doi/10.5281/zenodo.8364667

[41] C. Hajaj, N. Hason, and A. Dvir, "Less is more: Robust and novel features for malicious domain detection," *Electronics*, vol. 11, no. 6, p. 969, 2022.

[42] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 785–794.

[43] K. B. Abou Omar, "XGBoost and LGBM for Porto Seguro's Kaggle challenge: A comparison," *Preprint Semester Project*, 2018.

[44] L. Merrick and A. Taly, "The explanation game: Explaining machine learning models using shapley values," in *Machine Learning and Knowledge Extraction: 4th IFIP International CD-MAKE Conference, Dublin, Ireland, August 25–28, 2020, Proceedings 4*. Springer, 2020, pp. 17–38.