

Received July 14, 2020, accepted July 27, 2020, date of publication July 31, 2020, date of current version August 13, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3013494

Inline Detection of DGA Domains Using Side Information

RAAGHAVI SIVAGURU¹, (Member, IEEE), JONATHAN PECK^{1b,2,3}, FEMI OLUMOFIN⁴, ANDERSON NASCIMENTO¹, AND MARTINE DE COCK^{1b,2}

¹School of Engineering and Technology, University of Washington Tacoma, Tacoma, WA 98402, USA

²Department of Applied Mathematics, Computer Science and Statistics, Ghent University, 9000 Ghent, Belgium

³Data Mining and Modeling for Biomedicine, VIB Inflammation Research Center, 9052 Ghent, Belgium

⁴Infoblox, Santa Clara, CA 95054, USA

Corresponding author: Jonathan Peck (jonathan.peck@ugent.be)

This work was supported by the NVIDIA Corporation with the donation of the Titan Xp GPU. The work of Jonathan Peck was supported by the Fellowship of the Research Foundation Flanders (FWO).

ABSTRACT Malware applications typically use a command and control (C&C) server to manage bots to perform malicious activities. Domain Generation Algorithms (DGAs) are popular methods for generating pseudo-random domain names that can be used to establish a communication between an infected bot and the C&C server. In recent years, machine learning based systems have been widely used to detect DGAs. There are several well known state-of-the-art classifiers in the literature that can detect DGA domain names in real-time applications with high predictive performance. However, these DGA classifiers are highly vulnerable to adversarial attacks in which adversaries purposely craft domain names to evade DGA detection classifiers. In our work, we focus on hardening DGA classifiers against adversarial attacks. To this end, we train and evaluate state-of-the-art deep learning and random forest (RF) classifiers for DGA detection using side information that is harder for adversaries to manipulate than the domain name itself. Additionally, the side information features are selected such that they are easily obtainable in practice to perform inline DGA detection. The performance and robustness of these models is assessed by exposing them to one day of real-traffic data as well as domains generated by adversarial attack algorithms. We found that the DGA classifiers that rely on both the domain name and side information have high performance and are more robust against adversaries.

INDEX TERMS Adversarial machine learning, DGA detection, side information.

I. INTRODUCTION

Domain Generation Algorithms (DGAs) are subroutines that generate pseudo-random combinations of characters or words, and output domain name strings [1]. DGAs often use a seed input such as a number, which is embedded as part of the code, or a time-based element such as the system date, time etc., or a combination of both, to generate random strings. These strings are then concatenated with an available top level domain (TLD) to form domain names. The key idea behind DGAs is to generate the same set of domain names when executed by two different machines, such as by a botmaster and on an infected machine, at a given time. The botmaster registers one of the generated domain

names, while the infected machines systematically query the domains from the generated list until one of them is resolved. The domains from the list that have not been registered by the botmaster will typically result in an NXDomain (non-existent domain) response when queried, and can be discarded by the infected machine. This technique is often used by a command and control (C&C) center and an infected bot to establish communication and perform malicious activities as instructed by the C&C server.

Once communication between the infected machines and the botmaster has been established, the C&C server can issue commands to the bots to perform malicious activities such as distributed denial of service (DDoS) attacks, spamming, stealing sensitive information from the compromised machines, etc. In the past, malware authors used a predefined list of domain names, which was embedded in the malware,

The associate editor coordinating the review of this manuscript and approving it for publication was Mohammad Ayoub Khan ^{1b}.

TABLE 1. Overview of existing work on DGA detection.

Input	Retrospective	Inline
Domain name string	[2], [26]	[4], [5], [8], [15], [27] [6], [13], [14], [28], [29]
Side information features	[16], [30]–[33]	our work
Domain name string + side information features	[3], [18], [34]–[36] [10]–[12], [17]	our work

to communicate with the bots. This technique made it easy for the defenders to blacklist the malicious domain names and block further communication, effectively rendering the malware useless. To overcome this, modern C&Cs use DGAs to randomly generate domain names that are registered on the go, making them harder to detect. It is therefore important to identify the domains generated by DGAs and block them before they can be used to establish communication between the bot and the C&C center. There are several machine learning approaches proposed in the literature to address this issue including [2]–[9] and other work that we cite later in this paper. These well known state-of-the-art classifiers can be deployed in real-world DNS applications to detect DGA domain names and block them. While some work focuses on detecting DGAs from NXDomains [4], our work aims to detect DGAs from traffic to domains that have already been resolved.

Commonly used approaches for DGA detection can be categorized according to how fast they are able to flag malicious activity in DNS traffic. As illustrated in Table 1, some techniques work in a *retrospective* manner, in which past DNS traffic, which is logged over a certain window, is analyzed in batches to detect anomalies. Other techniques work *inline*, meaning that they can detect DGA domains as soon as they are queried. There are two ways in which inline DGA detection can happen:

- The domain first reaches the DGA classifier and if the classifier flags the domain as benign, then the query is passed to the DNS resolver to fetch the resolved IP address of the domain. However, if the classifier flags the domain as DGA, then the query will not be forwarded to the DNS resolver and it simply blocks the communication with that domain.
- The domain first queries through the DNS resolver; the DGA classifier uses the features learned from the DNS response to decide on whether the domain is DGA or not.

Our work fits into the second category of inline detection, where both the domain name and the side information features learned from the DNS query/response are used by the classifier for DGA detection. The side information features are carefully selected to allow *inline DGA detection in the broader sense*. In the strict sense, inline DGA detection means that the information required to determine whether a domain name is DGA or not is available from the DNS query data alone. A DNS resolver can use the strictly inline

DGA classifier’s decision to determine if it is safe or not to resolve the query. A less conventional version of inline detection, which we refer to as “inline DGA detection in the broader sense”, is one where data attributes from DNS responses are required (in addition to DNS queries). This means that the DNS resolver must resolve the query first, feed the information obtained to the DGA classifier, and then use the DGA classifier’s decision to determine if it is safe to get the DNS response to the client or not. As we observe in our experimental results, taking information from DNS responses into account improves the ability of DGA classifiers to correctly detect DGA domains among resolvable traffic. We note that any dependence on data requiring queries to additional sources, such as the WHOIS database (as used for instance in [10]–[12]), would disqualify the approach from inline detection, even in the broader sense.

Machine learning based approaches to detect DGA domain names in practice can also be categorized according to the information they leverage. One way is to train classifiers to detect DGA domain names using only the domain name string itself, see e.g. [4]–[6], [8], [13]–[15]. The alternative is to train the classifiers using context information such as the IP address of the domain, its geographic location, attributes from DNS response records etc. in addition to the domain name [3], [10], [12], [16]–[18]. In our work, we combine both approaches. The advantage of the former approach is that it does not require gathering of additional information, which may be expensive to collect in real time, and that it allows the defenders to detect the DGA domain names and block them even before they can be resolved. The advantage of the latter approach is that side information is a lot harder for the attacker to manipulate than the domain name string itself, making machine learning models trained on side information potentially more robust against adversarial attacks.

Adversarial machine learning is a research area focused on problems introduced by the use of machine learning techniques in adversarial environments in which an intelligent adversary attempts to exploit the weaknesses in such techniques [19]. The *adversarial attacks* of interest in this paper are *evasion attacks* in which an adversary uses artificially crafted instances, called *adversarial samples*, that are intentionally used to mislead a machine learning system and produce erroneous results. The goal of evasion attacks in the context of DGA detection is to generate domains that will be labeled as benign by the DGA classifier. The vulnerability

of a classifier against evasion attacks is measured in terms of DGA detection rate, which is the proportion of the adversarial samples predicted as malicious by the classifier. Lower DGA detection rates indicate high vulnerability of the classifier to the attack. There exists several evasion attacks against DGA classifiers such as CharBot [20], DeepDGA [21], DeceptionDGA [22], MaskDGA [23] and the DGAs (HMM & PCFG-based) proposed by [24]. CharBot and MaskDGA are black-box targeted evasion attacks that do not require any knowledge about the DGA classifier and are intended to generate samples that can evade detection by *any* classifier. On the other hand, DeceptionDGA is a white-box attack algorithm that uses the knowledge of features used by the DGA classifier to generate evading instances specific to a given classifier. Both types of attacks are found to be extremely powerful in generating domains that can evade detection by the DGA classifiers with high probability.

The main contributions of our work are:

- A comprehensive survey of lexical and side information features proposed in the literature on DGA detection.
- An experimental evaluation of the feasibility in collecting the features and their effectiveness when deployed for inline detection of DGAs in real streams of passive DNS traffic, which leads to a shortlist of features that are actually beneficial in practice.
- Experimental results that show how the side information features can make DGA classifiers more robust against adversarial attacks.

The paper is organized as follows. Section II gives an overview of related work in the fields of adversarial machine learning and DGA detection. Section III provides a detailed overview of side information features that can be extracted from DNS traffic to aid in the detection of DGA domains. In section IV, we list the 26 human engineered lexical features that are extracted manually from the domain name string in order to train the RF classifier for DGA detection. section V gives an overview of the different classifiers we will be studying and attempting to harden against adversarial attacks. Section VI describes the experimental setup and reports all of our empirical results. Finally, section VII concludes the work. All findings are based on the first author's master's thesis, the full version of which is available at the University of Washington [25].

II. RELATED WORK

Given the importance of being able to detect and block DGA domain related traffic, it comes as no surprise that the problem of automatic DGA detection has received a considerable amount of attention over the last decade. There are various ways in which existing DGA detection approaches differ from each other. As illustrated in Table 1, DGA detection can be categorized according to the kind of input that is required. Some techniques require just the *domain name string*, while other techniques require *side information*, or a combination of both. Both kinds of input have their own advantages and disadvantages. Methods that rely only on the domain

name string are popular because side information is typically harder to obtain. On the other hand, features extracted from side information are harder to manipulate, making methods based on them more robust against adversarial attacks. All the approaches presented in our paper perform *inline* DGA detection using domain name only, side information only and a combination of both domain name & side information features.

Furthermore, the classifiers can be trained in two ways to detect if a given domain name is generated by a DGA or not. The first technique is the *featureful approach*, where the classifier relies on human engineered features extracted from the domain names. The second technique is the *featureless approach*, where the classifier learns the features automatically during the training process. Classifiers that are based on deep learning architectures like Long Short-Term Memory (LSTM) [5], [8] and Convolutional Neural Network (CNN) models [6], [9] leverage the featureless approach, whereas models such as random forests (RFs) adopt the featureful approach. In our work, we will be using both featureful and featureless approaches to train random forest and deep learning classifiers for DGA detection.

We note that, as is clear from Table 1, side information features have so far only been used in retrospective approaches to DGA detection, while our focus is on inline detection instead. Still the existing work is very relevant to ours, as we can use many side information features that have been previously proposed in a retrospective setting for inline detection as well. A similar comment applies to techniques for domain name string classification. In Table 2 and Table 4 we provide a detailed overview of features that have been proposed, with references, and whether we retain them in our classifiers or not. A more detailed explanation about these features is provided in section III and section IV.

III. SIDE INFORMATION FEATURES

In this section we provide a detailed overview of side information features that can be extracted from DNS traffic to aid in the detection of DGA domains. An overview of all features is presented in Table 2, accompanied by a list of citations that illustrates the popularity of each kind of feature in the literature. The order of the side information features listed in Table 2 indicates the importance of those features in DGA detection as ranked by the Random Forest model (see section V). Not all features are equally easy to obtain in practice, and their contribution to the predictive accuracy of DGA classifiers varies. The last column of Table 2 indicates whether we retained the feature in our DGA-classifiers. Figure 1 shows a sample resource record from which the side information features are extracted. In figure 1, the attribute "name" represents the fully qualified domain name (FQDN), "ttl" represents the time-to-live of the DNS query, "type" represents the resource record type, "class" represents the class of resource record and "data" represents the resolved IP address. Below we give a more in-depth description of each kind of feature, and its typical use in the the literature

TABLE 2. Side information features.

Feature	Description	Reference	Retained
rrlen	Resource record length	[32]	✓
country	Country name that the domain maps to	[11], [12], [18], [34]	✓
ttl	Time-to-live of the DNS query	[32], [34]	✓
n_ip	Number of distinct IP addresses the domain maps to	[11], [12], [17], [18]	✓
qtype	Type of DNS packet requested	[32]	✓
rtype	Record type of the DNS response	[32]	✓
n_asn	Number of distinct ASNs the domain maps to	[33]	✓
subnet	Do all IPs belong to same subnet	[11], [12]	✓
n_countries	Number of distinct countries the domain maps to	[11], [12], [17], [18], [35]	✓
timestamp	Features derived from timestamp of the DNS query	[17], [18], [30]	✗
opcode	Kind of DNS query	[32]	✗
AA	Authoritative answer	[32]	✗
QDCOUNT	Number of entries in question section	[32]	✗
ANCOUNT	Number of resource records in answer section	[32]	✗
NSCOUNT	Number of name servers in authoritative section	[32]	✗
ARCOUNT	Number of resource records in additional record section	[32]	✗
RCODE	Response code	[32]	✗
rDNS	Reverse DNS query results	[11], [12], [17]	✗
TTL statistics	Mean, standard deviation etc. of time-to-live	[11], [12], [17], [18]	✗
n_domains	Number of distinct domains associated with the IP	[11], [12], [17], [18]	✗
n_queries	Number of queries for the domain and (domain, IP) pair	[18]	✗
WHOIS features	Registrar, domain creation/expiration date etc.	[10]–[12], [34]	✗

TABLE 3. TTL distribution (in seconds) for benign and DGA domains.

Type	Mean TTL	SD TTL	Median TTL
Benign	109,447	1,421,829	3,600
DGA	29,255	4,701,205	900

```

[{'name': 'junebugweddings.com.', 'ttl': 300, 'type': 1, 'class': 1, 'data': '104.27.191.148'}, {'name': 'junebugweddings.com.', 'ttl': 300, 'type': 1, 'class': 1, 'data': '104.27.190.148'}]
    
```

FIGURE 1. An example DNS resource record.

on DGA detection. Figure 2 shows a comparison of density plots for some of the side information features extracted from benign and DGA domain names, illustrating their predictive power. The different side information features are as follows:

- **rrlen**: This feature measures the length of the RData field, which is extracted directly from the DNS response resource record. The RData in a DNS response encompasses a list of resolved IP addresses, the time-to-live value of the query and the type of resource record.

- **country**: This feature refers to the geographic location that the resolved IP address maps to. If the DNS resource record contains multiple IP addresses, the country for each of the IP addresses is first identified. If all of the IP addresses belong to the same country, then this feature takes up that name. On the other hand, if any of the IP addresses map to a different location, then the value of this feature would be “multi-valued”. Alternatively, if the location could not be identified, then this feature takes the value “unknown”. This feature is then converted to categorical values that range between 0 and 185, which means that the domains in our dataset map to 184 different countries plus the values “multi-valued” and “unknown”.
- **ttl**: This feature represents the time-to-live value of the DNS query, which is the time interval that the resource record can be cached by the DNS resolver, and is directly obtained from the DNS response resource record. Table 3 compares the distribution of TTL values (in seconds), in terms of mean, standard deviation and median, for benign and DGA domains in our dataset (see section VI-A). It can be seen that DGA domains are in general far more short-lived than benign domains. For better visibility in figure 2, the density plot for TTL values are shown in hours instead of seconds.
- **n_ip**: This feature indicates the number of distinct IP addresses that are returned for the DNS domain lookup.

TABLE 4. Lexical features used by B-RF.

Feature	Description	Reference	Retained
domain_len	Domain name length	[4], [6], [11]–[14], [20], [27], [29]	✓
sld_len	Second level domain length	[13], [14], [20]	✓
tld_len	Top level domain length	[13], [14], [20]	✓
uni_domain	Domain Unique Characters length	[13], [14], [20]	✓
uni_sld	SLD Unique Characters length	[13], [14], [20]	✓
uni_tld	TLD Unique Characters length	[13], [14], [20]	✓
flag_dga	Has malicious TLD	[13], [14], [20], [27]	✓
tld_hash	TLD Hash	[6], [13], [14], [20]	✓
flag_dig	Starts with Digit	[6], [13], [14], [20]	✓
sym	Symbol ratio	[6], [13], [14], [20]	✓
hex	Hex ratio	[6], [13], [14], [20]	✓
dig	Digit Ratio	[4], [11]–[14], [17], [20], [29]	✓
vow	Vowel Ratio	[4], [6], [13], [14], [20], [29]	✓
con	Consonant Ratio	[13], [14], [20]	✓
rep_char_ratio	Ratio of Repeated Characters	[4], [14], [20]	✓
cons_con_ratio	Ratio of Consecutive Consonants	[4], [14], [20], [29]	✓
cons_dig_ratio	Ratio of Consecutive Digits	[4], [14], [20]	✓
tokens_sld	Number of tokens in SLD	[13], [14], [20], [27]	✓
digits_sld	Number of digits in SLD	[13], [14], [20], [27]	✓
ent	Entropy of characters in SLD	[4], [6], [13], [14], [20], [29]	✓
gni	Gini Index of characters in SLD	[6], [13], [14], [20]	✓
cer	Classification error of characters in SLD	[6], [13], [14], [20]	✓
2gram_med	2-Gram Median of characters in SLD	[6], [13], [14], [20]	✓
3gram_med	3-Gram Median of characters in SLD	[6], [13], [14], [20]	✓
2gram_cmed	2-Gram Circle Median of characters in SLD	[13], [14], [20]	✓
3gram_cmed	3-Gram Circle Median of characters in SLD	[13], [14], [20]	✓

It is manipulated directly by accessing the list of IPs contained in the RData field of the DNS response resource record.

- **qtype**: This feature represents the DNS query type that can be extracted from the question section of the DNS query. Figure 2 shows the different values for this features in our dataset.
- **rtype**: This feature represents the resource record type that can be extracted directly from the RData field in the DNS response resource record. Figure 2 shows the different values for this features in our dataset.
- **n_asn**: This feature indicates the number of distinct autonomous system numbers that the IP addresses map to. The ASN for a given IP address is obtained by using Python Geolite2 Maxmind API.¹
- **subnet**: This feature is a boolean value that represents if all the IP addresses belong to the same subnet. A value of 0 indicates that one or more of the IP addresses, returned in the DNS response, belong to a different subnet and value of 1 indicates that all the IP addresses map to the same subnet.
- **n_countries**: This feature represents the distinct number of countries that the resolved IP addresses map to. This feature has a very similar distribution when compared to the “n_asn” feature, which can be observed in figure 2.
- **timestamp**: The timestamp denotes the time at which the DNS query was issued by a host. This feature in itself may not be useful in detecting DGAs. Some of the past studies record all of the timestamps at which a particular domain name was queried and construct time-series data to analyze the periodicity at which the benign and DGA domains are queried [17], [30], whereas [18] computes the lifespan of a domain by subtracting the first and last seen timestamps of the domain name. Such approaches require access to past DNS traffic and hence are regarded as “retrospective”. Since we only focus on performing inline DGA detection in our work, we do not use the timestamp feature to perform DGA classification.
- **opcode**: This feature represents the kind of query such as standard query, inverse query, request for server status etc. In our dataset, all the domains being queried belong to *standard query* type and hence using this feature does not contribute in the prediction of DGA domain names.
- **aa**: This feature is a boolean flag which represents if the responding name server is an authority for the domain

¹<https://geoip2.readthedocs.io/en/latest/>

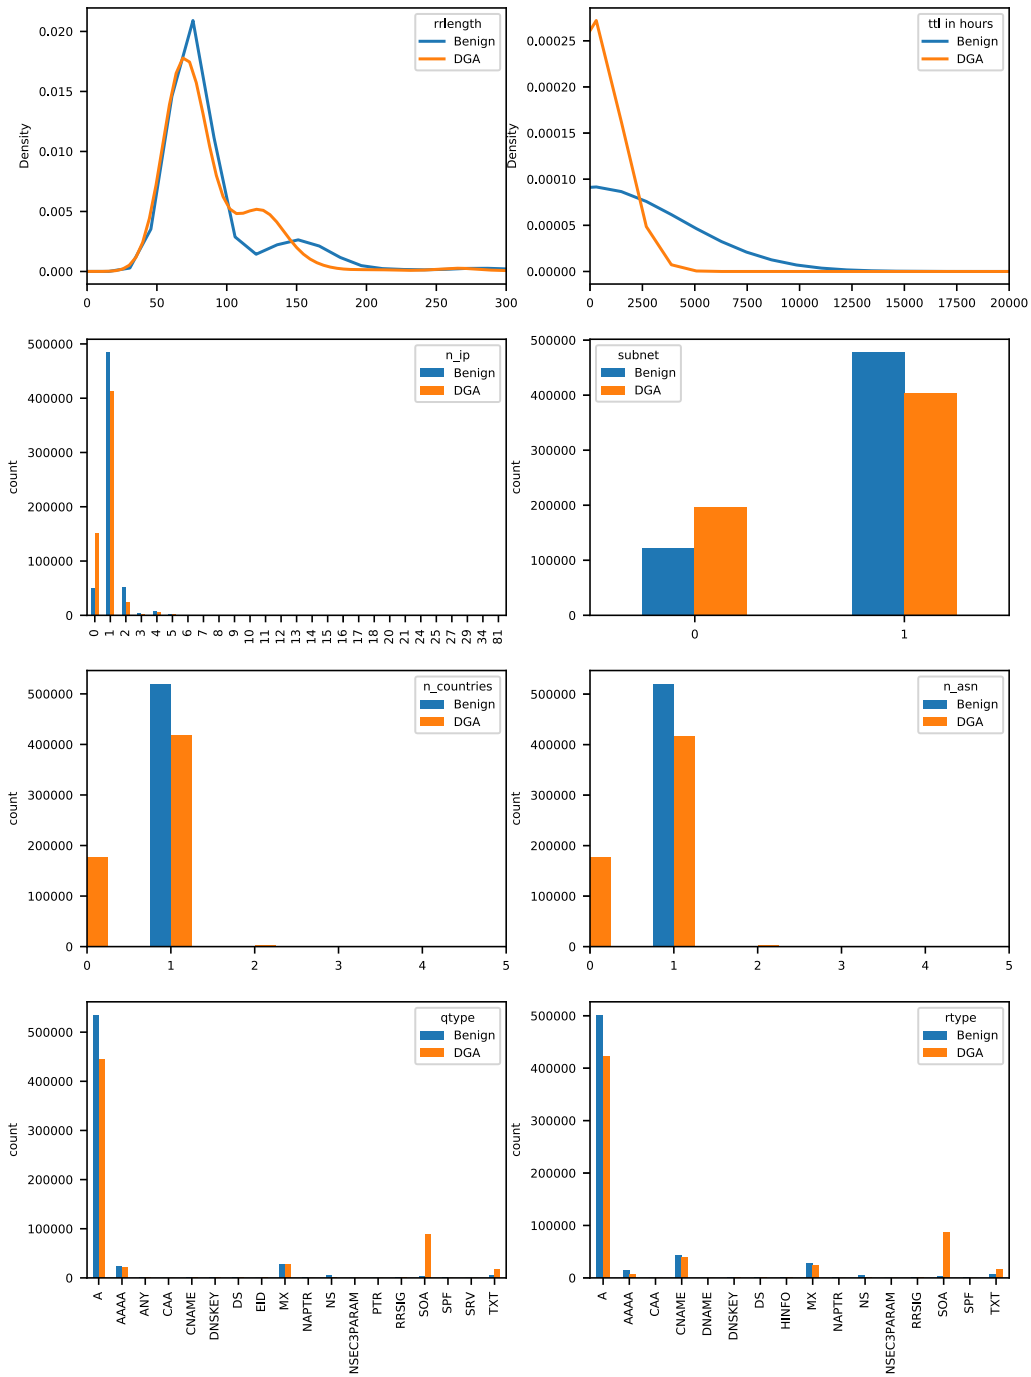


FIGURE 2. Comparison of values for side information features extracted from benign and DGA domains.

name being queried. The AA flag for all DNS responses in our dataset has the same value “True” and hence we do not leverage the AA flag information while training our DGA classifiers.

- **qdcount, ancount, nscout, arcount:** At this time, our DNS traffic collector do not capture this information & hence we do not use these features to train our model. However, it can be easily obtained from the DNS query and resource records.

- **rcode:** Since our dataset comprises of resolved domain names only, the rcode remains “0” for all the samples and hence we discard this information.
- **TTL statistics:** This refers to a collection of features such as standard deviation, mean, minimum, maximum etc. of all time-to-live values extracted from the DNS response. While these features are relevant in a retrospective approach that investigates a domain based on *all* DNS resource records related to it say during the

past 24 hours, it is not meaningful for fast inline DGA detection. Indeed, since all of the TTL values in a single response record have constant values, it would not add value to include these statistics as features.

- **n_domains:** This feature represents the number of distinct domain names that are mapped to a given IP address. In order to use this feature, one needs to maintain a bipartite graph that depicts the mapping for each (domain, IP) pair. Again, this method of performing graph inference is computationally intensive and does not contribute towards inline detection of DGA domains. Therefore we refrain from using this side information feature while training our DGA classifiers.
- **n_queries:** Similar to “timestamp” and “n_domains”, this feature also requires storing and fetching of information from past DNS traffic and hence n_queries cannot be used for inline detection of DGAs.
- **WHOIS features:** Extracting WHOIS features such as registrar, domain creation/expiration date etc. involves very expensive WHOIS queries. This affects the capability of the classifier to perform inline DGA detection on-the-go and hence we do not use any feature that require WHOIS queries.

IV. LEXICAL FEATURES

In this section we list the 26 human engineered lexical features that are extracted manually from the domain name string in order to train the RF classifier for DGA detection. Table 4 shows a list of the lexical features used in B-RF and details on how the feature values are calculated are given below:

- **domain_len:** This feature represents the length of the domain name, which is the number of characters in the SLD.TLD pair. For example, we refer “google.com” as the domain name, where “google” indicates the SLD (second level domain) and “com” indicates the TLD (top level domain). The value of the feature domain_len for the domain name “google.com” is 10.
- **sld_len:** This feature represents the number of characters in the second level domain.
- **tld_len:** This feature represents the number of characters in the top level domain.
- **uni_domain:** This feature represents the number of unique characters in the domain name, after removing special characters such as ‘.’ & ‘-’ from the domain name.
- **uni_sld:** This feature represents the number of unique characters in the second level domain, after removing special characters such as ‘.’ & ‘-’ from the SLD.
- **uni_tld:** This feature represents the number of unique characters in the top level domain, after removing special characters such as ‘.’ & ‘-’ from the TLD.
- **flag_dga:** This feature represents a boolean value (0 or 1) that indicates if the domain name contains any of the following TLDs, which are known to be frequently

associated with malicious activities²: “study”, “party”, “click”, “top”, “gdn”, “gq”, “asia”, “cricket”, “biz”, “cf”.

- **tld_hash:** This feature represents the hash value of top level domain.
- **flag_dig:** This feature represents a boolean value that indicates if the domain name starts with a digit/number (0-9).
- **sym:** This feature represents the ratio of number of special characters in the SLD to the total number of characters in SLD (sld_len).
- **hex:** This feature represents the ratio of number of hexadecimal characters (0-9 & a-f) in the SLD to the total number of characters in the SLD.
- **dig:** This feature represents the ratio of number of digits (0-9) in the SLD to the total number of characters in the SLD.
- **vow:** This feature represents the ratio of number of vowels (‘a’, ‘e’, ‘i’, ‘o’, ‘u’) in the SLD to the total number of characters in the SLD.
- **con:** This feature represents the ratio of number of consonants in the SLD to the total number of characters in the SLD.
- **rep_char_ratio:** This feature represents the ratio of number of characters that occurs more than once in the SLD to the total number of unique characters in the SLD.
- **cons_con_ratio:** This feature represents the ratio of consecutive consonants (such as “ct”, “fk”, “ns” etc.) to the length of the domain (domain_len).
- **cons_dig_ratio:** This feature represents the ratio of consecutive digits (such as “92”, “24”, “75” etc.) to the length of the domain (domain_len).
- **tokens_sld:** This feature represents the number of tokens in the SLD, where a token indicates sequence of characters separated by ‘-’.
- **digits_sld:** This feature represents the total number of digits in the SLD.
- **ent:** This feature represents the normalized entropy value of the characters in SLD and is calculated using the formula:

$$\text{ent} = \frac{\sum_{i=1}^n p_i \cdot \log_2(p_i)}{\log_2(\text{sld_len})}$$

where n represents the number of unique characters in the SLD and p_i represents the proportion between the frequency of the unique character c_i in the SLD to the total number of unique characters in the SLD.

- **gni:** This feature represents the Gini value of the characters in SLD and is calculated using the formula:

$$\text{gni} = 1 - \sum_{i=1}^n p_i^2$$

where n represents the number of unique characters in the SLD and p_i represents the proportion between the

²<https://www.spamhaus.org/statistics/tlds/>




FIGURE 3. DGA detection using LSTM.MI + RF model.

TABLE 5. Performance evaluation of DGA classifiers using 5-fold cross-validation.

Model	Features	Performance metrics	
		AUC@ 0.1%FPR	TPR@ 0.1%FPR
B-RF	DNS	53.23%	16.21%
	Lexical	89.78%	97.44%
	DNS + Lexical	98.19%	99.42%
LSTM.MI	Domain name string	94.47%	98.80%
LSTM.MI + B-RF	Domain name string + DNS	96.51%	99.89%
	Domain name string + DNS + Lexical	99.17%	99.91%

frequency of the unique character c_i in the SLD to the total number of unique characters in the SLD.

- **cer**: This feature represents the classification of error of characters in SLD, which is computed using the formula:

$$\text{cer} = 1 - \max_{i=1, \dots, n} p_i$$

where p_i represents the proportion between the frequency of the unique character c_i in the SLD to the total number of unique characters in the SLD.

- **2gram_med**: This feature represents the median of 2-gram frequencies in SLD.
- **3gram_med**: This feature represents the median of 3-gram frequencies in SLD.
- **2gram_cmed**: In order to compute this feature, the SLD of the domain is concatenated again with the SLD. (i.e) For example, if “google” is the SLD, a string such as “googlegoogle” is formed. The 2gram_med is then calculated on this newly formed string “googlegoogle” to obtain the value of this feature.
- **3gram_cmed**: In order to compute this feature, the SLD of the domain is concatenated again with the SLD. (i.e) For example, if “yahoo” is the SLD, a string such as “yahooyahoo” is formed. The 3gram_med is then calculated on this newly formed string “yahooyahoo” to obtain the value of this feature.

V. DGA CLASSIFIERS

We consider three different DGA classifiers in this work, which we detail below. We chose one model representative of the featureful approach (B-RF), one deep learning model which represents the featureless approach (LSTM.MI) and

finally a hybrid model which combines both approaches (LSTM.MI+B-RF).

A. B-RF

B-RF is a DGA classifier based on random forests. It consists of 100 trees and each tree is trained using a subset of the feature space to avoid overfitting. Entropy is used as the criterion to decide the split attribute while growing the trees in the random forest. There are 3 variants of B-RF classifier, each trained either on lexical features (as the RF classifier in [14]) or DNS features, or a combination of both lexical and DNS features. The performance of these variants of the B-RF classifier is listed in the first 3 rows of Table 5.

B. LSTM.MI

Woodbridge *et al.* [5] were the first to propose deep learning for DGA domain name detection. Their DGA classifier is a neural network consisting of an embedding layer, an LSTM layer, and a single node output layer with sigmoid activation. In this paper, we use the LSTM.MI model that was proposed recently by Tran *et al.* [8]. Its architecture is very similar to that of Woodbridge *et al.* [5]; the main distinction is that the LSTM.MI model is trained with a cost-sensitive learning algorithm that takes class imbalances into account. This allows the LSTM.MI approach to achieve slightly better results than the original LSTM approach (see [8], [14]). The 4th row in Table 5 shows the performance of the LSTM.MI classifier. It operates directly on the domain name string, instead of on lexical features extracted from it. Characters in the domain name are converted to lower case and are encoded with categorical values, ranging from 1 to 38, to represent ‘.’,

TABLE 6. Some examples of benign vs DGA domain names.

Benign domains (labeled 0)	DGA domains (labeled 1)
7ft4.com	vocom.eu
sgtobel.ch	leadhelp.net
intimvoronezh.net	1b6a95e6b5d4.com
essc-tabriz.com	corpncyeajsgeatkopoqs.info
konsaltbezopasnost.ru	kndydusmrlrofrcmfuayfmswrktyl.biz

‘-’, digits from 0 to 9 & characters from ‘a’ to ‘z’. All the domains in our data are fixed to a length of 77 characters, which is the length of the longest domain name in our dataset. Domains that are shorter than 77 characters are padded with zeroes in the left.

C. LSTM.MI+B-RF

The hybrid LSTM.MI+B-RF classifier combines both LSTM.MI and B-RF architectures by training a B-RF classifier with features listed in tables 2 and 4, in addition to the confidence score obtained from the LSTM.MI model for that domain name. The confidence score ranges between 0 and 1, signifying the probability of the domain being a DGA as predicted by the LSTM.MI classifier. The above workflow of DGA detection using LSTM.MI+B-RF setup is depicted in figure 3. The last two rows in Table 5 represent the performance of this DGA classifier.

VI. EXPERIMENTAL RESULTS

We trained the B-RF classifiers in sklearn [37], using the default settings for all hyperparameters not specified otherwise in section V-A. For the LSTM.MI classifier, we used the Keras [38] code³ made available by Tran *et al.* [8] without alteration. This code was written specifically for DGA detection, and we found that it worked well on our dataset, without the need to change the hyperparameter values. The experiments were conducted on a workstation with a GeForce GTX Titan Xp, 3840 cores, and 12GB RAM.

A. DATASET

In the first experiment, we train and evaluate the DGA classifiers from section V on a dataset with 600,000 DGAs (positive) and 600,000 benign (negative) samples.

Table 6 shows some examples of DGA & benign domains. The training data points originate from a real-time stream of passive DNS data, consisting of roughly 10-12 billion DNS queries per day collected from subscribers including ISPs (Internet Service Providers), schools, and businesses. From this traffic, the positive samples are collected by retaining resolved domain names that are listed in DGArchive,⁴ a blacklist containing known DGA domains [1]. Dictionary DGAs, which are human-readable DGA domains belonging to malware families such as suppbobox, gozi, matsnu and nymaim2 are discarded from the training set. This is because

these DGAs look more like benign domains and confuse the DGA classifiers [26]. Since this work is primarily aimed at measuring the impact of adversarial instances such as CharBot, we exclude samples from Dictionary DGAs.

The benign samples are collected based on a predefined set of heuristics as listed below:

- Domain name should have valid DNS characters only (digits, letters, dot and hyphen)
- Domain has to be resolved at least once for every day between June 01, 2019 and July 31, 2019.
- Domain name should have a valid public suffix
- Characters in the domain name are not all digits (after removing ‘.’ and ‘-’)
- Domain should have at most four labels (Labels are sequence of characters separated by a dot)
- Length of the domain name is at most 255 characters
- Longest label is between 7 and 64 characters
- Longest label is more than twice the length of the TLD
- Longest label is more than 70% of the combined length of all labels
- Excludes IDN (International Distribution Network) domains (such as domains starting with xn--)
- Domain must not exist in DGArchive

We note that it is common in earlier research on DGA classification to construct a training dataset with negative examples from whitelists such as Alexa [5], [8], [39]. Alexa⁵ ranks websites based on their popularity in terms of number of page views and number of unique visitors. For example, according to Alexa, the three highest ranked domain names in terms of popularity on 2020-07-08 are google.com, youtube.com, and tmall.com. In our previous research we observed that DGA classifiers trained on a dataset with domains pulled from whitelists and blacklists tend to do well when evaluated on a similar test dataset, but don’t fare well at all when deployed on real traffic [14]. We found that a whitelist such as Alexa, which consists of domains that are collected at the browser level, is not sufficiently representative of all non-malicious domain names that appear in real traffic, causing DGA classifiers trained on Alexa to yield many false positives. This is the main reason why in the current study we opt to use benign samples collected from real traffic with a predefined set of heuristics instead.

Both the DGA and benign domains in the dataset are collected from real-time passive DNS traffic that was observed

³<https://github.com/bkcs-hust/lstm-mi>

⁴<https://dgarchive.caad.fkie.fraunhofer.de/>

⁵https://www.alexa.com/topsites/category/Computers/Internet/Domain_Names

TABLE 7. Real traffic analysis of DGA classifiers on 66,440,662 (~66M) domains.

Model	LSTM.MI	B-RF	LSTM.MI+B-RF
Features	Domain name	DNS + Lexical	Domain name + DNS + Lexical
Out of the ~66M domains in real-traffic, number of domains flagged as DGA by the classifier	3,400,017	1,877,784	2,170,056
Out of the domains flagged as DGA by the classifier, number of domains found in DGArchive	1,151,750	1,149,689	1,150,116
Out of the domains flagged as DGA by the classifier, number of domains found in Alexa	1,626,232	1,717,638	1,420,319

in February 2019. The domains in the dataset are then pre-processed by following the two steps mentioned below:

- Retain only the SLD & TLD of the domain name and discard any 3LD (third level domain) or any other label if present. For example, for the domain name “www.google.com”, the 3LD which is “www” is removed and the SLD.TLD which “google.com” is retained.
- All the alphabetical characters in the domain name are converted to its corresponding lower case characters.

B. PERFORMANCE EVALUATION OF DGA CLASSIFIERS

The true positive rate (TPR) and false positive rate (FPR) for the DGA classifiers are calculated as follows:

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad \text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}}$$

where TP, TN, FP & FN represent the number of true positives, true negatives, false positives and false negatives respectively. The predictive performance of the classifiers is evaluated using 5-fold cross-validation for metrics such as TPR and Area Under the Receiver Operating Characteristic (ROC) Curve (AUC) as tabulated in Table 5. In cybersecurity applications, it is important to achieve high TPR for a very low FPR. This is because it is undesirable to block a large number of benign domains in real-world traffic as this interferes with users’ legitimate business. Hence all the reported metrics are thresholded at a very low FPR of 0.1%. We also obtain the ROC curve by plotting the TPR against the FPR of the classifiers and the AUC is subsequently obtained by taking the integral of the ROC curve. The AUC is a measure of how well the trained classifier can distinguish between the classes. Specifically, it can be interpreted as the probability that the classifier will output a higher score for a randomly chosen DGA domain than it would for a randomly chosen benign domain. An ideal classifier has an AUC score of 1, indicating it will always rank DGA domains higher than benign domains. This makes it possible to use the classifier to perfectly separate the classes via an appropriate threshold on its output scores. A classifier that just randomly guesses the outcome achieves an AUC of 0.5 and a classifier with AUC

0 has basically *inverted* all predictions, i.e. samples labeled as 0 are predicted as 1 by the classifier and vice versa. In addition to the AUC score, the AUC at a fixed FPR of 0.1% is also reported. This thresholded AUC represents the integral of the ROC curve for a FPR of 0 to 0.001.

There are several interesting observations to be made based on Table 5. First, looking at the AUC@1%FPR column, one can see that the predictive performance for inline DGA detection based on DNS features alone does not perform well: the B-RF/DNS based model achieves an AUC@1%FPR of only 53.23%. Second, when it comes to DGA detection based on the domain string alone, the deep learning approach (LSTM.MI) clearly outperforms the random forest approach (B-RF/Lexical) at 94.47% vs. 89.78%. This is fully in line with previous findings [8], [28]. Third, the most interesting and novel result from Table 5 is that the DGA classifiers, when trained with *both* lexical and side information features, have the best overall performance in terms of AUC score and TPR, namely 99.17% for the architecture from figure 3.

C. REAL TRAFFIC ANALYSIS

Next, we apply the best performing classifiers in Table 5 on one day of real traffic DNS traffic to evaluate their predictive performance in real-time. We collected a set of resolved domains that were observed on March 26, 2019 to perform this analysis. As part of pre-processing, the fully qualified domain names are validated against the heuristics mentioned in section VI-A, in order to maintain consistency with the training dataset. The domains that satisfy the heuristics are then retained in this experiment after discarding the third level domain (3LD/subdomain) from the domain name, if present. This resulted in a set consisting of 66,440,681 domains (contains duplicate domains with SLD.TLD pairs), out of which 1,159,662 domains were found in DGArchive and 14,653,217 domains were found in Alexa. There is also an overlap of 1,124,467 domains between the Alexa whitelist and DGArchive blacklist.

Table 7 shows a comparison of the number of domains that were flagged as DGA by the LSTM.MI, B-RF and

TABLE 8. Detection rate of CharBot domains as DGA.

Classifier	Features	DGA (CharBot) detection rate
B-RF	DNS	1.70 % \pm 0.24 %
	Lexical	3.80 % \pm 0.0 %
	Lexical + DNS	20.06 % \pm 0.56 %
LSTM.MI	Domain name string	8.00 % \pm 0.0 %
LSTM.MI+B-RF	Domain name string + DNS	14.98 % \pm 0.63 %
	Domain name string + Lexical + DNS	15.76 % \pm 0.53 %

LSTM.MI+B-RF classifier. The B-RF model (in Table 7) has the highest true positive rate among the 3 models being compared. Out of the 1.87M domains flagged as DGA by the classifier, approximately 61% were found in DGArchive. Although the LSTM.MI classifier catches the highest number of DGAs in real-traffic, the true positive rate is 34% which is 27% lower than the B-RF classifier. However, as seen in the last row of Table 7, the B-RF also has the highest number of false positives. This could be due to the fact that there is a large number of overlapping domains between Alexa and DGArchive as mentioned earlier in this section. A good workaround to reduce the number of false positives during the deployment is to check the flagged domains against Alexa before making the final decision.

D. DEFENSE AGAINST ADVERSARIAL ML

The use of side information is important in the context of adversarial ML because side information is a lot harder to manipulate than the domain name string itself [10]. In order to test this, we generated 1,000 DGA domains with CharBot [20], a simple DGA algorithm that was written specifically to evade existing DGA classifiers.

CharBot is a character-based DGA (hence the name) that generates domain names by randomly modifying two characters in well known benign domains from the Alexa top domain names. Like dictionary DGAs, the CharBot DGA has a list of strings embedded as part of the DGA code. In the case of dictionary DGAs this list is a dictionary of words that are combined in various ways to generate a domain name, while in the case of CharBot the list contains benign domain names that are perturbed slightly to generate a new domain name for malicious purposes. In both cases, the generated domain names exhibit properties that are very close to natural language, which makes them extremely difficult to distinguish from benign domain names.

Since, to the best of our knowledge, CharBot has not been deployed yet in the wild, we cannot collect side information for CharBot domains from real traffic. Instead, we pair up the CharBot domains with the DNS features obtained from 1,000 randomly sampled DGA domains in real traffic. To avoid any bias in the selection of DNS features for CharBot domains, we perform the random sampling for 5 trials and create 5 sets of CharBot DNS features. The lexical features extracted for CharBot are appended with the DNS features, which can then be exposed to DGA classifiers for detection of malicious domains. The idea here is to test if the DGA

classifiers trained on side information features are successful in detecting CharBot domains.

Table 8 shows the CharBot detection rate, which is the average proportion of CharBot domains that were flagged as DGA by the classifiers over the 5 randomized trials. Higher values of CharBot detection rate indicate that the classifier is more robust against new DGAs or adversarial attacks. As expected, the B-RF model trained on both lexical and side information features detects 20% of CharBot domains as DGA/malicious, which is 12% more than the LSTM.MI model. This clearly indicates that the use of side information features to train the DGA classifier makes it more robust against adversarial samples like CharBot domains, when compared to classifiers that rely only on the domain name for DGA detection.

While the inclusion of side information features is clearly beneficial, still $\sim 80\%$ of CharBot domains remain undetected. This means that effective inline detection of CharBot domains remains a largely unsolved problem for which existing techniques from the literature are not sufficient.

VII. CONCLUSION

In this paper, we evaluated state-of-the-art classifiers for inline DGA detection using side information features that are easily obtained from DNS query and response. Results from tables 5 and 8 show that using side information in addition to the domain name to train classifiers not only improves the predictive performance, but also makes it more robust against adversaries like CharBot, when compared to the classifiers that use just the domain name to detect DGAs. Additionally, the side information features in our approach are carefully chosen to perform lightweight *inline* detection of DGA domains, and do not rely on external sources such as WHOIS for feature extraction.

REFERENCES

- [1] D. Plohmann, K. Yakdan, M. Klatt, J. Bader, and E. Gerhards-Padilla, "A comprehensive measurement study of domain generating malware," in *Proc. 25th USENIX Secur. Symp.*, 2016, pp. 263–278. [Online]. Available: https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_plohmann.pdf
- [2] M. Antonakakis, R. Perdisci, Y. Nadji, N. Vasiloglou, S. Abu-Nimeh, W. Lee, and D. Dagon, "From throw-away traffic to bots: Detecting the rise of DGA-based malware," in *Proc. USENIX Secur. Symp.*, vol. 12, 2012, pp. 491–506. [Online]. Available: <https://www.usenix.org/system/files/conference/usenixsecurity12/sec12-final127.pdf>
- [3] S. Schiavoni, F. Maggi, L. Cavallaro, and S. Zanero, "Phoenix: DGA-based botnet tracking and intelligence," in *Proc. Int. Conf. Detection Intrusions Malware, Vulnerability Assessment*. Cham, Switzerland: Springer, 2014, pp. 192–211. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-08509-8_11

- [4] S. Schüppen, D. Teubert, P. Herrmann, and U. Meyer, "FANCI: Feature-based automated nxdomain classification and intelligence," in *Proc. 27th USENIX Secur. Symp.*, 2018, pp. 1165–1181. [Online]. Available: <https://www.usenix.org/system/files/conference/usenixsecurity18/sec18-schuppen.pdf>
- [5] J. Woodbridge, H. S. Anderson, A. Ahuja, and D. Grant, "Predicting domain generation algorithms with long short-term memory networks," 2016, *arXiv:1611.00791*. [Online]. Available: <http://arxiv.org/abs/1611.00791>
- [6] B. Yu, D. L. Gray, J. Pan, M. D. Cock, and A. C. A. Nascimento, "Inline DGA detection with deep networks," in *Proc. IEEE Int. Conf. Data Mining Workshops (ICDMW)*, Nov. 2017, pp. 683–692. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8215728>
- [7] J. Koh and B. Rhodes, "Inline detection of domain generation algorithms with context-sensitive word embeddings," in *Proc. 2018 IEEE Int. Conf. Big Data*, Dec. 2018, pp. 2965–2970. [Online]. Available: <https://arxiv.org/abs/1811.08705>
- [8] D. Tran, H. Mac, V. Tong, H. A. Tran, and L. G. Nguyen, "A LSTM based framework for handling multiclass imbalance in DGA botnet detection," *Neurocomputing*, vol. 275, pp. 2401–2413, Jan. 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0925231217317320>
- [9] J. Saxe and K. Berlin, "Expose: A character-level convolutional neural network with embeddings for detecting malicious URLs, file paths and registry keys," 2017, *arXiv:1702.08568*. [Online]. Available: <https://arxiv.org/abs/1702.08568> and https://www.researchgate.net/publication/314115440_eXpose_A_Character-Level_Convolutional_Neural_Network_with_Embeddings_For_Detecting_Malicious_URLs_File_Paths_and_Registry_Keys
- [10] R. R. Curtin, A. B. Gardner, S. Grzonkowski, A. Kleymenov, and A. Mosquera, "Detecting DGA domains with recurrent neural networks and side information," in *Proc. 14th Int. Conf. Availability, Rel. Secur.*, 2019, pp. 1–10. [Online]. Available: <https://arxiv.org/pdf/1810.02023.pdf>
- [11] T. Chin, K. Xiong, C. Hu, and Y. Li, "A machine learning framework for studying domain generation algorithm DGA-based malware," in *Proc. Int. Conf. Secur. Privacy Commun. Syst.* Cham, Switzerland: Springer, 2018, pp. 433–448. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-01701-9_24
- [12] Y. Li, K. Xiong, T. Chin, and C. Hu, "A machine learning framework for domain generation algorithm DGA-based malware detection," *IEEE Access*, vol. 7, pp. 32765–32782, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8631171>
- [13] C. Choudhary, R. Sivaguru, M. Pereira, B. Yu, A. C. Nascimento, and M. De Cock, "Algorithmically generated domain detection and malware family classification," in *Proc. Int. Symp. Secur. Comput. Commun.* Singapore: Springer, 2018, pp. 640–655. [Online]. Available: https://link.springer.com/chapter/10.1007/978-981-13-5826-5_50
- [14] R. Sivaguru, C. Choudhary, B. Yu, V. Tymchenko, A. Nascimento, and M. D. Cock, "An evaluation of DGA classifiers," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2018, pp. 5058–5067. [Online]. Available: <http://faculty.washington.edu/mdcock/papers/rsivaguru2018a.pdf>
- [15] B. Yu, J. Pan, J. Hu, A. Nascimento, and M. De Cock, "Character level based detection of DGA domain names," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2018, pp. 4168–4175. [Online]. Available: <http://faculty.washington.edu/mdcock/papers/byu2018a.pdf>
- [16] S. Yadav, A. K. K. Reddy, A. L. N. Reddy, and S. Ranjan, "Detecting algorithmically generated domain-flux attacks with DNS traffic analysis," *IEEE/ACM Trans. Netw.*, vol. 20, no. 5, pp. 1663–1677, Oct. 2012. [Online]. Available: <https://ieeexplore.ieee.org/document/6151233>
- [17] L. Bilge, S. Sen, D. Balzarotti, E. Kirida, and C. Kruegel, "Exposure: A passive DNS analysis service to detect and report malicious domains," *ACM Trans. Inf. Syst. Secur.*, vol. 16, no. 4, pp. 1–28, Apr. 2014, doi: 10.1145/2584679.
- [18] P. Lison and V. Mavroudis, "Neural reputation models learned from passive DNS data," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2017, pp. 3662–3671. [Online]. Available: <https://www.nr.no/directdownload/1515491568/neuralreputation-plison.pdf>
- [19] Y. Vorobeychik and M. Kantarcioglu, "Adversarial machine learning," *Synthesis Lectures Artif. Intell. Mach. Learn.*, vol. 12, no. 3, pp. 1–169, 2018. [Online]. Available: <http://lib.uok.ac.ir:8080/multiMediaFile/151136546-4-1.pdf>
- [20] J. Peck, C. Nie, R. Sivaguru, C. Grumer, F. Olumofin, B. Yu, A. Nascimento, and M. De Cock, "CharBot: A simple and effective method for evading DGA classifiers," *IEEE Access*, vol. 7, pp. 91759–91771, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8756038>
- [21] H. S. Anderson, J. Woodbridge, and B. Filar, "DeepDGA: Adversarially-tuned domain generation and detection," in *Proc. ACM Workshop Artif. Intell. Secur.*, 2016, pp. 13–21. [Online]. Available: <https://arxiv.org/abs/1610.01969>
- [22] J. Spooen, D. Preuveeneers, L. Desmet, P. Janssen, and W. Joosen, "Detection of algorithmically generated domain names used by botnets: A dual arms race," in *Proc. 34th ACM/SIGAPP Symp. Appl. Comput.* New York, NY, USA: Association Computing Machinery, Apr. 2019, pp. 1902–1910. [Online]. Available: <https://dl.acm.org/doi/10.1145/3297280.3297467>
- [23] L. Sidi, A. Nadler, and A. Shabtai, "MaskDGA: A black-box evasion technique against DGA classifiers and adversarial defenses," 2019, *arXiv:1902.08909*. [Online]. Available: <http://arxiv.org/abs/1902.08909>
- [24] Y. Fu, L. Yu, O. Hambolu, I. Ozelcik, B. Husain, J. Sun, K. Sapra, D. Du, C. T. Beasley, and R. R. Brooks, "Stealthy domain generation algorithms," *IEEE Trans. Inf. Forensics Security*, vol. 12, no. 6, pp. 1430–1443, Jun. 2017. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/7852496>
- [25] R. Sivaguru, "Hardening inline DGA classifiers against adversarial attacks," M.S. thesis, School Eng. Technol., Univ. Washington, Tacoma WA, USA, 2019.
- [26] M. Pereira, S. Coleman, B. Yu, M. De Cock, and A. Nascimento, "Dictionary extraction and detection of algorithmically generated domain names in passive DNS traffic," in *Proc. Int. Symp. Res. Attacks, Intrusions, Defenses*. Cham, Switzerland: Springer, 2018, pp. 295–314. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-00470-5_14
- [27] A. Joshi, L. Lloyd, P. Westin, and S. Seethapathy, "Using lexical features for malicious URL detection—A machine learning approach," 2019, *arXiv:1910.06277*. [Online]. Available: <http://arxiv.org/abs/1910.06277>
- [28] B. Yu, J. Pan, D. Gray, J. Hu, C. Choudhary, A. C. Nascimento, and M. De Cock, "Weakly supervised deep learning for the detection of domain generation algorithms," *IEEE Access*, vol. 7, pp. 51542–51556, 2019. [Online]. Available: <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8691763>
- [29] Z. Wang, Z. Jia, and B. Zhang, "A detection scheme for DGA domain names based on SVM," in *Proc. Int. Conf. Math., Modeling, Simulation Algorithms (MMSA)*, 2018, pp. 1–7. [Online]. Available: <https://www.atlantis-press.com/proceedings/mmsa-18/25894313>
- [30] J. Kwon, J. Lee, H. Lee, and A. Perrig, "PsyBoG: A scalable botnet detection method for large-scale DNS traffic," *Comput. Netw.*, vol. 97, pp. 48–73, Mar. 2016. [Online]. Available: https://ccs.korea.ac.kr/pds/COMNET16_KWON.pdf
- [31] M. Antonakakis, R. Perdisci, W. Lee, N. Vasiloglou, and D. Dagon, "Detecting malware domains at the upper DNS hierarchy," in *Proc. USENIX Secur. Symp.*, vol. 11, 2011, pp. 1–16. [Online]. Available: https://www.usenix.org/legacy/event/sec11/tech/full_papers/Antonakakis.pdf
- [32] L. Watkins, S. Beck, J. Zook, A. Buczak, J. Chavis, W. H. Robinson, J. A. Morales, and S. Mishra, "Using semi-supervised machine learning to address the big data problem in DNS networks," in *Proc. IEEE 7th Annu. Comput. Commun. Workshop Conf. (CCWC)*, Jan. 2017, pp. 1–6. [Online]. Available: <https://ieeexplore.ieee.org/document/7868376>
- [33] I. Khalil, T. Yu, and B. Guan, "Discovering malicious domains through passive DNS data graph analysis," in *Proc. 11th ACM Asia Conf. Comput. Secur. (ASIA CCS)*, 2016, pp. 663–674. [Online]. Available: <https://dl.acm.org/doi/10.1145/2897845.2897877>
- [34] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker, "Beyond blacklists: Learning to detect malicious Web sites from suspicious URLs," in *Proc. 15th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2009, pp. 1245–1254. [Online]. Available: <http://csweb.ucsd.edu/jtma/papers/beyondbl-kdd2009.pdf>
- [35] M. Singh, M. Singh, and S. Kaur, "Detecting bot-infected machines using DNS fingerprinting," *Digit. Invest.*, vol. 28, pp. 14–33, Mar. 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/abs/pii/S174228761830272X>

- [36] J. Abbink and C. Doerr, "Popularity-based detection of domain generation algorithms," in *Proc. 12th Int. Conf. Availability, Rel. Secur.*, Aug. 2017, pp. 1–8. [Online]. Available: <https://www.cyber-threat-intelligence.com/publications/WMR2017-DGA.pdf>
- [37] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Oct. 2011.
- [38] F. Chollet (2015) *Keras*. [Online]. Available: <https://github.com/fchollet/keras>
- [39] P. Lison and V. Mavroudis, "Automatic detection of malware-generated domains with recurrent neural models," 2017, *arXiv:1709.07102*. [Online]. Available: <http://arxiv.org/abs/1709.07102>




RAAGHAVI SIVAGURU (Member, IEEE) received the B.Tech. degree in information technology from Anna University, India, in 2014, and the M.S. degree in computer science and systems from the University of Washington Tacoma, USA, in 2019. Her M.S. thesis was on improving the predictive performance of DGA detection classifiers and evaluating hardening techniques that use side information to defend against adversaries. She is currently a Software Engineer with Microsoft, Redmond, WA, USA.




JONATHAN PECK received the B.Sc. degree in computer science and the M.Sc. degree in mathematical informatics from Ghent University, Belgium, in 2015 and 2017, respectively. He is currently pursuing the Joint Ph.D. degree with the VIB Inflammation Research Center, Ghent, Belgium, and the Department of Applied Mathematics, Computer Science and Statistics. His research was sponsored by the Fellowship of the Research Foundation Flanders (FWO). His research interest includes improving the robustness of machine learning models to adversarial manipulations.



FEMI OLUMOFIN received the Ph.D. degree in computer science from the University of Waterloo, Canada. He is currently a Senior Member with the Data Science and Analytics Team, Infoblox, San Francisco. He has made contributions to research and development. His research interests include privacy enhancing technologies, security, applied cryptography, big data analytics, and machine learning.



ANDERSON NASCIMENTO received the B.S. degree in electrical engineering from the University of Brasilia, Brazil, in 1998, and the M.S. and Ph.D. degrees in information and communication engineering from The University of Tokyo, Japan, in 2001 and 2004, respectively. He was a Permanent Member with the Nippon Telegraph and Telecom Cryptography Research Group, Japan. He was also a Faculty Member with the University of Brasilia. He is currently an Endowed Associate Professor of information security and information technology with the School of Engineering and Technology, University of Washington Tacoma, USA. His research interests include cryptography, information security, privacy, and machine learning applications.



MARTINE DE COCK received the M.S. and Ph.D. degrees in computer science from Ghent University, Belgium, in 1998 and 2002, respectively. She was a Research Assistant and a Postdoctoral Fellow with the Scientific Research, Flanders, a Visiting Scholar with the BISC Group, University of California, Berkeley, CA, USA, a Visiting Scholar with the Knowledge Systems Laboratory, Stanford University, USA, and an Associate Professor with the Department of Applied Mathematics, Computer Science and Statistics, Ghent University. She is currently a Professor with the School of Engineering and Technology, University of Washington Tacoma, USA, and a Guest Professor with Ghent University. She has over 180 peer-reviewed publications in international journals and conferences on artificial intelligence, data mining, machine learning, information retrieval, web intelligence, and logic programming. Her current research interests include privacy-preserving machine learning, cybersecurity, and data analytics to improve the quality of healthcare. She is a program committee member of numerous international conferences. She has co-organized the KDDCup, in 2013. She served as an Associate Editor for the IEEE TRANSACTIONS ON FUZZY SYSTEMS.

• • •