

Received May 29, 2019, accepted June 26, 2019, date of publication July 5, 2019, date of current version July 25, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2927075

# CharBot: A Simple and Effective Method for Evading DGA Classifiers

JONATHAN PECK<sup>1,2</sup>, CLAIRE NIE<sup>3</sup>, RAAGHAVI SIVAGURU<sup>3</sup>, CHARLES GRUMER<sup>3</sup>, FEMI OLUMOFIN<sup>4</sup>, BIN YU<sup>4</sup>, ANDERSON NASCIMENTO<sup>3</sup>, AND MARTINE DE COCK<sup>1,3</sup>

<sup>1</sup>Department of Applied Mathematics, Computer Science and Statistics, Ghent University, 9000 Ghent, Belgium

<sup>2</sup>Data Mining and Modeling for Biomedicine, VIB Inflammation Research Center, 9052 Ghent, Belgium

<sup>3</sup>School of Engineering and Technology, University of Washington at Tacoma, Tacoma, WA 98402, USA

<sup>4</sup>Infoblox, Santa Clara, CA 95054, USA

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

The work of J. Peck was supported by the Research Foundation Flanders (FWO).

**ABSTRACT** Domain generation algorithms (DGAs) are commonly leveraged by malware to create lists of domain names, which can be used for command and control (C&C) purposes. Approaches based on machine learning have recently been developed to automatically detect generated domain names in real-time. In this paper, we present a novel DGA called CharBot, which is capable of producing large numbers of unregistered domain names that are not detected by state-of-the-art classifiers for real-time detection of the DGAs, including the recently published methods FANCI (a random forest based on human-engineered features) and LSTM.MI (a deep learning approach). The CharBot is very simple, effective, and requires no knowledge of the targeted DGA classifiers. We show that retraining the classifiers on CharBot samples is not a viable defense strategy. We believe these findings show that DGA classifiers are inherently vulnerable to adversarial attacks if they rely only on the domain name string to make a decision. Designing a robust DGA classifier may, therefore, necessitate the use of additional information besides the domain name alone. To the best of our knowledge, the CharBot is the simplest and most efficient black-box adversarial attack against DGA classifiers proposed to date.

**INDEX TERMS** Adversarial machine learning, domain generation algorithms, supervised learning.

## I. INTRODUCTION

The purpose of distributing malware is often to extract sensitive information from victim machines or to use them for disseminating spam. To achieve this, botmasters need to be able to communicate with the infected machines, which is done via command-and-control (C&C) servers. The use of a fixed pool of C&C servers is not attractive, however, since these servers may be taken offline or blacklisted. Therefore, malware authors design *domain generation algorithms* or DGAs to automatically create many domain names that are likely to be unregistered and hence available for the malware to establish a communication channel [1]. A DGA makes use of a *seed*, i.e., some random number that is accessible to both the botmaster and the malware on the infected machines.

Possible seeds include the current date, trending topics on Twitter, weather forecasts, etc. Once this seed has been fixed,

The associate editor coordinating the review of this manuscript and approving it for publication was Irene Amerini.

the botmaster, as well as all of the infected machines, can generate the same list of domains. The botmaster registers one of these domains and waits for the malware to successfully resolve a DNS query against it. From that point on, communication can take place. Should the C&C server ever be taken offline or have its domain blacklisted, this process can simply be restarted and a new C&C server can be established.

An extensive amount of research in the past decade has been devoted to the development of methods for detection of domains generated by DGAs [2]–[5]. These methods can be roughly divided into two classes: classifiers that detect DGAs based solely on the domain name itself; and classifiers that use some sort of context information, such as IP addresses of the source, traffic, and query patterns by the infected machines. Our focus in this paper is on the first kind of classifiers, i.e. techniques that can detect DGA domains in real-time based on the domain name string. These systems are particularly attractive since additional information beyond the domain name string can be expensive to acquire. It might

also simply not be available due to privacy concerns. Another significant advantage of systems that perform DGA detection based solely on the domain name is their potential use in real-time systems, blocking malicious domains before they are actually resolved. Accordingly, much research has been carried out to prevent this type of C&C communication using systems that can detect in real-time whether a domain name is likely generated by a DGA or not [2], [4]–[13].

Such DGA classifiers need to be sufficiently robust so that they can still reliably detect DGA domains even when the DGAs start generating lists from seeds that were not seen during training. Existing work in this area is comprised both of methods that make use of human-engineered features as well as deep learning techniques which learn to extract relevant features automatically. We show in this paper that both kinds of methods are inherently vulnerable to simple attacks and hence the use of side information may be crucial to developing robust DGA classifiers. Specifically, we introduce a new and effective DGA called *CharBot*. It is a simplistic character-based DGA (hence the name) that generates domain names by randomly modifying two characters in well known benign domains collected from the Alexa top domain names.<sup>1</sup> We find that the domains CharBot generates are almost always unregistered, hence available for C&C communication.

To demonstrate CharBot’s capabilities, we attack two types of recently proposed prototypical DGA classifiers that are considered state-of-the-art at the time of this writing: a random forest (RF) model called *FANCI* based on human-engineered features extracted from the domain name [11] and a deep neural network (DNN) model called *LSTM.MI* [5]. We also test a RF approach called *B-RF* based on the features proposed in [12]. We train these models on data sets consisting of benign and malicious domain names. The benign names originate from the Alexa top domain names. For the malicious domains, we use the OSINT Bambenek Consulting feeds.<sup>2</sup> We find that the domain names generated by CharBot go largely undetected by all these state-of-the-art DGA classifiers.

We attempt to harden the classifiers against CharBot by incorporating samples from it in the training data sets and retraining the models. Although this strategy does increase the detection rates, they are still not high enough to be practical. We also try retraining using samples generated by DeepDGA — a state of the art generative model for malicious domain names [14] — as well as the DeceptionDGA by Spooren *et al.* [15], but we find that this does not adequately help with detecting CharBot. CharBot is much simpler than both DeepDGA and DeceptionDGA: DeepDGA is a deep learning approach, whereas CharBot performs only simple string manipulations; DeceptionDGA is designed to evade classifiers based on human-engineered features. By contrast,

CharBot is fully black-box: it does not require any details of the models being attacked.

CharBot works by corrupting domain names from the Alexa top domains, so it is natural to ask whether the domains it generates can also be used to successfully attack DGA classifiers that do not depend on Alexa for training. To answer this question, we investigate whether the DGA classifiers can be hardened by replacing the Alexa data set by an alternative data set of benign domains during training. To this end, we use a data set of domain names that occurred in real DNS traffic, weakly labeled according to heuristic rules [7]. We find that training on this different data set yields approximately the same results as when training on Alexa. This supports the idea that CharBot attacks are transferable across models and data sets.

These findings expose a dangerous weakness in modern DGA classifiers: they can be circumvented using a simple algorithm and they cannot be easily trained to detect it well. We speculate that this weakness is inherent in any model that relies solely on domain name strings to perform DGA classification. CharBot works by introducing a small number of typographical errors in benign domain names from the Alexa data set. As such, the statistical properties of the names it generates will be almost identical to those of the Alexa domains. This makes it nearly impossible for a classifier to draw any significant distinction between Alexa names and CharBot names. Moreover, any other set of legitimate domains that should be accepted by a classifier with high probability could in principle be used instead of Alexa by a CharBot attack. Therefore, we do not believe these attacks can be mitigated without relying on additional side information. Such information might include the IP addresses the domains resolve to, how many times the domains were queried and when, etc. This has been explored in other works already [3], [10], [16]–[18]. To our knowledge, we are the first to expose this type of weakness in DGA classifiers that do not use side information. We would, therefore, recommend that the community focuses its research efforts on DGA classifiers that utilize side information and not just rely on the domain name string by itself.

The rest of this paper is structured as follows. Section II gives an overview of related work in the field of adversarial machine learning. Section III details the CharBot algorithm. Section IV describes the data sets we used for the experiments. Section V outlines our experiments and discusses their results, as well as several ways we could defend against CharBot attacks. Section VII concludes the work and lists some possibilities for future research.

## II. RELATED WORK

Machine learning approaches that leverage the domain name string for DGA detection can be categorized into two groups: so-called “featureful” methods that rely on human defined lexical features extracted from the domain names, such as domain name length, vowel-character ratio, bigrams, etc. [2], [11], [16] and “featureless” methods in which the

<sup>1</sup><https://alexa.com/topsites>. Accessed: 2019-02-10.

<sup>2</sup><http://osint.bambenekconsulting.com/feeds/>. Accessed: 2019-02-10.

automatic discovery of good features is part of the overall machine learning model training process, as a form of representation learning [4]–[7], [9], [19]. Popular kinds of classifiers used in the featureful approach for DGA detection are logistic regression and tree ensemble methods, while the featureless approach relies on the use of deep neural networks, namely Long Short-Term Memory (LSTM) networks and Convolutional Neural Networks (CNN). Most papers about the featureless approach include a featureful approach as a baseline method [4]–[8], [12], [13], [19], and the featureless approach is typically reported to yield better, more accurate results.

A natural response of malware authors to machine learning classifiers for DGA detection is to try to purposely craft domain names that will be mislabeled as benign by the classifiers. This kind of evasion attack is studied as part of the broader field of *adversarial machine learning* (AML) [20]. In this setting, an intelligent adversary aims to exploit weaknesses in a machine learning model in order to obtain desired (illegitimate) outcomes. A prototypical example is that of spam classification, where the adversary attempts to craft spam e-mails that evade detectors while still achieving the desired results. Seminal contributions in this area include the work of Dalvi *et al.* [21] as well as the papers by Lowd and Meek [22], [23], who study classical machine learning algorithms such as linear classifiers, naive Bayes, support vector machines and maximum entropy filters. More recent works primarily study AML for deep neural networks [24]–[26].

A recent innovation in the area of deep learning and generative modeling, is the *Generative Adversarial Network* or GAN, first proposed by Goodfellow *et al.* [27]. In the GAN framework, a generative model is trained by pitting it against an adversary. The adversary is a discriminative model whose goal is to discern whether a given sample came from the data generating distribution or from the generative model. The generator is trained to maximize the loss of the discriminator, so the GAN training procedure corresponds to a two-player minimax game. Ideally, when the training converges, the generator should recover the data generating distribution and the discriminator should not be able to do any better than random guessing.

GANs have found several uses in cybersecurity by now. Anderson *et al.* [14] proposed *DeepDGA*, which is a generative model for DGA domains trained using a GAN. They find that adding samples from DeepDGA to the training data of deep learning based DGA classifiers improves their performance against unseen malware families, aiding generalization of the models when insufficient training data is available. In the field of password security, Hitaj *et al.* [28] have proposed *PassGAN*, another generative model trained in the GAN framework. PassGAN learns to capture the distribution of human passwords and is able to surpass state of the art tools for password guessing. Hu and Tan [29] recently proposed *MalGAN*, a GAN with which they are able to construct malware samples that can bypass black-box machine

learning methods. Their attack is particularly striking because they are able to reduce malware detection rates to almost zero without requiring direct access to the detectors they aim to evade. Moreover, they found that explicitly retraining the detectors on MalGAN samples is ineffective: MalGAN can easily be adapted to take this retraining into account, bypassing the retrained models again with almost 100% success. With CharBot, we achieve similar (and, in several cases, better) results with a much simpler approach that can actually be incorporated within a piece of malware, in contrast to deep-learning based methods which are usually too large or too computationally intensive.

Several authors have recently looked into the automatic generation of URLs for phishing. To this end, Bahnsen *et al.* [30] create a text consisting of known phishing URLs from PhishTank<sup>3</sup> and use it to train an LSTM for text generation, i.e. given a small seed sentence, predict the next characters iteratively. They report that this technique generates examples that are not detected by their own LSTM based phishing URL classifier [31]. Anand *et al.* [32] trained a GAN – containing a character based LSTM as part of its architecture – to generate synthetic phishing URLs to augment the training data for feature-based phishing URL detection classifiers. The problem they address is the class imbalance in typical training data sets, which contain many more examples of benign URLs than of phishing URLs. Instead of adding all generated phishing URLs as positive examples to their training data, they first map the generated URLs to their corresponding feature vectors, and select “representative samples” based on Euclidean distance in this feature space. In a similar vein to Anand *et al.*, Burns and Heath [33] train a GAN on OpenPhish,<sup>4</sup> PhishTank and DNS-BH<sup>5</sup> data sources to develop synthetic phishing domains. They compare a random forest classifier trained on Alexa and Umbrella<sup>6</sup> data sets to models that were augmented with samples generated by the GAN. They find that the augmented models appear to have consistently higher test set accuracy than the original classifier.

URLs intended for phishing are quite different in nature than DGA domains for C&C purposes. Indeed, to be successful, phishing URLs need to deceive humans, which requires them to be as indistinguishable as possible from benign URLs to the human observer. DGA domain names used for C&C purposes are not intended at all to be read by human users. DGA domain names are successful if they can evade DGA classifiers and have not been previously registered, i.e. they should be available for the botmaster to register. To the best of our knowledge, so far Anderson *et al.* are the only ones who have looked into generative modeling of DGA domain names [14]. Although their results are significant, we show in this work that classifiers which have been adversarially

<sup>3</sup><https://www.phishtank.com/>. Accessed: 2019-02-08.

<sup>4</sup><https://openphish.com/>. Accessed: 2019-02-08.

<sup>5</sup><http://www.malwaredomains.com/>. Accessed: 2019-02-08.

<sup>6</sup><http://s3-us-west-1.amazonaws.com/umbrella-static/index.html>. Accessed: 2019-02-08.

trained using DeepDGA remain vulnerable to simple attacks such as the CharBot algorithm we propose in section III.

The CharBot algorithm is a *black-box targeted evasion attack* that works against tree ensembles and neural networks. “Black-box” refers to the fact that our CharBot DGA does not require details of the classifiers in order to work: it can attack any model trained on any data set and succeed with high probability. It is a targeted attack because we want the classifiers to output a specific class in response to our DGA samples, namely *benign*. Untargeted attacks, on the other hand, merely aim to change the classification to *any* class other than the original; for example, an untargeted attack would also count a change from benign to malicious as a success, whereas in our scenario that would be unacceptable. Finally, CharBot is an evasion attack because it occurs at test time, when the model is already trained and deployed. This is in contrast to *poisoning attacks* which occur at training time and work by corrupting samples in the training data set in order to deliberately introduce weaknesses into the model [20]. The CharBot attack itself is inspired by *typosquatting*, a well-known technique used by phishers and social engineers [34]. Typosquatting involves taking a legitimate domain and introducing a few typographical errors that are unlikely to be noticed by human users (e.g., changing *google.com* into *g0ogle.com*). Whereas *a priori* one would think that typosquatting exploits an inherently human vulnerability, we make the surprising discovery that state-of-the-art DGA classifiers are actually vulnerable to such techniques as well.

Similarly to our work here, Spooren *et al.* [15] developed *DeceptionDGA*, a novel DGA which incorporates knowledge of the features used by a DGA classifier in order to attack it. They report significant reductions in predictive accuracy for the FANCI model as well as the Endgame LSTM by Woodbridge *et al.* [4]. The DeceptionDGA algorithm is more complicated than CharBot, requiring knowledge of the underlying model in order to deploy it. Despite this difference in complexity, the detection rates we observe for CharBot in our experiments are comparable to those of DeceptionDGA.

We also wish to acknowledge the concurrent work of [35] who describe *MaskDGA*, a black-box technique for evading DGA classifiers that is similar to CharBot. MaskDGA makes use of a surrogate model as well as a list of DGA domains. It uses these data to craft character-level perturbations of the malicious domains such that they are no longer recognized by the surrogate model. Similarly to our own results, the authors find that such techniques are highly effective at reducing the accuracy of state-of-the-art DGA classifiers. They also make the recommendation that DGA classifiers should rely on additional side-information whenever this is possible in order to mitigate adversarial attacks.

### III. CHARBOT

CharBot is a character-based DGA intended to show how successful a simplistic DGA based on small perturbations can be at evading detection by state-of-the-art classifiers.

Without loss of generality, throughout this paper we consider domains consisting of a second-level domain (SLD) and a top-level domain (TLD), separated by a dot, as in e.g. *wikipedia.org*. CharBot requires the following inputs:

- A list of legitimate domain names. In our case, ten thousand Alexa domains with a second-level domain (SLD) length of six or greater are used.
- A list of top-level domains (TLDs).
- A date to be used as a seed for pseudorandomization.

With these inputs, CharBot (1) selects a domain from the provided list, (2) selects two characters from the SLD, and (3) selects two replacement characters. The replacement characters are chosen from an equal distribution of DNS-valid characters — the alphanumeric characters and the dash — and the algorithm ensures the characters selected from the SLD are different from the replacement characters. Finally, CharBot (4) appends a TLD to the new domain by selecting one of the following: *com, at, uk, pl, be, biz, co, jp, cz, de, eu, fr, info, it, ru, lv, me, name, net, nz, org, us*. Pseudocode is given in algorithm 1.

---

#### Algorithm 1 CharBot

---

**Data:** a list of SLDs  $D$ , a list of TLDs  $T$ , a seed  $s$

**Result:** a DGA domain

- 1 Initialize the pseudorandom generator with the seed  $s$ .
  - 2 Randomly select a SLD  $d$  from  $D$ .
  - 3 Randomly select two indices  $i$  and  $j$  so that  $1 \leq i, j \leq |d|$ .
  - 4 Randomly select two replacement characters  $c_1$  and  $c_2$  from the set of DNS-valid characters.
  - 5 Set  $d[i] \leftarrow c_1$  and  $d[j] \leftarrow c_2$ .
  - 6 Randomly select a TLD  $t$  from  $T$ .
  - 7 **return**  $d.t$
- 

A DGA is successful if it can generate many unique domains that have not yet been registered and which are not flagged by DGA classifiers as malicious. CharBot draws its replacement characters from a uniform distribution. Therefore, the more characters we replace, the more the generated domains resemble random strings. This increases the detection rate by DGA classifiers, so we aim to keep the number of replacement characters minimal. We tested several choices for the number of characters to be replaced. We found that two characters strike an appropriate balance between the rate of detection by DGA classifiers and the probability that a domain is already registered: with two characters, domains are flagged slightly more often but almost all domains are unregistered (see table 1); when replacing only a single character, detection rates go down but more domains turn out to be registered already.

Adversarial attacks such as CharBot are always accompanied by an *adversarial cost function*  $c(x, \tilde{x})$  which describes the cost associated with perturbing an “ideal” sample  $x$  into a sample  $\tilde{x}$  that the adversary can actually use. For image classification, it is common to use  $\ell_p$  distances for this

purpose [24], [26], [36]. However, in our context, the cost of perturbing a correctly classified benign domain  $x$  into a malicious domain  $\tilde{x}$  that is classified as benign must be measured differently, as we are working in a discrete input space (text) instead of a continuous one (images). Specifically, it makes sense to define our cost function as follows:

$$c(x, \tilde{x}) = \begin{cases} d_L(x, \tilde{x}) & \text{if } \tilde{x} \text{ is unregistered,} \\ \infty & \text{otherwise.} \end{cases}$$

Here,  $d_L$  denotes the *Levenshtein distance* or *edit distance* [37]. The cost function  $c(x, \tilde{x})$  increases with the number of edits (insertions, deletions, and substitutions) required to transform  $x$  into  $\tilde{x}$ , as each edit makes the attack more detectable by DGA classifiers. However, there is an infinitely large cost associated with generating a domain that is already registered, since such domains cannot be used by the attacker at all and may cause the malware to malfunction. CharBot was designed to minimize this cost function efficiently and as simply as possible.

We note that obvious extensions of the CharBot algorithm are possible. For example, we could additionally implement insertions of new characters and deletions of existing characters. The number of characters might also be chosen adaptively based on some heuristic instead of fixed in advance. However, we chose to limit ourselves to only substituting a fixed number of characters since this simple strategy already gives us very good results. Moreover, simpler attacks are likely to be preferred by attackers and therefore constitute a greater security concern.

The only obstacle to the deployment of CharBot in real malware might be its size. We implemented CharBot in 17,983 bytes of Python code. The Alexa data set it requires takes up 145,008 additional bytes, although the public availability of this data set means it could be downloaded on the fly. Therefore, we would need at most 162,991 bytes for a full implementation of CharBot with Alexa included. By comparison, the DeepDGA algorithm [14] requires to embed in the malware a trained machine learning model that takes up at least 6,539,192 bytes. This is about 40× larger than CharBot. We therefore feel that file size is no obstacle to deploying CharBot in real malware.

#### IV. DATA SETS

We use three different kinds of data in our experiments:

##### A. ALEXA

The top 1 million unique domain names from Alexa.<sup>7</sup> Alexa ranks websites based on their popularity in terms of the number of page views and number of unique visitors. It only retains the websites' SLD and TLD, aggregating across any subdomains. For example, according to Alexa, the five highest ranked domain names in terms of popularity on 2019-02-06 are *google.com*, *youtube.com*, *facebook.com*, *baidu.com* and *wikipedia.org*. It is generally assumed that

<sup>7</sup><https://www.alexa.com>, Accessed 2019-02-08.

the top 1 million domain names in the Alexa ranking are “benign” domain names in the sense that they were not created by a DGA. Of course, this does not mean that the domain is “benign” in the sense of not being used for malicious activity. Indeed, there is reason to believe that a significant number of Alexa top ranked domains are used for malicious purposes [38], but this is not the problem we are considering here. In our setting, we only consider a domain to be “malicious” if it was generated by a DGA.

##### B. BAMBENEK

1 million unique DGA domain names from the Bambenek Consulting feeds<sup>8</sup> for 3 different days, namely Jun 24, Jul 22, Jul 23, 2017. These feeds contain DGA domain names from specific malware families that were observed in real traffic on those days. Such domain names can be collected by reverse engineering a known malware family, generating lists of domain names with the reverse engineered malware, and checking which of these domain names also occur in real traffic.

##### C. QNAME

1 million unique domain names originating from a real-time stream of passive DNS data that consists of roughly 10-12 billion DNS queries per day collected from subscribers including ISPs (Internet Service Providers), schools, and businesses. We annotated this stream based on a set of heuristic filtering rules following [7]. Specifically, we labeled as benign all domains that have been resolved at least twice, never resulted in an NXDomain response and span more than 30 days. Here, span is defined as the number of days between the first and last successfully resolved query for a given domain. We randomly sampled 1 million such domains that appeared in DNS traffic between September 2015 and August 2018. This data set is *weakly labeled* since the heuristic filtering rules do not guarantee that the domains are actually benign or malicious; however, we believe it to be a useful approximation.

The Alexa and Qname data sets serve as our negative (benign) examples, whereas Bambenek serves as our set of positive (malicious) examples. Alexa and Qname have precisely 537 domains in common, which is a negligible number compared to the total sizes of the data sets, therefore making Qname a good data set to test transferability of CharBot.

We refer to the combination of Alexa and Bambenek data as *AlexaBamb* and similarly for *QnameBamb*. These data sets consist of 2 million samples each, 1 million per class.

#### V. EXPERIMENTS

We perform experiments on two DGA classifiers that are considered state of the art at the time of this writing: FANCI [11] and LSTM.MI [5], as well as a third model we call B-RF based on the work by [12]. All classifiers are trained to label a domain name as either benign (negative) or

<sup>8</sup><http://osint.bambenekconsulting.com/feeds/>, Accessed 2019-02-08

TABLE 1. Adversarial data sets.

DGA	Data Set	Seeds Used	# Unique Synthetic Domains	# Unregistered Domains (out of 500 sampled)
CharBot	Training	2018-12-04	100,000	500 (100%)
	Testing	2019-01-01	10,000	500 (100%)
DeepDGA	Training	2018-12-04	100,000	499 (99.8%)
	Testing	2019-01-01	10,000	499 (99.8%)
DeceptionDGA	Training	N/A	100,000	494 (98.8%)
	Testing	N/A	10,000	494 (98.8%)

malicious (positive). We find that the best results overall are achieved with the deep learning based LSTM.MI approach, followed by the random forest-based B-RF approach, and finally the random forest-based FANCI method. The difference in predictive accuracy between the various approaches is substantial. The results hold across the AlexaBamb and QnameBamb data sets (see table 4, and a more detailed discussion in section V-D).

To arrive at the results, we train on the AlexaBamb data set as well as on the QnameBamb data set with a 80%/20% train/test split for each, reporting the true positive rate (TPR), the partial area under the ROC curve (AUC) and the fraction of samples from CharBot, DeepDGA and DeceptionDGA which the models were able to detect (see table 4 and table 5). All of these metrics are reported at FPRs of 0.1% and 1%.<sup>9</sup> The AUC@0.1%FPR is the integral of the ROC curve from FPR = 0 to FPR = 0.001 on the test data, and similarly for the AUC@1%FPR. We repeat all experiments on the original models as well as the models after adversarial retraining.

To perform the adversarial retraining, we utilized the data sets shown in table 1. Specifically, we used CharBot and DeepDGA with different seeds to generate training and testing data sets. The training data sets were used to augment the original training data of the classifiers; the testing data sets were used to verify their performance. For DeceptionDGA, Spoores et al. [15] supplied a list of 150,000 domains generated by their algorithm from which we sampled our training and testing data. Note that, based on a random sample of 500 domains,<sup>10</sup> CharBot has the highest fraction of unregistered domains (100%), followed by DeepDGA (99.8%) and DeceptionDGA (98.8%).

The experiments on the QnameBamb data set are intended to investigate the transferability of CharBot. All CharBot domain names used in the experiments (see table 1) are created by CharBot by corrupting domain names from the Alexa data set. This might leave DGA classifiers that are trained on AlexaBamb extra vulnerable to CharBot attacks. A natural question to ask is whether CharBot can also successfully bypass DGA classifiers that were trained on a data

<sup>9</sup>A low false positive rate is very important in deployed DGA detection systems because blocking legitimate traffic is highly undesirable. The threshold of 0.1% FPR was chosen because this rate is often used by real-world models in practice, whereas 1% is the largest FPR that could still be useful.

<sup>10</sup>We limited ourselves to a random sample of 500 domains to avoid getting blocked by ISPs.

set different from Alexa, one CharBot has no access to. To test this, we trained LSTM.MI, FANCI, and B-RF on the QnameBamb data and reported the same statistics as for AlexaBamb.

Below we give a brief description of the LSTM.MI, FANCI, and B-RF classifiers, followed by detailed results (section V-D) and a discussion of possible countermeasures for defending against small perturbations attacks such as CharBot (section VI).

#### A. LSTM.MI

Woodbridge *et al.* [4] 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.* [5]. Its architecture is very similar to that of Woodbridge *et al.* [4]; 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 [5], [12]). The code for training the LSTM.MI model is publicly available.<sup>11</sup>

#### B. FANCI

The FANCI classifier recently proposed by Schüppen *et al.* [11] is a random forest (RF) classifier designed to classify NXDomains as benign (bNXD) or malicious (mAGD). NXDomains, or Non-Existent Domains, are domains that can not be resolved. DGAs generate hundreds or even thousands of domains every day, only very few of which are actually registered by the botmaster. That means that almost all queries for DGA generated domains by infected machines will result in an NXDomain response by the local DNS server, so it is reasonable to attempt to detect DGA activity by analyzing NXDomains.

To this end, the FANCI classifier leverages 21 manually defined features, extracted from the domain name string. The 21 features can be divided into structural, linguistic, and statistical categories (see table 2). The FANCI RF model is comprised of 9 decision trees, of which 7 use the Gini coefficient as the measure of impurity and the other 2 use

<sup>11</sup><https://github.com/bkcs-hust/lstm-mi>. Accessed: 2019-02-08.

**TABLE 2. Features used by FANCI and B-RF. (\*) For these features, FANCI uses dot free public-suffix-free domain. (\*\*) For these features, FANCI uses public-suffix-free domain.**

#	Feature	FANCI	B-RF
1	Domain name length	✓	✓
2	Second level domain length	✗	✓
3	Top level domain length	✗	✓
4	Domain Unique Characters length	✗	✓
5	SLD Unique Characters length	✗	✓
6	TLD Unique Characters length	✗	✓
7	Has malicious TLD	✗	✓
8	Has Valid TLD	✓	✗
9	TLD Hash	✗	✓
10	Contains Digits	✓	✗
11	Starts with Digit	✗	✓
12	Underscore Ratio*	✓	✗
13	Symbol ratio	✗	✓
14	Hex ratio	✗	✓
15	Digit Ratio*	✓	✓
16	Vowel Ratio*	✓	✓
17	Consonant Ratio	✗	✓
18	Ratio of Repeated Characters*	✓	✓
19	Ratio of Consecutive Consonants*	✓	✓
20	Ratio of Consecutive Digits*	✓	✓
21	Number of tokens in SLD	✗	✓
22	Number of digits in SLD	✗	✓
23	Entropy*	✓	✓
24	Gini Index	✗	✓
25	Classification error of characters	✗	✓
26	N-Gram Distribution*	✓	✗
27	2-Gram Median	✗	✓
28	3-Gram Median	✗	✓
29	2-Gram Circle Median	✗	✓
30	3-Gram Circle Median	✗	✓
31	Number of Subdomains**	✓	✗
32	Subdomain Length Mean**	✓	✗
33	Has www Prefix	✓	✗
34	Contains Single-Character Subdomain**	✓	✗
35	Is Exclusive Prefix Repetition	✓	✗
36	Contains TLD as Subdomain**	✓	✗
37	Ratio of Digit-Exclusive Subdomains**	✓	✗
38	Ratio of Hexadecimal-Exclusive Subdomains**	✓	✗
39	Contains IP Address**	✓	✗
40	Alphabet Cardinality*	✓	✗

**TABLE 3. FANCI features not expected to have any effect in our experiments.**

#	Feature
31	Number of Subdomains
33	Has www Prefix
34	Contains Single-Character Subdomain
35	Is Exclusive Prefix Repetition
36	Contains TLD as Subdomain


entropy. Each tree takes between 2 to 18 features. The source code of the FANCI classifier is available on GitHub.<sup>12</sup>

The domain names used in our experiments contain only SLDs and TLDs (see section IV). As such, it is expected that a number of features used in the FANCI model would not make a distinction between malicious and benign examples. Table 3 lists the FANCI features that are not expected to have any effect.

**C. B-RF**

B-RF [12] is a random based DGA detection classifier that is trained on 26 manually engineered features as indicated in table 2. There is some overlap between the features used by FANCI and those used by B-RF. For instance, both make use of the domain name length, digit and vowel ratio, ratio of repeated characters, etc. Some features are used by FANCI but not by B-RF, such as whether the domains have valid TLDs or whether they contain digits. Other features like 2-gram median and 3-gram median are only used by B-RF.

B-RF consists of 100 trees and each tree is trained using a subset with a maximum of 20 features. Entropy is used as the criterion to decide the split attribute while growing the trees in the random forest.



**FIGURE 1. ROC curves for the classifiers trained on the AlexaBamb data set.**

**D. RESULTS**

The predictive performance metrics are summarized in table 4 for false positive rates<sup>13</sup> of 0.1% and 1%. Figure 1 shows ROC curves for the different models on the AlexaBamb data. We plot the ROC curve only for FPRs between 0 and 0.01, as higher FPRs are meaningless in practice. We conclude from these results that the deep learning approach does better than the RF approaches, which is in line with what has been reported before in the literature [4]–[7], [13]. Among the RF models, B-RF outperforms FANCI significantly. We found that this improvement was not due to the number of trees, as decreasing the number of trees used by B-RF from 100 to 9 (as in FANCI) still yielded superior performance for B-RF. We therefore believe this difference in performance is caused by the different feature sets.

We were unable to establish a classification threshold that achieves 0.1% FPR for FANCI. Therefore, in reporting FANCI results, we only consider FPR = 1%. We believe this is due to the fact that [11] used proprietary data to filter classification outcomes which increased accuracy. Our use of a different post-filtering data set may be the

<sup>13</sup>One can argue that even a FPR of 0.1% is still too high to be useful in practice. While this can certainly be true depending on the application, note that lower FPRs can only make our results better as the models will necessarily have lower TPR and lower detection rates for CharBot.

<sup>12</sup><https://github.com/fanci-dga-detection/fanci>. Accessed: 2019-02-08.

TABLE 4. Performance metrics of LSTM.MI, FANCI and B-RF on the different data sets.

Classifier	Data set	FPR=0.001		FPR=0.01	
		TPR	AUC	TPR	AUC
LSTM.MI	AlexaBamb	96.79%	94.91%	99.27%	98.89%
	AlexaBamb + CharBot	95.50%	95.35%	98.89%	98.67%
	AlexaBamb + DeepDGA	96.65%	96.44%	99.20%	99.00%
	AlexaBamb + DeceptionDGA	95.17%	95.05%	98.54%	98.46%
	QnameBamb	81.98%	83.37%	98.98%	96.68%
	QnameBamb + CharBot	82.91%	84.98%	98.51%	96.48%
	QnameBamb + DeepDGA	83.22%	83.98%	98.85%	96.50%
	QnameBamb + DeceptionDGA	84.66%	85.57%	98.61%	96.82%
FANCI	AlexaBamb	—	—	74.46%	80.46%
	AlexaBamb + CharBot	—	—	72.49%	78.71%
	AlexaBamb + DeepDGA	—	—	73.98%	80.02%
	AlexaBamb + DeceptionDGA	—	—	73.84%	80.18%
	QnameBamb	—	—	74.13%	79.06%
	QnameBamb + CharBot	—	—	72.13%	77.80%
	QnameBamb + DeepDGA	—	—	72.89%	78.19%
	QnameBamb + DeceptionDGA	—	—	73.65%	78.86%
B-RF	AlexaBamb	85.72%	82.93%	94.72%	94.67%
	AlexaBamb + CharBot	84.62%	79.30%	93.89%	93.75%
	AlexaBamb + DeepDGA	85.81%	80.94%	94.39%	94.35%
	AlexaBamb + DeceptionDGA	84.51%	78.62%	93.78%	93.73%
	QnameBamb	82.75%	73.85%	96.88%	94.52%
	QnameBamb + CharBot	83.03%	74.98%	96.37%	94.29%
	QnameBamb + DeepDGA	82.64%	73.87%	96.39%	94.26%
	QnameBamb + DeceptionDGA	82.98%	76.26%	96.26%	94.56%

TABLE 5. Detection rates of the different DGAs.

Classifier	Data set	FPR=0.001			FPR=0.01		
		CharBot	DeepDGA	DeceptionDGA	CharBot	DeepDGA	DeceptionDGA
LSTM.MI	AlexaBamb	5.58%	33.98%	4.02%	15.50%	39.53%	12.74%
	AlexaBamb + CharBot	55.19%	92.54%	19.69%	81.08%	98.44%	47.67%
	AlexaBamb + DeepDGA	12.39%	98.35%	7.34%	12.39%	98.35%	7.34%
	AlexaBamb + DeceptionDGA	23.59%	88.71%	40.29%	52.18%	96.66%	71.52%
	QnameBamb	15.25%	6.18%	16.61%	31.90%	19.51%	37.73%
	QnameBamb + CharBot	52.67%	42.48%	34.45%	81.96%	85.90%	66.27%
	QnameBamb + DeepDGA	27.84%	94.51%	24.25%	43.28%	99.61%	47.33%
	QnameBamb + DeceptionDGA	30.45%	15.97%	37.74%	53.31%	37.97%	24.25%
FANCI	AlexaBamb	—	—	—	3.05%	6.33%	1.66%
	AlexaBamb + CharBot	—	—	—	22.26%	12.12%	2.64%
	AlexaBamb + DeepDGA	—	—	—	6.45%	83.17%	2.08%
	AlexaBamb + DeceptionDGA	—	—	—	4.27%	6.57%	2.33%
	QnameBamb	—	—	—	21.43%	5.37%	46.85%
	QnameBamb + CharBot	—	—	—	48.44%	14.20%	49.62%
	QnameBamb + DeepDGA	—	—	—	45.13%	77.88%	50.11%
	QnameBamb + DeceptionDGA	—	—	—	44.75%	13.77%	50.45%
B-RF	AlexaBamb	1.69%	6.61%	1.38%	27.59%	23.97%	22.37%
	AlexaBamb + CharBot	1.84%	9.19%	1.20%	64.33%	34.06%	41.22%
	AlexaBamb + DeepDGA	4.54%	46.55%	2.86%	31.80%	84.12%	26.86%
	AlexaBamb + DeceptionDGA	2.00%	8.23%	1.34%	33.14%	24.51%	32.53%
	QnameBamb	18.80%	2.99%	41.67%	61.05%	22.57%	62.31%
	QnameBamb + CharBot	43.47%	12.54%	49.18%	85.82%	60.70%	79.33%
	QnameBamb + DeepDGA	44.04%	33.08%	49.80%	65.97%	98.84%	67.79%
	QnameBamb + DeceptionDGA	39.97%	10.53%	47.86%	62.29%	23.30%	67.74%

reason for the difference (note that the authors of the FANCI paper use accuracy whereas we use AUC). For AlexaBamb and its augmented datasets, we obtained between 90.7%

and 91.4% accuracy. For QnameBamb and its augmented datasets, we obtained between 92.5% and 93.16% accuracy. This is not too far from the 93.7% that was reported in [15]



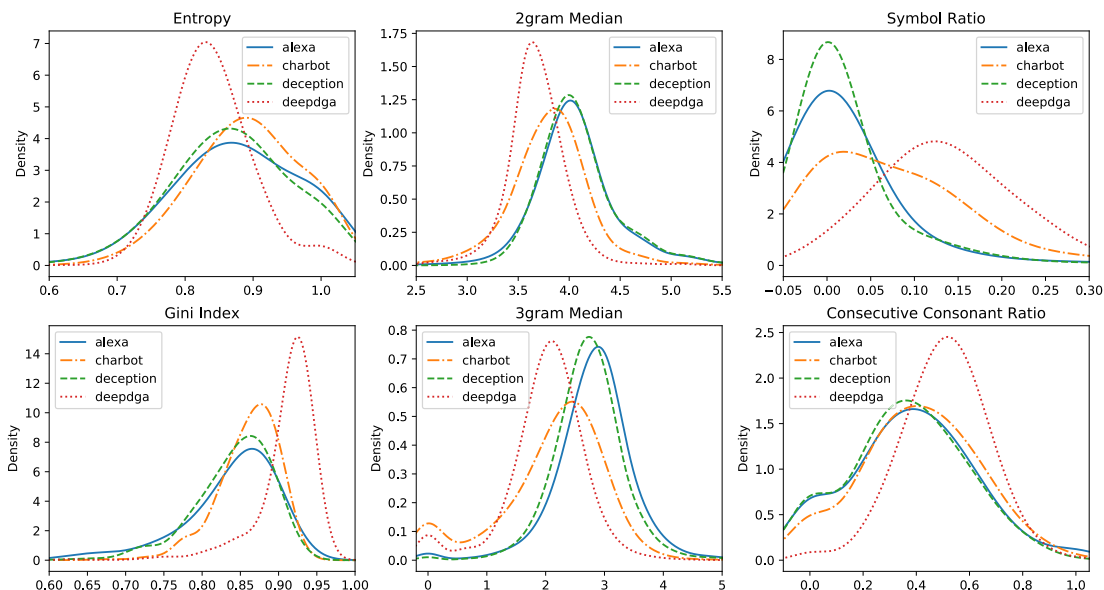


FIGURE 2. Kernel density estimation of the feature distributions of Alexa and adversarial domain names.

where the authors also attempted to replicate the FANCI results.

All models fail to adequately detect CharBot and DeceptionDGA domains even when explicitly trained on them. The LSTM.MI model succeeds in detecting DeepDGA close to 99% of the time with adversarial training, but the other models generally fail at detecting DeepDGA as well. Training on Qname instead of Alexa makes a significant difference, both in predictive performance as well as detection rate: the models have lower predictive accuracy when trained on Qname, but they are better able to detect CharBot domains. At the 0.1% FPR, however, these detection rates are nowhere near high enough to be useful in practice. FANCI is unable to properly detect CharBot at 1% FPR, whereas LSTM.MI and B-RF sometimes manage to obtain over 80% detection rate here. This is not a very useful result, however, since 1% FPR is considered too high to be practical. Therefore, at a low FPR, the domains generated by CharBot can be said to be *transferable* across different models and data sets in the sense that CharBot can fool models that have vastly different architectures and are not trained on Alexa. Combined with its simplicity, speed and small size, this makes CharBot an ideal DGA for use in malware in the wild.

The success of CharBot may be explained as follows. The algorithm works by taking the Alexa list of benign domains — which most DGA classifiers would overwhelmingly classify as such — and introduces a small number of typographical errors. The statistical properties of the domains generated by CharBot are therefore likely almost identical to those of Alexa, causing a low detection rate. The transferability may be explained by noting that even though Alexa and Qname are different data sets, they still capture the same underlying distribution: namely, that of benign domains.

This closeness in distribution is most likely shared among all sufficiently large corpora of benign domains, allowing CharBot to fool any DGA classifier that only takes the domain name string into account. We test this hypothesis by performing kernel density estimation on the feature distributions of the Alexa domains and the adversarial domains. The results are plotted in figure 2. The Entropy and Gini index features are standard impurity measures for decision trees. The other features are:

- 2gram Median. This feature takes the median frequency from the list of 2gram frequencies for the given SLD. Bigram frequencies are collected from the Python package called wordfreq.<sup>14</sup>
- 3gram Median. 3gram median is similar to 2gram median except that it returns the median frequency from the list of trigram frequencies for the given SLD.
- Symbol ratio. This feature defines the ratio of non-alphabetical characters in the SLD, which includes digits and special characters.
- Consecutive Consonant Ratio. This feature defines the ratio of consecutive consonants in the SLD.

From the plots, we observe that the feature distributions of CharBot domains are much closer to those of Alexa than the distributions of DeepDGA are. However, DeceptionDGA is more similar to Alexa than CharBot is, although the difference is very small in some cases. Nevertheless, CharBot gets quite close to Alexa, which explains why it is so successful in fooling DGA classifiers. It also shows that defending against CharBot may be very difficult, potentially requiring a very high FPR. Figure 2 also provides insights into what

<sup>14</sup><https://pypi.org/project/wordfreq/1.1/>. Accessed: 2019-02-14.

parts of CharBot may be improved to yield an even more effective DGA:

- The entropy curve of CharBot can be made more similar to that of Alexa domains by using a different replacement character distribution. Currently, we are using the uniform distribution, which has the highest possible entropy. Switching to a lower entropy distribution may improve the performance of CharBot, although this would need to be carefully balanced against the probability of a generated domain already being registered.
- The random replacement of two characters caused the 2-gram distributions of CharBot to differ from those of the Alexa domains. We can overcome this weakness by replacing neighboring characters with 2-grams that occur frequently in Alexa. A similar line of reasoning applies to 3-grams.
- The symbol ratio distributions can be made more similar by drawing  $c_1$  and  $c_2$  from the same letters or digit sets of the original domain. For example, replace a digit with another random digit and not with a letter.

**TABLE 6. Statistics of the domain name lengths for each data set.**

Data Set	Mean	Standard Deviation
Alexa	14.30	4.70
Bambenek	21.80	6.42
Qname	23.99	7.97
CharBot	14.27	4.02
DeepDGA	28.52	9.06
DeceptionDGA	10.20	3.90

Investigating the lengths of the domain names that were generated vs. those that are present in the Alexa and Qname data sets (see table 6), we find that CharBot names are close in length to Alexa names (which is to be expected), but Qname, Bambenek and DeepDGA domains are significantly longer on average, whereas DeceptionDGA are significantly shorter. This difference in lengths may contribute to the detection rates: when training on Alexa, CharBot domains are similar in length whereas DeepDGA domains are longer like the Bambenek domains. By contrast, when training on Qname, domains are longer on average, which aids detection of CharBot (although the difference is not very large).

## VI. COUNTERMEASURES

We consider a few options for defending against attacks such as CharBot:

### A. COMPARING INCOMING DOMAINS TO ALEXA

The simplest defense against CharBot would be to take the domain in question and compare it to the full Alexa list. If the domain is equal to one found in the Alexa list save for one or two replaced characters, the domain is flagged as malicious. However, the Alexa data set contains one million samples, so this approach of computing the Hamming distance of input domains on the fly may not be practical. We can

make this computation even harder by modifying CharBot to perform deletions and insertions, forcing the use of the edit distance [37] rather than the Hamming one. Practical implementations can reduce lookup time by pre-computing noisy versions of the Alexa list into a compact data structure such as a Bloom filter [39]. However, this approach is marred by a combinatorial explosion of possible corrupted domain names based on the Alexa data set: if we let  $n$  be the size of the CharBot data set,  $\ell$  be the average length of a domain name,  $k$  be the number of edits CharBot introduces and  $m$  be the size of the replacement alphabet, then the number of possible domains CharBot can generate is given approximately by

$$n \binom{\ell}{k} (m-1)^k.$$

For  $n = 10,000$ ,  $\ell = 16$ ,  $m = 40$  and  $k = 2$  this yields 1,825,200,000 possible domains. Besides, this defense can also easily be defeated by simply using a different legitimate data set instead of Alexa for generating domain names.

### B. INCREASING THE CAPACITY OF THE MODELS

Using more complicated classification models may allow them to find a meaningful separation between Alexa and CharBot domains. However, this would require careful feature engineering for featureful models and increase the computational burden of both model training and inference. Given that practical DGA classifiers need to be regularly retrained to keep up with new malware and they need to process many domains in real-time, this may not be feasible. Nevertheless, this may be an option worth exploring in future work.

### C. WHITE-BOX ADVERSARIAL TRAINING

Our adversarial training procedure in this paper has consisted of generating a list of adversarial domains once and then augmenting the training data with them. However, adversarial training is usually done iteratively: at every iteration of training, the current batch of training samples is augmented with adversarially generated set specifically for the model at that particular stage [24], [25]. This requires a *white-box* attack which is able to take the model parameters into account. Adversarial attacks have mostly been considered in the image domain, although there is some work on text classification [40], [41]. Making use of this recent body of work on white-box adversarial training for text classification may allow us to improve the detection rate of CharBot.

### D. USING SIDE INFORMATION

Perhaps the most realistic defense against attacks like CharBot would be to use additional information besides the domain name string alone. For instance, the IP addresses the domain maps to, how often the domain was queried and when, etc. There have been several works investigating the use of such information in DGA classification [3], [10], [16]–[18]. A fruitful avenue for future work could be to test whether these classifiers are more resilient to CharBot.

## VII. CONCLUSION

We have proposed CharBot, a simple and efficient DGA. We have shown CharBot to be effective at both generating large amounts of unregistered domain names as well as fooling three DGA classifiers: FANCI, LSTM.MI and B-RF. We also compared CharBot to DeepDGA and DeceptionDGA, two state-of-the-art domain generation algorithms. The domain names generated by CharBot were more likely to be unregistered than those generated by DeepDGA or DeceptionDGA. Moreover, adversarial retraining using CharBot, DeepDGA or DeceptionDGA did not result in adequate detection of CharBot domains names.

Our DGA is the very first example of a black-box adversarial machine learning attack against DGA classifiers that is not based on Generative Adversarial Networks. We show that simply introducing small perturbations to a set of legitimate domains is good enough and such advanced techniques are unnecessary. We believe this highlights a dangerous weakness of modern DGA classifiers, namely their vulnerability to extremely simple attacks that make no use of sophisticated machine learning techniques. CharBot is an algorithm that could be realistically used in malware in the wild to circumvent state of the art DGA classifiers, making it a real threat. We speculate that this vulnerability is actually *inherent* to any classifier that relies only on the domain name string to perform DGA classification. The CharBot DGA is similar to dictionary DGAs: both have 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 altered 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.

Machine learning models that attempt to do DGA classification based only on the domain name itself, such as the ones considered in this paper, might not be sufficient to detect a DGA like CharBot. The result highlights the need for ML models that exploit additional context features such as the IP-addresses that the domains are mapped to, or temporal access patterns (e.g. how often the domain was requested, and when) [3], [16]–[18], as was done successfully for dictionary DGAs [10].

For future work, we focus on defending DGA classifiers against simple attacks such as CharBot. The avenues we are investigating to achieve this include performing white-box adversarial training as well as augmenting the model inputs with side information that is more difficult to manipulate.

## REPRODUCIBILITY

To foster reproducibility of our results, we are open to sharing all of our code as well as data sets of CharBot samples upon request.

## ACKNOWLEDGEMENT

The authors would like to thank the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. They thank Bobby Filar for making the code of the original DeepDGA algorithm available to them [14] and Jan Spooren for providing them with domain names generated by DeceptionDGA [15].

## REFERENCES

- [1] D. Plohmann, K. Yakdan, M. Klatt, J. Bader, and E. Gerhards-Padilla, "A comprehensive measurement study of domain generating malware," in *Proc. USENIX Security Symp.*, 2016, pp. 263–278. [Online]. Available: [https://www.usenix.org/system/files/conference/usenixsecurity16/sec16\\_paper\\_plohmann.pdf](https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_plohmann.pdf)
- [2] M. Antonakakis, R. Perdisci, Y. Nadji, N. Vasiloglou, II, S. Abu-Nimeh, W. Lee, and D. Dagon, "From throw-away traffic to bots: Detecting the rise of DGA-based malware," in *Proc. USENIX Security Symp.*, vol. 12, 2012, pp. 491–506. [Online]. Available: <https://www.usenix.org/system/files/conference/usenixsecurity12/sec12-final127.pdf>
- [3] 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/abstract/document/6151233>
- [4] 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: <https://arxiv.org/abs/1611.00791>
- [5] 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/pii/S0925231217317320>
- [6] P. Lison and V. Mavroeidis, "Automatic detection of malware-generated domains with recurrent neural models," 2017, *arXiv:1709.07102*. [Online]. Available: <https://arxiv.org/abs/1709.07102>
- [7] 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/document/8215728>
- [8] 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. 6th Int. Symp. Secur. Comput. Commun. (SSCC)*, in Communications in Computer and Information Science. Singapore: Springer, 2018, pp. 640–655. [Online]. Available: [https://link.springer.com/chapter/10.1007/978-981-13-5826-5\\_50](https://link.springer.com/chapter/10.1007/978-981-13-5826-5_50)
- [9] J. J. Koh and B. Rhodes, "Inline detection of domain generation algorithms with context-sensitive word embeddings," in *Proc. IEEE Int. Conf. Big Data*, Dec. 2018, pp. 2966–2971. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8622066>
- [10] 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. 21st Int. Symp. Res. Attacks, Intrusions, Defenses (RAID)*, 2018, pp. 295–314. [Online]. Available: [https://link.springer.com/chapter/10.1007/978-3-030-00470-5\\_14](https://link.springer.com/chapter/10.1007/978-3-030-00470-5_14)
- [11] S. Schüppen, D. Teubert, P. Herrmann, and U. Meyer, "FANCI: Feature-based automated nxdomain classification and intelligence," in *Proc. USENIX Secur. Symp.*, 2018, pp. 1165–1181. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity18/presentation/schuppen>
- [12] R. Sivaguru, C. Choudhary, B. Yu, V. Tymchenko, A. Nascimento, and M. De Cock, "An evaluation of DGA classifiers," in *Proc. IEEE Int. Conf. Big Data*, 2018, pp. 5051–5060. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8621875>
- [13] B. Yu, J. Pan, J. Hu, A. Nascimento, and M. De Cock, "Character level based detection of DGA domain names," in *Proc. IEEE World Congr. Comput. Intell.*, Jul. 2018, pp. 4168–4175. [Online]. Available: <http://faculty.washington.edu/mdcock/papers/byu2018a.pdf>
- [14] 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://dl.acm.org/citation.cfm?id=2996767>

- [15] J. Spooen, D. Preuveneers, 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.*, 2019, pp. 1902–1910. [Online]. Available: <https://irias.kuleuven.be/2361222?lmo=0>
- [16] 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*, 2014, pp. 192–211. [Online]. Available: [https://link.springer.com/chapter/10.1007/978-3-319-08509-8\\_11](https://link.springer.com/chapter/10.1007/978-3-319-08509-8_11)
- [17] 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://www.sciencedirect.com/science/article/abs/pii/S1389128615004843>
- [18] P. Lison and V. Mavroeidis, "Neural reputation models learned from passive DNS data," in *Proc. IEEE Int. Conf. Big Data*, 2017, pp. 3662–3671. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8258361>
- [19] 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>
- [20] Y. Vorobeychik and M. Kantarcioglu, *Adversarial Machine Learning* (Synthesis Lectures on Artificial Intelligence and Machine Learning), vol. 12, no. 3. San Rafael, CA, USA: Morgan & Claypool, 2018, pp. 1–169.
- [21] N. Dalvi, P. Domingos, S. Sanghai, and D. Verma, "Adversarial classification," in *Proc. 10th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2004, pp. 99–108. [Online]. Available: <http://s33043.gridserver.com/files/teaching/stanford/2008/readings/AdversarialClassification%20DalviEtAl%20KDD04.pdf>
- [22] D. Lowd and C. Meek, "Adversarial learning," in *Proc. 11th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2005, pp. 641–647. [Online]. Available: [https://www.researchgate.net/profile/Daniel\\_Lowd/publication/221654486\\_Adversarial\\_learning/links/00b7d522999ea43eb4000000.pdf](https://www.researchgate.net/profile/Daniel_Lowd/publication/221654486_Adversarial_learning/links/00b7d522999ea43eb4000000.pdf)
- [23] D. Lowd and C. Meek, "Good word attacks on statistical spam filters," in *Proc. CEAS*, 2005. [Online]. Available: [http://www.utdallas.edu/~muratk/courses/dmsec\\_files/125.pdf](http://www.utdallas.edu/~muratk/courses/dmsec_files/125.pdf)
- [24] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," 2014, *arXiv:1412.6572*. [Online]. Available: <https://arxiv.org/abs/1412.6572>
- [25] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," 2017, *arXiv:1706.06083*. [Online]. Available: <https://arxiv.org/abs/1706.06083>
- [26] A. Raghunathan, J. Steinhardt, and P. S. Liang, "Semidefinite relaxations for certifying robustness to adversarial examples," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 10877–10887. [Online]. Available: <https://papers.nips.cc/paper/8285-semidefinite-relaxations-for-certifying-robustness-to-adversarial-examples.pdf>
- [27] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680. [Online]. Available: <http://papers.nips.cc/paper/5423-generative-adversarial-nets>
- [28] B. Hitaj, P. Gasti, G. Ateniese, and F. Perez-Cruz, "PassGAN: A deep learning approach for password guessing," 2017, *arXiv:1709.00440*. [Online]. Available: <https://arxiv.org/abs/1709.00440>
- [29] W. Hu and Y. Tan, "Generating adversarial malware examples for black-box attacks based on GAN," 2017, *arXiv:1702.05983*. [Online]. Available: <https://arxiv.org/abs/1702.05983>
- [30] A. C. Bahnsen, I. Torroledo, L. D. Camacho, and S. Villegas, "DeepPhish: Simulating malicious AI," in *Proc. IEEE APWG Symp. Electron. Crime Res.*, 2018, pp. 15–17. [Online]. Available: [https://albahnsen.com/wp-content/uploads/2018/05/deepphish-simulating-malicious-ai\\_submitted.pdf](https://albahnsen.com/wp-content/uploads/2018/05/deepphish-simulating-malicious-ai_submitted.pdf)
- [31] A. C. Bahnsen, E. C. Bohorquez, S. Villegas, J. Vargas, and F. A. González, "Classifying phishing URLs using recurrent neural networks," in *Proc. IEEE APWG Symp. Electron. Crime Res.*, Apr. 2017, pp. 1–8. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/7945048>
- [32] A. Anand, K. Gorde, J. R. A. Moniz, N. Park, T. Chakraborty, and B.-T. Chu, "Phishing URL detection with oversampling based on text generative adversarial networks," in *Proc. IEEE Int. Conf. Big Data*, Dec. 2018, pp. 1167–1176. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8622547>
- [33] J. Burns and E. Heath, "Using generative adversarial networks to harden phishing classifiers," in *Proc. FloCon*, New Orleans, LA, USA, 2019. [Online]. Available: <https://flocon2019.sched.com/event/GXW1/using-generative-adversarial-networks-to-harden-phishing-class>
- [34] J. Szurdi, B. Kocso, G. Cseh, J. Spring, M. Felegyhazi, and C. Kanich, "The long 'taile' of typosquatting domain names," in *Proc. 23rd USENIX Secur. Symp. (USENIX Security)*, 2014, pp. 191–206. [Online]. Available: <https://www.usenix.org/system/files/conference/usenixsecurity14/sec14-paper-szurdi.pdf>
- [35] 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: <https://arxiv.org/abs/1902.08909>
- [36] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, "DeepFool: A simple and accurate method to fool deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2574–2582. [Online]. Available: [https://www.cv-foundation.org/openaccess/content\\_cvpr\\_2016/papers/Moosavi-Dezfooli\\_DeepFool\\_A\\_Simple\\_CVPR\\_2016\\_paper.pdf](https://www.cv-foundation.org/openaccess/content_cvpr_2016/papers/Moosavi-Dezfooli_DeepFool_A_Simple_CVPR_2016_paper.pdf)
- [37] V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals," *Sov. Phys.-Dokl.*, vol. 10, no. 8, pp. 707–710, 1966. [Online]. Available: <https://nymity.ch/sybilhunting/pdf/Levenshtein1966a.pdf>
- [38] P. Royal, "Quantifying maliciousness in alexa top-ranked domains," in *Proc. BlackHat*, Seattle, WA, USA, 2012. [Online]. Available: <https://media.blackhat.com/ad-12/Royal/bh-ad-12-quantifying-royal-slide.pdf>
- [39] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Commun. ACM*, vol. 13, no. 7, pp. 422–426, Jul. 1970. [Online]. Available: <http://crystal.uta.edu/~mcguigan/cse6350/papers/Bloom.pdf>
- [40] Z. Gong, W. Wang, B. Li, D. Song, and W.-S. Ku, "Adversarial texts with gradient methods," 2018, *arXiv:1801.07175*. [Online]. Available: <https://arxiv.org/abs/1801.07175>
- [41] J. Ebrahimi, A. Rao, D. Lowd, and D. Dou, "HotFlip: White-box adversarial examples for text classification," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics*, vol. 2, 2018, pp. 31–36. [Online]. Available: <http://www.aclweb.org/anthology/P18-2006>



robustness of machine learning models to adversarial manipulations.


**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 Department of Applied Mathematics, Computer Science and Statistics and the VIB Inflammation Research Center, Ghent, Belgium. His research is sponsored by a fellowship of the Research Foundation Flanders (FWO) and focuses improving the



**CLAIRE NIE** received the B.Eng. degree in chemical engineering from McGill University, Canada, in 2012. She is currently pursuing the M.S. degree in computer science and systems with the University of Washington Tacoma, USA. She was a Financial Auditor of Technology Companies, Silicon Valley. During her graduate studies, she carried out research on training random forest classifiers for the detection of algorithmically generated domain names.



**RAAGHAVI SIVAGURU** received the B.Tech. degree in information technology from Anna University, India, in 2014. She is currently pursuing the M.S. degree in computer science and systems with the University of Washington Tacoma, USA. Her research interests include improving the predictive performance of DGA detection classifiers and conducting an evaluation of hardening techniques that use side information to defend against adversaries.



**CHARLES GRUMER** received the B.A. degree in psychology from the University of Washington Seattle, in 2015. He is currently pursuing the M.S. degree in computer science and systems with a specialty in data science, the University of Washington Tacoma. He is slated to graduate at the end of 2019, pending the completion and successful defense of his thesis, which focuses on white box adversarial example generation for DGA classifiers. He is also a Data Science Intern with Infoblox.



**FEMI OLUMOFIN** received the Ph.D. degree in computer science from the University of Waterloo, Canada. He is currently a Senior Member of the data science and analytics team with Infoblox in the San Francisco Bay Area. He has made contributions to research and development in the areas of privacy enhancing technologies, security, applied cryptography, big data analytics, and machine learning.



**BIN YU** received the Ph.D. degree in electronic engineering from Tsinghua University, China. He was a Postdoctoral Fellow with the Pattern Recognition and Image Processing Lab, Michigan State University, USA, and an Associate Professor with Beijing Jiaotong University, China. He was with many high tech companies in Silicon Valley at a senior leadership positions and a led projects of machine learning and artificial intelligence for internet search, medical imaging, computer vision,

and e-commerce. He is currently a Chief Data Scientist, pioneered big data analytics with Infoblox, Santa Clara, USA, to detect malicious DNS traffic, using deep learning, and artificial intelligence techniques to keep pace with fast changing malware evolution. He has a rich experience in both academia and industry for more than 25 years and has published more than 50 peer-reviewed papers and patents in artificial intelligence, deep learning, machine learning, image processing, and cybersecurity. He served as a Senior Member for the IEEE Computer Society.



**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, and a Faculty Member with the University of Brasilia, Brazil. He is currently the 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 in these areas.

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 in these areas.



**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, 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 150 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 also a Program Committee Member of numerous international conferences. She co-organized the KDDCup2013. She has served as an Associate Editor for the IEEE TRANSACTIONS ON FUZZY SYSTEMS.

• • •