On the Use of DGAs in Malware: An Everlasting Competition of Detection and Evasion

Jan Spooren
imec-Distrinet, KU Leuven
Celestijnenlaan 200A
3001 Heverlee, Belgium
jan.spooren@cs.kuleuven.be

Davy Preuveneers
imec-Distrinet, KU Leuven
Celestijnenlaan 200A
3001 Heverlee, Belgium
davy.preuveneers@cs.kuleuven.be

Lieven Desmet
imec-Distrinet, KU Leuven
Celestijnenlaan 200A
3001 Heverlee, Belgium
lieven.desmet@cs.kuleuven.be

Peter Janssen
EURid VZW, Belgium
Woluwelaan 150
1831 Diegem, Belgium
peter.janssen@eurid.eu

Wouter Joosen
imec-Distrinet, KU Leuven
Celestijnenlaan 200A
3001 Heverlee, Belgium
wouter.joosen@cs.kuleuven.be

ABSTRACT

Malware typically makes use of *Domain Generation Algorithms* (DGAs) as a mechanism to contact their *Command and Control* server. In recent years, different approaches to automatically detect generated domain names have been proposed, based on machine learning. The first problem that we address is the difficulty to systematically compare these DGA detection algorithms due to the lack of an independent benchmark. The second problem that we investigate is the difficulty for an adversary to circumvent these classifiers when the machine learning models backing these DGA-detectors are known. In this paper we compare two different approaches on the same set of DGAs: classical machine learning using manually engineered features and a ‘deep learning’ recurrent neural network. We show that the deep learning approach performs consistently better on all of the tested DGAs, with an average classification accuracy of 98.7% versus 93.8% for the manually engineered features. We demonstrate that the deep learning solution yields better results even when only 10,000 malicious samples are available. We also show that one of the dangers of manual feature engineering is that DGAs can adapt their strategy, based on knowledge of the features used to detect them. To demonstrate this, we use the knowledge of the used feature set to design a new DGA which makes the Random Forest classifier powerless with a classification accuracy of 57.3%. The deep learning classifier is also (albeit less) affected, reducing its accuracy to 78.9%.

CCS Concepts

- Security and privacy → Malware and its mitigation;
- Computing methodologies → Neural networks; Classification and regression trees;

Keywords

Malware Detection, Domain Generation Algorithms, Machine Learning

1. INTRODUCTION

The Internet connects billions of devices, from servers and personal computers to tablets, mobile phones, household appliances, and many more. Malicious actors are constantly scanning the internet for vulnerable devices which could be compromised, or are tricking users into unknowingly installing malware on their devices.

Figure 1: Bot using DGA to connect to a C&C Server

Once this malware is present on a machine, it can be used to attack other machines, send unsolicited or phishing e-mails, eavesdrop on communication, steal e-mail addresses, encrypt the contents of the machine requesting from the

user a ransom for the ability to decrypt, and many more malicious schemes. Large pools [18] of infected machines, called botnets [4] exist, which are controlled from *Command and Control* (C&C) servers (as depicted in Figure 1).

To prevent those C&C servers from being shut down or made unreachable, malware typically uses *Domain Generation Algorithms* (DGAs) [16] in order to create a new set of pseudo-random domain names every certain period of time. Many random domains are generated and then tried by the malware, only one of which the botnet owner must register in order for the botnet to be able to successfully reconnect to its C&C server. This makes taking down botnets a difficult task.

Conversely, recognizing algorithmically generated domain names can help in detecting infected hosts, as well as flag domain name registrations made with the intent to control a botnet. Machine learning classifiers that can successfully distinguish an algorithmically generated domain name from a human-created domain name are therefore useful tools to security researchers, law enforcement and network operations staff.

In the domain of security and privacy, deep neural networks have demonstrated their ability to autonomously find and extract relevant features as well as improved classification accuracies when compared to traditional machine learning methods [11, 17, 5]. Also for the detection of DGAs, there is recent related work proposing solutions based on both traditional machine learning and deep learning methods. However, malicious adversaries can exploit these AI classification methods too to evade detection of their malware. These observations reflect the dual arms race. The first race relates to further improving advanced deep learning methods to extract better features for offensive or defensive purposes. The second one is the cat and mouse game of adversaries trying to circumvent new defensive measures developed by security specialists, irrespective of the type of machine learning methods being used.

The first challenge that we address in this work is that existing detection methods are difficult to compare. Not only do they rely on different data sets and ground truths to assess their classification accuracy, assumptions made by some of these related works regarding the ground truth of malicious versus benign domain names may not always be valid. Both concerns make it difficult to interpret and compare the accuracy of DGA detection methods. The second challenge that we focus on in this work is that adversaries become smarter too in circumventing these DGA detection methods.

The main contributions of this work are as follows:

1. We independently benchmark and compare two state-of-the-art DGA detection methods: (a) FANCI [20], a state-of-the-art Random Forest based classifier recently presented at USENIX Security 2018 using manually crafted features, and (b) Woodbridge’s LSTM [27], a deep recurrent neural network based classifier. We compare the two approaches, both as binary classifiers for detecting a particular DGA and as multiclass classifiers, predicting the DGA family a domain name was generated with.

2. We investigate how the quantity of available training data influences the classification accuracy for the 5 best and the 5 worst detectable DGAs which were evaluated.
3. We illustrate the arms race using FANCI as a detection method and designing a new DGA that exploits knowledge about FANCI’s feature extraction to circumvent detection. We also evaluate how well our new DGA is able to fool Woodbridge’s LSTM for which the inherent features that characterize a DGA are more implicit.

In section 2, we briefly summarize relevant related work, elaborating on an approach using a Random Forest with manually crafted features. In section 3, we explain the used methodology, the collection of the ground truth data (section 3.1) and the compared classifiers (section 3.2). We present the results of this comparison in section 4. In section 5, we show that knowledge of the used featureset can be abused by DGAs to circumvent detection, by building a DGA which achieves exactly this. Finally, in section 7, we state our main conclusions.

2. RELATED WORK

Many different approaches were taken to detect DGA - generated domain names: Yadav et al. [28] used DNS queries and responses for detecting malware. They used the entropy in the observed unigram and bigram statistics on the character sequences present in domain names to distinguish (often difficult to pronounce) random character strings from human-created (and usually better pronounceable) domain names. Antonakakis et al. [3] use a clustering approach on the length, level of randomness and character frequency distribution, including the n -gram distribution (with $n = 1, \dots, 4$) of observed domain names that have been queried by overlapping sets of hosts. In a second phase, a *Hidden Markov Model* (HMM) is used for determining the likelihood that a domain name is a C&C domain generated by a particular DGA.

Mowbray and Hagen [12] described a procedure to detect even previously unseen DGAs from DNS query data, by spotting an unusual number of DNS requests for an unusual distribution in domain name length. While their method proved successful, it is doubtful if using domain name length only as a feature for detection will suffice when DGAs become more advanced.

Schiavoni et al. [19] created a system called Phoenix for detecting DGA-based botnets, which uses filtering using linguistic features (detecting meaningful subwords, n -gram normality scores and statistical linguistic features for the English language), followed by a clustering step and finally, a fingerprinting step to extract invariant properties of a DGA and subsequently label the collected domain names.

In an excellent paper by Woodbridge et al. [27], *Long Short-Term Memory* (LSTM) networks [8] are proposed to detect algorithmically-generated domain names. Woodbridge and his team at EndGame, Inc. show that LSTMs are far superior to Random Forests with manually engineered features, logistic regression with bigram features and HMM classifiers.

Pereira et al. [14] proposed a novel *WordGraph* method for detection of dictionary-based DGAs using graph-theory. They showed that random forests using lexical features perform badly on dictionary-based DGAs, likely due to the good resemblance of the features to manually chosen domain names. Their method outperforms convolutional neural nets for dictionary-based DGAs, which they have found to work well too on these DGAs, presumably due to the neural net’s ability to learn interesting soft n -grams, which, in a sense, allows them to memorize dictionaries, too.

Recently, a paper by Schüppen et al. [20] presented FANCI at USENIX Security 2018, carrying no reference to any work using deep learning, in which manually engineered features were used, in combination with a Random Forest to detect DGA-generated domains in NXDomain-failed DNS queries (i.e., queries to non-existent domain names).

In this work, we have re-implemented and re-evaluated two state-of-the-art DGA detection methods, the first using deep learning, the other using traditional machine learning, i.e. respectively the LSTM method presented in [27] and the FANCI Random Forest classifier presented in [20], and compared the performance of both approaches on the same data set. Such a systematic comparison is necessary as (1) different works rely on dissimilar methods to obtain the ground truth used to train the machine learning classifiers, and (2) previous research has shown that some works make invalid assumptions about the truthfulness of their ground truth data. Both solutions will be tested against the newly designed DGA.

In an earlier version of this paper [22], presented at the 34th ACM/SIGAPP Symposium for Applied Computing, we first presented *deception_dga*. Our DGA has meanwhile been evaluated in a paper by Jonathan Peck et al. [13] where also a new DGA named *CharBot* is proposed. *CharBot* is based on a dictionary of 10,000 valid domain names, of which two characters are replaced by random characters. The paper also evaluates the new DGA as well as *deception_dga* on FANCI and an LSTM-based classifier. Another very recent new academic DGA is *MaskDGA* [21], which starts with regular DGA domains and adds small perturbations in order to thwart DGA detection classifiers.

3. METHODOLOGY

To compare the classification performance of two classifiers, we first need to obtain a ‘ground truth’ data set: in this case, a collection of domain names which are known with certainty to be algorithmically generated or not.

3.1 Ground truth

Ground truth data for both *malicious* (DGA) and *benign* (non-DGA) domain names is by the very nature of this research harvested from different sources:

Malicious ground truth was obtained from DGArchive [15], a service offered by Fraunhofer FKIE and maintained by Daniel Plohmman. Since only a limited number of top level domains (TLD) are available and the particular TLD used by a domain generation algorithm does not have any influence on classification accuracy, we removed the TLD

from the DGA domain names. We selected all DGAs in DGArchive for which 100,000 or more unique recorded domain names were available when disregarding the TLD from the domain name. This resulted in a list of 26 DGAs shown in Table 1.

Table 1: Domain Generation Algorithms Evaluated

DGA	Length		Comments
banjori_dga	7 → 26	†	banking trojan
chinad_dga	16		
conficker_dga	5 → 11	†	worm
cryptolocker_dga	12 → 18	†	ransomware
dnschanger_dga	10		
dyre_dga	34		banking malware
emotet_dga	16	†	information stealer
gameover_dga	19 → 28		banking info stealer
gameover_p2p	11 → 32	†	botnet Zeus
gozi_dga	12 → 24	†	information stealer
locky_dga	5 → 17		ransomware
murofet_dga	7 → 16	†	
murofetweekly_dga	32 → 47		
necurs_dga	7 → 25		
nymaim_dga	5 → 13		ransomware
padcrypt_dga	16 → 19		ransomware
proslifean_dga	6 → 12		
pushdo_dga	8 → 12		
pykspa_dga	6 → 12	†	
qadars_dga	12		
qakbot_dga	8 → 25		
ramnit_dga	8 → 19		file infector
ranbyus_dga	14 → 17		
rovnix_dga	18		
symmi_dga	13 → 20		
virut_dga	6		file infector

†: Malware not active anymore in 2018.

For **benign ground truth**, many authors choose the Alexa [1] top- n lists of most accessed web sites. However, recent work by Le Pochat et al. [10] has shown the Alexa top- n lists to be easily manipulatable. In fact, it turns out the Alexa lists already contain DGA-generated domain names. Moreover, the list of most popular domain names is very likely biased towards shorter, easier to pronounce and to remember domain names, compared to the average registered domain name. Therefore, the authors of this work have chosen to use the list of the first 100,000 registered domain names from a well known TLD in 2016, which were filtered according to the following criteria:

- The domain names did not end up on any of the following black lists: *Google’s Safe Browsing List* [7], *Spamhaus DBL blacklist* [24], the *SURBL blacklist* [23].
- No internationalized domain names were used (since at the time of writing none of the DGAs use internationalized (IDN) domain names – which would make classification of internationalized domain names trivial)
- The domain names were not known in DGArchive.

Figure 2: Feed-forward vs. recurrent neural networks

For each DGA, we constructed the ground truth data as a shuffled mix of 100,000 benign and 100,000 malicious domain names, using the same benign ground truth data for each DGA evaluated.

3.2 DGA Classifiers

In this section, we briefly describe the main characteristics of the two state-of-the-art DGA classifiers we used for our experiments.

3.2.1 The Woodbridge LSTM

We implemented the Long Short-Term Memory (LSTM) neural network as proposed in the aforementioned work by Woodbridge et al. [27].

An LSTM is a specific type of recurrent neural network (RNN). RNNs are frequently used for recognizing or predicting patterns in sequential data. Contrary to feed forward neural networks, RNNs have an internal short-term memory to persist important things about the input they receive. They achieve this by copying the output and looping it back into the network, as depicted in Figure 2. This characteristic enables RNNs to create a deeper understanding of a sequence and its context, and to predict what is coming next. LSTMs extend RNNs by enabling such networks to remember their inputs over a longer period of time, thereby extending their memory capacity beyond two time steps. A cell in an LSTM has a state that can be read, written or reset via a set of programmable gates. These gates modulate both the input connection and recurrent connection with a value between 0 and 1 allowing the current state to remain the same between time steps, to be forgotten in the next time step or a combination thereof.

LSTMs are a good fit to recognize DGAs as they can learn and generalize the pattern generation process of many DGAs without the need to manually craft higher level features based on the raw input. The authors also argue that LSTMs work as a black box making it difficult for adversaries to circumvent a classifier without the same training set. The DGA classifier proposed by Woodbridge et al. consists of following sequential layers:

- An **embedding layer**, which converts variable length sequences of domain name characters to a fixed length, zero-padded array of features.
- The **LSTM layer**, which receives its input from the embedding layer with a dimensionality of 38 (encoding 26 characters, 10 digits, the dash character and an end token), and generates an output dimension of 128.
- A **dropout layer** of 0.5, preventing over-fitting.
- A dense **output layer**, with one output dimension followed by a sigmoid activation function for the binary classifier or an output layer with the number of dimensions corresponding with the number of predicted classes (the number of DGAs to be predicted, plus the benign class) followed by a *softmax* activation function for the multiclass classifier.

To ensure the validity of our classification results, we re-used the implementation available on the EndGame Inc. GitHub repository [9].

We evaluate the performance of the classifier for each DGA separately, using *5-fold cross validation*: 4/5 of the data is used to train the network in 10 epochs, using a batch size of 128; the remaining 1/5 of the data is then used for testing the trained network. This is repeated 4 more times with different folds of the data, discarding the previously trained network, thus ensuring testing data was never used for training.

3.2.2 Random Forest and FANCI features

We implemented the FANCI feature extraction as described in the aforementioned work by Schüppen et al. [20] and fed the 41 extracted features¹ to a Random Forest of 100 trees, each constructed considering 6 random features. We again used 5-fold cross validation to evaluate its performance.

In this comparative work, we wish to compare the efficacy of the deep-learning LSTM network, which will find its own features from the input data, versus the manually engineered features from the FANCI system in detecting algorithmically generated domain names, purely based on the domain name character sequences. The FANCI system not only detects DGA generated domain names by examining domain name character sequences, but also looks at other features, harvested from NXDomain (non-existing domain) DNS queries. E.g., FANCI feature 5, *Has Valid TLD* is a binary feature, indicating that the domain name has a valid top level domain. It is clear that no DGA will ever output domain names with an invalid TLD, since those domain names would never

¹FANCI uses 21 features, however, feature #20 is a feature vector of 21 values, thereby resulting in 41 values

Table 2: FANCI features that are not used in the comparison

No	Feature	Output
2	Number of Subdomains	int
3	Subdomain Length Mean	real
4	Has www Prefix	bool
5	Has Valid TLD	bool
6	Contains Single-Character Subdomain	bool
8	Contains TLD as Subdomain	bool
9	Ratio of Digit-Exclusive Subdomains	real
10	Ratio of Hex-Exclusive Subdomains	real
11	Underscore Ratio	real

resolve and are therefore useless. This characteristic is a useful feature in the FANCI system, since it can easily detect human typing errors in the TLD-part of domain names, which are therefore clearly never DGA-generated domain names. Our methodology to fairly compare the two classifiers on their ability to detect an algorithmically generated domain name, consists of feeding both with a mix of 50% generated domain names by a known DGA and 50% benign ground truth data, both of which carry no TLD nor subdomains. Therefore, a number of features used by the FANCI system will in practice not be used in our comparative test. These features are listed in Table 2. It is therefore to be expected that the classification scores for FANCI as listed in this work will be lower than those published in the FANCI paper [20].

4. DGA CLASSIFICATION RESULTS

In this section we will discuss the results of the classifiers introduced in section 3.2, first implemented as a set of 26 binary classifiers, one for each DGA (section 4.1), then as multiclass classifiers, which will make a prediction in a single classifier if the domain name was generated, and if so, by which DGA (section 4.2).

4.1 Binary classification

The classification results for each of the tested DGAs are listed for both classifiers in Table 3. For all of the tested DGAs, the LSTM-classifier yielded a better or identical *True Positive Rate* (TPR), *False Positive Rate* (FPR), *Precision* and *Accuracy*. On average, the FPR of the Random Forest classifier is over 4 times higher than the LSTM. The average accuracy of the LSTM-classifier is 98.7%, while the average Random Forest classifier reaches 93.7%.

The standard deviation on the FPR is 3 times higher for the Random Forest classifier, and the standard deviation for the accuracy is 4 times higher for the Random Forest classifier, indicating that the LSTM not only yields better classification results, but also a higher consistency over different DGAs. Notably, the Random Forest with the FANCI features perform considerably poorer on `gozi_dga`, `locky_dga`, `nymaim_dga`, `pushdo_dga` and `pykspa_dga`.

One possible explanation for the better consistency of the LSTM classifier is that the LSTM network learns its used classification features automatically during the training process, whereas the Random Forest needs manually engineered

features, which may be better or worse suited for different DGAs.

4.2 Multiclass classification

For training the two presented classifiers as multiclass classifiers, we create a ground truth set, consisting of 10,000 samples for each DGA and 10,000 benign samples, resulting in 270,000 training instances. We trained the LSTM using 30 epochs and used a Random Forest, bagging with 100 iterations. We again used stratified 5-fold cross validation for the evaluation of both classifiers.

The overall classification accuracy was 72.8% for the LSTM, versus 63.8% for the Random Forest. The per-DGA classification results are presented in Table 4, and heatmap representations of the confusion matrices for both classifiers are shown in Figure 3.

Some notable observations:

- Although the overall classification accuracy is significantly higher for the LSTM, the detection of the benign class is significantly poorer for the LSTM than for the Random Forest. We attribute this to the fact that the Random Forest feature set was specifically crafted to detect DGA samples out of benign samples, whereas the LSTM needs to learn its feature set in a very imbalanced dataset, consisting of only 3,7% benign samples, where there is typically more similarity between different DGA variants than between the benign samples. We expect this performance to improve when training with more benign samples. This hypothesis is confirmed by increasing the benign ground truth to 50,000 samples (16.1% of the total number of samples), which indeed increases the TPR of the benign class to 90.9 % and also slightly increases the overall classification accuracy to 74.7%.
- Both classifiers are having difficulty in correctly classifying a number of DGAs: `conficker_dga`, `locky_dga`, `nekurs_dga`, `nymaim_dga`, `qakbot_dga` and `ramnit_dga` have very low True Positive Rates. When examining the confusion matrices in Figure 3, it shows that e.g. `locky_dga`(11) is frequently being misclassified as `cryptolocker_dga` (4), `proslkefan_dga` (17) and `ranbyus_dga` (23). These DGAs all use a nearly identical character frequency distribution, as is illustrated for `locky_dga` and `cryptolocker_dga` in Figure 4, indicating that they are closely related, making misclassification between these DGAs likely.

In section 4.1, we mentioned that a weakness of manually engineered features is the fact that these features may be better or worse performing, depending on the DGA they are trying to classify. There is, however, another danger in the use of manually engineered features: knowledge of which features are tracked can be used to the advantage of DGA developers, who can modify the Domain Generation Algorithm to become less easily detectable.

In the following section, we will do exactly this: we will construct a new DGA, with the knowledge of the FANCI feature set to make it less easily detectable.

Table 3: Comparison of LSTM and Random Forest performance on individual DGAs. The Min, Average, Max and StdDev exclude our specially crafted *deception_dga* and *deception2_dga*. Better scores are indicated in boldface.

DGA	LSTM				Random Forest			
	TPR	FPR	Prec	Acc	TPR	FPR	Prec	Acc
banjori_dga	1.000	0.000	1.000	1.000	0.950	0.050	0.951	0.950
chinad_dga	0.999	0.001	0.999	0.999	0.998	0.002	0.998	0.998
conficker_dga	0.967	0.048	0.952	0.959	0.871	0.129	0.871	0.871
cryptolocker_dga	0.992	0.004	0.996	0.994	0.936	0.064	0.937	0.936
dnschanger_dga	0.994	0.012	0.988	0.991	0.964	0.036	0.965	0.964
dyre_dga	1.000	0.000	1.000	1.000	1.000	0.000	1.000	1.000
emotet_dga	0.997	0.002	0.998	0.998	0.986	0.014	0.986	0.986
gameover_dga	1.000	0.000	1.000	1.000	1.000	0.000	1.000	1.000
gameover_p2p	1.000	0.000	1.000	1.000	0.989	0.011	0.989	0.989
gozi_dga	0.957	0.044	0.956	0.957	0.857	0.143	0.859	0.857
locky_dga	0.976	0.028	0.973	0.974	0.862	0.138	0.863	0.862
murofet_dga	0.997	0.002	0.998	0.997	0.960	0.040	0.960	0.960
murofetweekly_dga	1.000	0.000	1.000	1.000	1.000	0.000	1.000	1.000
necurs_dga	0.984	0.021	0.979	0.981	0.879	0.121	0.880	0.879
nymaim_dga	0.968	0.049	0.952	0.960	0.864	0.136	0.864	0.864
padcrypt_dga	1.000	0.000	1.000	1.000	0.992	0.008	0.992	0.992
proslikefan_dga	0.968	0.038	0.963	0.965	0.880	0.120	0.880	0.880
pushdo_dga	0.997	0.013	0.987	0.992	0.865	0.135	0.875	0.865
pykspa_dga	0.972	0.035	0.966	0.969	0.857	0.143	0.857	0.857
qadars_dga	1.000	0.000	1.000	1.000	0.994	0.006	0.994	0.994
qakbot_dga	0.987	0.016	0.984	0.985	0.905	0.095	0.905	0.905
ramnit_dga	0.986	0.016	0.984	0.985	0.883	0.117	0.883	0.883
ranbyus_dga	0.996	0.002	0.998	0.997	0.973	0.027	0.973	0.973
rovnix_dga	1.000	0.001	0.999	0.999	0.998	0.002	0.998	0.998
symmi_dga	0.989	0.020	0.980	0.984	0.936	0.064	0.936	0.936
virtut_dga	0.991	0.033	0.968	0.979	0.975	0.025	0.976	0.975
<i>deception_dga</i>	<i>0.871</i>	<i>0.160</i>	<i>0.844</i>	<i>0.855</i>	<i>0.599</i>	<i>0.401</i>	<i>0.599</i>	<i>0.599</i>
<i>deception2_dga</i>	<i>0.815</i>	<i>0.237</i>	<i>0.775</i>	<i>0.789</i>	<i>0.573</i>	<i>0.427</i>	<i>0.576</i>	<i>0.573</i>
Min	0.957	0.000	0.952	0.957	0.857	0.000	0.857	0.857
Average	0.989	0.015	0.985	0.987	0.937	0.063	0.938	0.937
Max	1.000	0.049	1.000	1.000	1.000	0.143	1.000	1.000
StdDev	0.013	0.017	0.016	0.014	0.055	0.055	0.054	0.055

(a) LSTM

(b) Random Forest

Figure 3: Confusion matrix heat maps for the multiclass LSTM and RF classifiers

Table 4: Comparison of LSTM and Random Forest performance when performing multi-class classification on the same set of DGAs. Better scores are indicated in boldface.

DGA	No	LSTM			Random Forest		
		TPR	FPR	Prec	TPR	FPR	Prec
benign	0	0.762	0.008	0.785	0.978	0.000	0.987
banjori_dga	1	0.998	0.000	0.996	0.830	0.015	0.674
chinad_dga	2	0.994	0.000	0.999	0.990	0.000	0.994
conficker_dga	3	0.187	0.015	0.318	0.247	0.023	0.289
cryptolocker_dga	4	0.730	0.043	0.397	0.465	0.039	0.314
dnschanger_dga	5	0.984	0.029	0.562	0.915	0.029	0.548
dyre_dga	6	1.000	0.000	1.000	1.000	0.000	1.000
emotet_dga	7	0.947	0.014	0.723	0.819	0.023	0.576
gameover_dga	8	1.000	0.000	1.000	1.000	0.000	0.998
gameover_p2p	9	0.942	0.003	0.923	0.808	0.009	0.784
gozi_dga	10	0.920	0.004	0.887	0.677	0.011	0.712
locky_dga	11	0.030	0.006	0.159	0.027	0.009	0.101
murofet_dga	12	0.744	0.010	0.732	0.297	0.021	0.356
murofetweekly_dga	13	0.999	0.000	1.000	1.000	0.000	1.000
necurs_dga	14	0.163	0.006	0.531	0.100	0.011	0.250
nymaim_dga	15	0.082	0.007	0.303	0.065	0.011	0.191
padcrypt_dga	16	0.999	0.000	0.999	0.923	0.005	0.867
proslkefan_dga	17	0.524	0.053	0.276	0.401	0.043	0.263
pushdo_dga	18	0.970	0.003	0.934	0.765	0.020	0.596
pykspa_dga	19	0.373	0.014	0.503	0.134	0.020	0.208
qadars_dga	20	0.985	0.000	0.995	0.976	0.001	0.983
qakbot_dga	21	0.372	0.004	0.786	0.131	0.014	0.268
ramnit_dga	22	0.145	0.012	0.325	0.083	0.012	0.216
ranbyus_dga	23	0.887	0.026	0.568	0.749	0.027	0.520
rovnix_dga	24	0.990	0.000	0.999	0.986	0.000	0.990
symmi_dga	25	0.963	0.002	0.939	0.866	0.011	0.746
virut_dga	26	0.979	0.022	0.635	0.997	0.021	0.642
Min		0.030	0.000	0.159	0.027	0.000	0.101
Average		0.728	0.010	0.714	0.638	0.014	0.595
Max		1.000	0.053	1.000	1.000	0.043	1.000
StdDev		0.347	0.014	0.280	0.367	0.012	0.312

(a) locky_dga

(b) cryptolocker_dga

Figure 4: Similar character distributions for locky_dga and cryptolocker_dga

5. ATTACKING THE MANUALLY CRAFTED FEATURESET

The goal of this exercise is to demonstrate that it is possible to create a new DGA, which we will name `deception_dga`, which takes the features of a classifier into account in order to circumvent detection. An adversary creating such a DGA will need to implement the feature set used by the classifier, obtain a classifier implementation, and obtain the benign ground truth.

We used the following approach:

- We implemented the FANCI feature set of 41 features as defined in the FANCI paper [20] in a Python implementation. However, as mentioned in section 3.2.2, the features listed in Table 2 are unused in our comparison, the goal of which was to compare two different machine learning techniques objectively on their ability to detect algorithmically generated domain names, purely based on domain name character sequences. Therefore, in the following, we will examine the effect of influencing a certain number of the 32 features actually used.
- We used the WEKA[26] workbench and its implementation of the Random Forest to test the performance of our newly built DGA.
- An adversary creating a new DGA will require benign ground truth data for testing too: we assume an adversary would probably use the publicly available Alexa [1] list. We used the top 1 million sites of September 16th 2018, and used 100,000 unique domain names (excluding the TLD), starting with domain 100,001 (`senenews`) and ending with the domain in position 203,824 (`teatrebarcelona`)².
- We generate 100,000 domain names with our DGA with a given seed as the malicious ground truth, add the 100,000 Alexa-domains as the benign ground truth, shuffle these records and perform classification on this data set with 5-fold cross validation.

With the results of the classification, adversaries can iteratively improve their DGAs (as illustrated in Figure 5), until a desired suboptimal classification result is achieved, therefore effectively thwarting the DGA detection classifier. While creating our `deception_dga` DGA, we only used data from the Alexa list, mimicking the information available to a DGA author. However, for fairness of comparison, we evaluated the classifier performance of the Random Forest using the FANCI feature set and the Woodbridge LSTM on domains generated by `deception_dga`, using the benign ground truth which we used for evaluating the two classifiers in section 4.

5.1 The development of `deception_dga`

As mentioned, the development of the DGA used an iterative approach, with the goal of reducing the accuracy of the DGA detection classifiers. The different versions of the DGA are described below.

²Some domains appear multiple times in the Alexa list, with different TLDs. Those duplicates were skipped.

Figure 5: Iterative DGA development process

5.1.1 Version 0 - baseline

We created a baseline version, which would generate domain names with a length uniformly distributed between 13 and 18 characters, using the alphabet of ‘a’ - ‘z’, also uniformly distributed.

- Accuracy of detection: Random Forest: 0.987 LSTM: 1.000
- FANCI features influenced: 0/32

5.1.2 Version 1 - Modelling domain name length

FANCI’s first feature is the domain name length. This version of the DGA varies the length of the domain names generated, to make it follow the same statistical length distribution as observed in the Alexa benign ground truth. `deception_dga` now outputs domains with a length varying between 5 and 42 characters.

- Accuracy of detection: Random Forest: 0.873 LSTM: 0.993
- FANCI features influenced: 1/32

5.1.3 Version 2 - Adding vowel ratio

On top of version 1, this version also varies the vowel distribution in generated domain names, as observed in the Alexa ground truth for domain names of a specific length.

- Accuracy of detection: Random Forest: 0.715 LSTM: 0.949
- FANCI features influenced: 2/32

5.1.4 Version 3 - Using character probability distribution

This version abandons generation of characters with a uniform distribution between ‘a’ and ‘z’, and instead uses the character probability distribution found in the domain names recorded in the Alexa list. The full domain name alphabet ‘a’ - ‘z’, ‘0’ - ‘9’ and ‘-’ is used. Using the character distribution of the Alexa domains, implies that the *VowelRatio* feature is automatically modeled according to the Alexa ground truth, as well as the *DigitRatio*, *Alphabet Cardinality* and *Entropy* FANCI features, as well as its list of seven 1-gram features (*mean*, *stddev*, *min*, *max*, *median*, *lower_quartile* and *upper_quartile*). This change did not have a very big impact though: classification accuracy of the Random Forest even slightly increased, while the LSTM accuracy dropped, albeit significantly higher still than the Random Forest accuracy.

- Accuracy of detection: Random Forest: 0.722 LSTM: 0.929
- FANCI features influenced: 11/32

5.1.5 Version 4 - Using character probability, given the previous character

Further in the direction of version 3, we used the overall character probability distribution for the first character (for which we only allowed ‘a’ - ‘z’, no dash or digits). All subsequent characters use the distribution frequency as recorded in the Alexa ground truth for a character, given a specific previous character. This approach essentially models bigram-probabilities. This version of the DGA tries to present expected values to the following FANCI features: *Domain Name Length*, *Contains Digits*, *Vowel Ratio*, *Digit Ratio*, *Alphabet Cardinality*, *Ratio of Repeated Characters*, *Ratio of Consecutive Consonants*, *Ratio of Consecutive Digits*, *Entropy*, the seven 1-gram features and the seven 2-gram features.

- Accuracy of detection: Random Forest: 0.599 LSTM: 0.855
- FANCI features influenced: 23/32

This version of `deception_dga` was submitted as `deception_dga` to the DGA Archive dataset of algorithmically generated domain names.

5.1.6 Version 5 - Using unigram, bigram and trigram probability distributions

Further extending version 4, we now took the approach to model the full unigram distribution for the first character, bigram distribution for the second character and trigram distribution for all subsequent characters. As source for the statistical distribution, we took a slightly larger subset from the Alexa list, starting again from domain 100,001, but harvesting 250,000 unique domains, without TLD or subdomain, which were not internationalized domain names. This version of the DGA tries to present expected values to the following FANCI features: *Domain Name Length*, *Contains*

Figure 6: Change of classification accuracy when a certain number of features is actively ‘emulated’ by the DGA.

Digits, *Vowel Ratio*, *Digit Ratio*, *Alphabet Cardinality*, *Ratio of Repeated Characters*, *Ratio of Consecutive Consonants*, *Ratio of Consecutive Digits*, *Entropy*, the seven 1-gram features, the seven 2-gram features and the seven 3-gram features.

- Accuracy of detection: Random Forest: 0.573 LSTM: 0.789
- FANCI features influenced: 30/32

This version was submitted as `deception2_dga` to the DGA Archive dataset of algorithmically generated domain names.

5.2 Results

An overview of the effect in classification accuracy when different features are being ‘emulated’ by the DGA is shown in Figure 6³. As a general trend, as more features are emulated by the DGA, the Random Forest classifier accuracy starts to degrade. There is one remarkable exception, which is when switching from `deception_dga` version 3 to 4, the number of emulated features increases from 2 to 10, however the Random Forest classification accuracy slightly increases from 0.715 to 0.722. One possible explanation might be that *Domain Name Length* and *Vowel Ratio* are very heavy weighing features in the decision process of the Random Forest and the *Vowel Ratio* distribution produced when using the character frequency distribution of the Alexa domains may not be as exact as when the *Vowel Ratio* distribution is specifically enforced while generating domain names.

The new DGA also gives the LSTM a hard time, dropping its accuracy down to 0.789, which is the lowest recorded so far for the LSTM. Increasing the number of epochs on the LSTM from 10 to 15, increases the accuracy to 0.816.

5.3 Practical feasibility

We evaluate the practical usability of `deception_dga` for versions 4 and 5, which both break the Random Forest using the FANCI feature set:

³Please, note that the y-axis scale starts at 0.5, since accuracy below 0.5 is meaningless for binary classifiers.

5.3.1 Version 4 (deception_dga)

This version of `deception_dga`, described in section 5.1.5 consists of 535 lines of Python code, including comments and whitespace. This also includes the statistical parameters it collected from the Alexa benign ground truth, which we believe will not change much over time and need never be updated in the code. It does not call (import) any libraries and in fact implements its own *linear congruential generator* for generating random numbers, given a particular seed.

The code, using a single thread on a 2.6 GHz Intel Core i5 processor, generates over 6000 random domain names per second and is a perfectly usable DGA. It completely breaks the Random Forest using the FANCI feature set, since clearly, a *False Positive Rate* of 40% and a *True Positive Rate* of 60% prevent any useful deployment of this classifier.

5.3.2 Version 5 (deception2_dga)

This version, `deception2_dga`, described in section 5.1.6 consists of 163 lines of Python code, including comments and whitespace. The unigram, bigram and trigram statistics it uses can be kept in pickle files using a total of 452,938 bytes. The algorithm can generate 20 random domain names per second, using a single thread on a 2.6 GHz Intel Core i5 processor, making also `deception2_dga` a realistically usable DGA.

Since `deception2_dga` draws its unigram, bigram and trigram statistics from an Alexa-derived list of 250,000 unique domain names, we are interested in possible collisions, i.e., domain names generated by the DGA, which are actually part of this source list of Alexa-derived domains. Malware will typically be unaffected by collisions i.e., by domain names that have already been registered legitimately, since the goal of using a DGA is precisely to generate a large number of domains to become resilient to domain name seizure. Were the registration of one of those generated domain names to result in the malware to be unable to reach its C&C server, then the DGA would entirely defeat its purpose. Clearly, malware will only decide it found its new C&C server when it finds a domain name that can be resolved *and* which responds in the expected way on the expected TCP or UDP port. Our interest in collisions is therefore purely academic. One might expect some collisions, since the domains are generated along statistical probabilities for characters, given previous characters.

However, out of 100,000 generated domain names, only 57 domain names (0.057%) are found in the Alexa source list.

6. OTHER REFLECTIONS AND OBSERVATIONS

We demonstrated that the countermeasures used to evade detection with one specific classifier (Random Forest classifier for which the input features were documented), also affected the classification accuracy of another state-of-the-art LSTM-based classifier, albeit to a lesser extent. Although Woodbridge et al. claimed that LSTMs make it difficult for adversaries to circumvent a classifier without the same training set, our experiments showed that this claim turned out

Figure 7: Classification results when training on fewer samples - 5 worst detectable DGAs

Figure 8: Classification results when training on fewer samples - 5 best detectable DGAs. Random Forest results with full markers and dashed connection lines.

to be not entirely correct. Indeed, adversarial examples that are designed to mislead machine learning models at test-time often transfer, i.e. the same adversarial example fools more than one model. This transferability property has been explored before by Tramèr et al. [25]. An alternative evaluation path that could further confirm this hypothesis would have been to build a new DGA that aims to evade detection by the LSTM based detector, and then evaluating how well it is also able to circumvent the Random Forest based detector.

In the previous experiments, we have used a state-of-the-art Random Forest based DGA classifier and built a new DGA that circumvented this particular classifier. One can argue that given the fact that we explained the countermeasures used by the new DGA to evade detection, security specialists could in turn use this knowledge to further improve their classifiers. Building upon the previous observation, one could automate this arms race between adversaries and defenders with a deep learning method known as Gen-

erative Adversarial Networks (GAN) [6], as used in DeepDGA [2]. This class of unsupervised machine learning relies on two deep learning networks, where (1) the generative neural network produces new candidates with the objective to increase the error rate and (2) the discriminative neural network aims to distinguish malicious and benign instances thereby decreasing the error rate. During the training phase, each neural network further improves its own objective until an equilibrium is reached. The adversary would use the generative neural network to produce new domain names, whereas the defender would use the discriminative neural network to improve detection.

6.1 Amount of malicious training samples

In this work, we have selected a set of DGAs based on the number of available samples of 100,000 samples or more. This begs the question as to what the LSTM classification results would be, when less samples would be available to train the model on. We therefore repeated the classification using less training samples, both on the 5 most difficult to predict DGAs (Figure 7) and the 5 best predictable DGAs (Figure 8). Training and classification was performed with a random subset of the available training data, using a balanced training dataset (equal amounts of malicious and benign domain names), and again 5-fold cross validation. We increased the number of training epochs inversely proportional to the size of the training set, thereby ensuring that the amount of training of the LSTM network remained constant. We did not use an early stopping criterion, since this appeared to make the classification accuracy more variable at the highly zoomed-in accuracy scale.

Figure 7 only shows the LSTM classification results, since for none of these hard to classify DGAs, the Random Forest classifier reaches 90% classification accuracy – even for the full 100,000 malicious samples. Figure 8 shows the classification accuracy for the best predictable DGAs; the hollow markers and full connection lines refer to the LSTM classification accuracy, while the solid markers and broken connection lines refer to the Random Forest classifier, using the FANCI feature set. Please note that the scales of both axes in Figure 8 differ from the scale used in Figure 7!

For the difficult to classify DGAs, we observe that 50,000 malicious samples are sufficient to obtain good classification accuracy (less than 1% classification accuracy drop, compared to training with 100,000 malicious samples); but even when training on only 10,000 malicious samples, the classification accuracy is still acceptable ($> 92\%$) and greatly outperforms the Random Forest.

For the best detectable DGAs, we observe that 10,000 malicious samples are sufficient to obtain an LSTM classification accuracy better than 99.8% and when training on 2,500 malicious samples, the LSTM still outperforms the Random Forest.

7. CONCLUSIONS

In this work, we have compared two state of practice approaches in machine learning for the detection of algorithmically generated domain names, when a sufficient number of DGA-generated domain names is available for training.

We establish that the deep learning approach consistently outperforms the Random Forest with manually engineered features, where the classification accuracy over a set of 26 real-world DGAs averages 98.7% for the neural net, versus 93.7% for the Random Forest. A significantly lower standard deviation for the neural net accuracy over the different evaluated DGAs also indicates better classification consistency over the different DGAs.

Moreover, when training the two classifiers with a subset of the available malicious ground truth for the 5 best detectable and 5 worst detectable DGAs, we observe that even for only 10,000 available generated domain names, the deep learning approach outperforms the Random Forest.

Additionally, we investigated the concern that knowledge of manually engineered features may be abused by DGA developers to create new versions of DGAs which will become less detectable. To this end, we developed a new DGA, using the knowledge of the feature set used by the FANCI system. This DGA indeed renders the Random Forest classifier useless, with its classification accuracy reduced to 57.3%. The deep learning classifier is also affected by the approach of this new DGA, reducing its accuracy to 78.9%.

As future work, we consider the use of the generative neural network of a Generative Adversarial Network (GAN) to strengthen the circumvention capabilities of our DGA against existing deep learning based DGA detectors, while also evaluating whether there is any classification accuracy improvement by the discriminative neural network of this GAN.

8. ACKNOWLEDGMENTS

This research is partially funded by the Research Fund KU Leuven. The authors would like to thank Daniel Plohmman for his great work on the DGArchive and for making the results of his work available to the research community.

The authors gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

9. REFERENCES

- [1] ALEXA INTERNET, INC. Alexa. <https://www.alexa.com/topsites/>, 2018. [Online; accessed 16-September-2018].
- [2] ANDERSON, H. S., WOODBRIDGE, J., AND FILAR, B. Deepdga: Adversarially-tuned domain generation and detection. In *Proceedings of the 2016 ACM Workshop on Artificial Intelligence and Security* (New York, NY, USA, 2016), AISec '16, ACM, pp. 13–21.
- [3] ANTONAKAKIS, M., PERDISCI, R., NADJI, Y., VASILOGLOU, N., ABU-NIMEH, S., LEE, W., AND DAGON, D. From throw-away traffic to bots: Detecting the rise of dga-based malware. In *Presented as part of the 21st USENIX Security Symposium (USENIX Security 12)* (Bellevue, WA, 2012), USENIX, pp. 491–506.
- [4] COOKE, E., JAHANIAN, F., AND MCPHERSON, D. The zombie roundup: Understanding, detecting, and disrupting botnets.

- [5] DU, M., LI, F., ZHENG, G., AND SRIKUMAR, V. Deeplog: Anomaly detection and diagnosis from system logs through deep learning. In *CCS* (2017), ACM, pp. 1285–1298.
- [6] GOODFELLOW, I. J., POUGET-ABADIE, J., MIRZA, M., XU, B., WARDE-FARLEY, D., OZAI, S., COURVILLE, A., AND BENGIO, Y. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2* (Cambridge, MA, USA, 2014), NIPS’14, MIT Press, pp. 2672–2680.
- [7] GOOGLE. Google Safe Browsing List. <https://developers.google.com/safe-browsing/>, 2017. [Online; accessed August-2017].
- [8] HOCHREITER, S., AND SCHMIDHUBER, J. Long short-term memory. *Neural Comput.* 9, 9 (Nov. 1997), 1735–1780.
- [9] HYRUM ANDERSON. Endgame GitHub. https://github.com/endgameinc/dga_predict/, 2018. [Online; accessed 20-August-2018].
- [10] LE POCHAT, V., VAN GOETHEM, T., TAJALIZADEHKHOOB, S., KORCZYNSKI, M., AND JOOSEN, W. Tranco: A research-oriented top sites ranking hardened against manipulation. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS 2019)* (Feb. 2019).
- [11] LI, Z., ZOU, D., XU, S., OU, X., JIN, H., WANG, S., DENG, Z., AND ZHONG, Y. Vuldeepecker: A deep learning-based system for vulnerability detection. In *NDSS* (2018), The Internet Society.
- [12] MOWBRAY, M., AND HAGEN, J. Finding domain-generation algorithms by looking at length distribution. In *25th IEEE International Symposium on Software Reliability Engineering Workshops, ISSRE Workshops, Naples, Italy, November 3-6, 2014* (2014), IEEE Computer Society, pp. 395–400.
- [13] PECK, J., NIE, C., SIVAGURU, R., GRUMER, C., OLUMOFIN, F., YU, B., NASCIMENTO, A., AND COCK, M. D. Charbot: A simple and effective method for evading dga classifiers, 2019.
- [14] PEREIRA, M., COLEMAN, S., YU, B., DECOCK, M., AND NASCIMENTO, A. Dictionary extraction and detection of algorithmically generated domain names in passive dns traffic: 21st international symposium, raid 2018, heraklion, crete, greece, september 10-12, 2018, proceedings. 295–314.
- [15] PLOHMANN, D. DGArchive. <https://dgarchive.caad.fkie.fraunhofer.de/>, 2018. [Online; accessed 10-September-2018].
- [16] PLOHMANN, D., YAKDAN, K., KLATT, M., BADER, J., AND GERHARDS-PADILLA, E. A comprehensive measurement study of domain generating malware. In *25th USENIX Security Symposium (USENIX Security 16)* (Austin, TX, 2016), USENIX Association, pp. 263–278.
- [17] RIMMER, V., PREUVENEERS, D., JUÁREZ, M., VAN GOETHEM, T., AND JOOSEN, W. Automated feature extraction for website fingerprinting through deep learning. *CoRR abs/1708.06376* (2017).
- [18] ROSSOW, C., ANDRIESSE, D., WERNER, T., STONE-GROSS, B., PLOHMANN, D., DIETRICH, C. J., AND BOS, H. Sok: P2pwned - modeling and evaluating the resilience of peer-to-peer botnets. *2013 IEEE Symposium on Security and Privacy* (2013), 97–111.
- [19] SCHIAVONI, S., MAGGI, F., CAVALLARO, L., AND ZANERO, S. Phoenix: Dga-based botnet tracking and intelligence. In *Detection of Intrusions and Malware, and Vulnerability Assessment* (Cham, 2014), S. Dietrich, Ed., Springer International Publishing, pp. 192–211.
- [20] SCHÜPPEN, S., TEUBERT, D., HERRMANN, P., AND MEYER, U. FANCI : Feature-based automated nxdomain classification and intelligence. In *27th USENIX Security Symposium (USENIX Security 18)* (Baltimore, MD, 2018), USENIX Association, pp. 1165–1181.
- [21] SIDI, L., NADLER, A., AND SHABTAI, A. Maskdga: A black-box evasion technique against dga classifiers and adversarial defenses, 2019.
- [22] SPOOREN, J., PREUVENEERS, D., DESMET, L., JANSSEN, P., AND JOOSEN, W. Detection of algorithmically generated domain names used by botnets: A dual arms race. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing* (New York, NY, USA, 2019), SAC ’19, ACM, pp. 1916–1923.
- [23] SURBL.ORG. SURBL - URI Reputation Data. <http://www.surbl.org/>, 2017. [Online; accessed August-2017].
- [24] THE SPAMHAUS PROJECT LTD. The Domain Block List. <https://www.spamhaus.org/dbl/>, 2017. [Online; accessed August-2017].
- [25] TRAMÈER, F., PAPERNOT, N., GOODFELLOW, I., BONEH, D., AND MCDANIEL, P. The space of transferable adversarial examples. *arXiv* (2017).
- [26] WITTEN, I. H., FRANK, E., HALL, M. A., AND PAL, C. J. *Data Mining, Fourth Edition: Practical Machine Learning Tools and Techniques*, 4th ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2016.
- [27] WOODBRIDGE, J., ANDERSON, H. S., AHUJA, A., AND GRANT, D. Predicting domain generation algorithms with long short-term memory networks. *CoRR abs/1611.00791* (2016).
- [28] YADAV, S., REDDY, A. K. K., REDDY, A. L. N., AND RANJAN, S. Detecting algorithmically generated domain-flux attacks with dns traffic analysis. *IEEE/ACM Transactions on Networking* 20, 5 (Oct 2012), 1663–1677.