

Sharing FANCI Features: A Privacy Analysis of Feature Extraction for DGA Detection

Benedikt Holmes
RWTH Aachen University
Aachen, Germany
holmes@itsec.rwth-aachen.de

Arthur Drichel
RWTH Aachen University
Aachen, Germany
drichel@itsec.rwth-aachen.de

Ulrike Meyer
RWTH Aachen University
Aachen, Germany
meyer@itsec.rwth-aachen.de

Abstract—The goal of Domain Generation Algorithm (DGA) detection is to recognize infections with bot malware and is often done with help of Machine Learning approaches that classify non-resolving Domain Name System (DNS) traffic and are trained on possibly sensitive data. In parallel, the rise of privacy research in the Machine Learning world leads to privacy-preserving measures that are tightly coupled with a deep learning model’s architecture or training routine, while non deep learning approaches are commonly better suited for the application of privacy-enhancing methods outside the actual classification module. In this work, we aim to measure the privacy capability of the feature extractor of feature-based DGA detector FANCI (Feature-based Automated Nxdomain Classification and Intelligence). Our goal is to assess whether a data-rich adversary can learn an inverse mapping of FANCI’s feature extractor and thereby reconstruct domain names from feature vectors. Attack success would pose a privacy threat to sharing FANCI’s feature representation, while the opposite would enable this representation to be shared without privacy concerns. Using three real-world data sets, we train a recurrent Machine Learning model on the reconstruction task. Our approaches result in poor reconstruction performance and we attempt to back our findings with a mathematical review of the feature extraction process. We thus reckon that sharing FANCI’s feature representation does not constitute a considerable privacy leakage.

Keywords—Data privacy; Intrusion detection; Machine learning.

I. INTRODUCTION

Machine Learning (ML) has had great success in solving advanced data-driven problems and its application also yields great performance for solving the Domain Generation Algorithm (DGA) classification problem. Instead of using static IP-addresses or domain names, bots use DGAs to generate pseudo-random domain names and then query the Domain Name System (DNS) to obtain the IP address of their command and control server. The botnet herder knows the DGA generation scheme and is therefore able to register a subset of the generated domains, while the connection is now more difficult for the defender to block. Most of the bot’s queries result in non-existing domain (NXD) responses as only the domain names that are registered in advance are resolved to valid IP addresses. ML classifiers can be trained to separate benign NXDs, e.g., caused by typos or misconfigured software, from DGA generated domains. Thereby, DGA activities can be detected even before bots receive instructions from the herder.

For reasons such as the availability, diversity, or size of data, it is uncommon that ML models are trained solely on a large data set obtained from a single source. On the other hand, collecting or sharing sensitive data is a privacy-concern. For ML-based DGA detection on NX traffic, the malicious training samples are publicly sourced, e.g., obtained from DGArchive [1], while samples of benign NXDs are often locally collected and can contain privacy-sensitive information as their disclosure may allow drawing conclusions about sensitive activity on the network, e.g., usage of a particular software or end-user browsing. Deep Learning (DL) is designed to allow models to directly receive raw data as input and therefore privacy-preserving measures are often coupled with the training routine. Non-DL approaches are commonly preceded by a feature extraction stage that performs a data transformation with the goal of reducing size of the data while increasing expressiveness by compressing data samples to finite and fixed length vectors. Whether such transformation can also yield a sufficiently abstract data representation able to hide sensitive information is our main research focus.

In this work, we thus take a step back from the advances in DL-based DGA detection and reconsider a simpler, feature-based DGA detection approach and evaluate its practicability towards privacy-preserving intelligence sharing: FANCI (Feature-based Automated Nxdomain Classification and Intelligence) [2] is the first feature-based classifier that achieves significant performance in DGA detection while only considering few hand-crafted features. Complementing the research on its classification performance [2], we investigate whether FANCI’s public feature extractor is prone to malicious inversion. More concretely, we ask whether knowledge of FANCI features threatens the disclosure of the original domain names as the latter could be reconstructable from features. If the feature extraction process can be deemed inversion-resilient, then this allows the risk-free publication of sensitive NX data in the form of FANCI’s feature representation and would thus enable to provide data-privacy in otherwise privacy-concerning sharing tasks, e.g., (1) collaborative learning approaches in which many parties join their data or (2) classification outsourcing in which DGA detection is offered as a service.

Our approaches exhibit poor reconstruction performance even when provided with real-world data samples. Consequently, we believe that FANCI’s feature extractor is hard to

invert, which motivates low-risk publication of feature vector sets for aforementioned sharing scenarios.

The work is structured as follows: Sections II & III elaborate on relevant related work and preliminaries such as FANCI and its feature extractor. Section IV gives a mathematical review of the feature extraction process to assess the limitations of any reconstruction approach. Then, these insights are used to motivate the subsequent data-driven approach, detailed in Section V & VI, in which a DL model is trained to learn a reconstruction mapping based on three large real-world NX data sets. These allow us to assess whether a reconstructor trained on one data set may perform well on another data set at test time. Results, quantified by a normalized edit distance, are presented in Section VII followed by a discussion in Section VIII. Finally, Section IX concludes the paper with an outlook on future work.

II. RELATED WORK

We briefly give an overview of DGA detection methodologies and position ourselves in the research area of ML privacy.

A. DGA Detection

A variety of different DGA detection techniques have been proposed in the past, which can broadly be divided into context-less [2]–[6] and context-aware approaches [7]–[12]. Context-less approaches only use information that can be extracted from a single domain name to determine whether a domain name is benign or malicious while context-aware approaches use additional contextual information to improve classification performance. Previous studies suggest that context-less approaches achieve state-of-the-art detection performance while being less resource intensive and less privacy invasive than context-aware approaches [2]–[4][6].

The context-less approaches can further be divided into feature-based classifiers such as random forests or support vector machines (e.g., [2]), and feature-less classifiers such as recurrent, convolutional, or residual neural networks [3][4][6]. The former group uses domain knowledge to extract hand-crafted features from a single domain name prior to classification. The latter group of approaches consists of DL classifiers that learn to extract valuable features on their own, yet require many training samples.

The main object under study is the context-less and feature-based DGA detector FANCI [2] that comprises a feature extractor and implements a random forest classification module.

B. Privacy in Machine Learning

ML has become the main suspect of privacy research investigating threats and defenses regarding models' and training procedures' natural information leakage of the consumed sensitive training data (e.g., [13]–[18]). The attack class that relates closest to our work is Model Inversion [14][19] which, for a given model output, iteratively searches for the best fit input candidate based on some likelihood maximization scheme, e.g., by misusing loss and gradient of a neural network model. In our work however, the object under study is not a

probability-outputting classifier, but rather just a data transformation module. While the general goal of our studied threat and that of Model Inversion are aligned (finding a suitable input for a given output), our work more closely matches the terminology of [17]: There, the term *reconstruction* specifies malicious inversions of the feature extraction stage with the goal to map features back to raw training data samples.

III. PRELIMINARIES

This section briefly introduces FANCI's feature extractor and presents the concept of Sequence-to-Sequence learning, which we leverage as reconstruction tool later in the study.

A. FANCI'S Feature extractor

We utilize the most recent open-source implementation of FANCI's feature extractor [20], which extracts 15 structural, 8 linguistic, and 22 statistical features from domain names as listed in Table I. For some features, an according footnote highlights that the implementation deviates from the definition in the original paper [2], e.g., `contains_ipv4_addr` should also regard IPv6 addresses. The feature extraction recognizes 39 unique characters (letters a-z, digits 0-9 and special characters dot, hyphen and underscore). For our study, we flatten the representation of the feature vector: Feature `number_of_subdomains` is represented as a one-hot encoded vector and clips the number of sub-domains at value 4. Although it is in fact just one feature, we keep the representation via four values. Similarly, we also view each entry of the one-, two-, and three-gram distribution vectors as single feature. Thereby, our feature count differs slightly from the one presented in the original work [2] and we end up with 45-component feature vectors.

As marked in Table I, many of FANCI's features are computed on the Dot-free public-Suffix-Free (DSF) part of the domain which excludes both dot characters and the public valid suffix, which is usually the Top-Level-Domain (TLD). The validity of a suffix is determined by checking against a predefined list that is included in the feature extractor.

For the rest of this work, we formally refer to the feature extractor as a function $E: S \rightarrow F$ mapping domains from S to the feature space F , where S is the set of strings over the 39-character alphabet with lengths up to 253.

B. Sequence-to-Sequence Learning

Sequence-to-Sequence learning (Seq2Seq) encompasses encoder-decoder models that solve ML tasks related to mapping variable-length input sequences to variable-length output sequences [21][22]. Usually, both the encoder and decoder of a Seq2Seq architecture utilize recurrent units to process the variable-length sequences and work together as follows: The encoder consumes and compresses the input sequence to a fixed-length state while the decoder is trained to create the target sequence from this state. A common use case is machine language translation on token sequences (i.e., words or characters). To train the model, bounds of the output sequences must be encoded in some fashion such that the token-wise decoding

TABLE I
FEATURES EXTRACTED FOR FANCI

Feature	Name	Type	Choices	Normalized by
1	length	integer	250	253
2-5	number_of_subdomains ^a	integer	1	1
6	subdomain_lengths_mean ^a	rational	250	length
7	contains_wwwdot	binary	2	1
8	has_valid_tld	binary	2	1
9	one_char_subdomains ^a	binary	2	1
10	prefix_repetition	binary	2	1
11	contains_tld_as_infix ^a	binary	2	1
12	only_digits_subdomains ^a	binary	2	1
13	only_hex_subdomains_ratio ^a	rational	250+1	1
14	underscore_ratio ^a	rational	250+1	1
15	contains_ipv4_addr ^{a,c}	binary	2	1
16	contains_digits ^a	binary	2	1
17	vowel_ratio ^{a,b}	rational	250+1	1
18	digit_ratio ^{a,b}	rational	250+1	1
19	char_diversity ^{a,b,c}	rational	250	1
20	alphabet_size ^{a,b}	integer	38	38
21	ratio_of_repeated_chars ^{a,b}	rational	38+1	1
22	consecutive_consonant_ratio ^{a,b}	rational	250+1	1
23	consecutive_digits_ratio ^{a,b}	rational	250+1	1
24,31,38	for $n \in \{1, 2, 3\}$: n -grams_std ^{a,b}	rational	1	$\left(\begin{array}{l} \text{features of} \\ n\text{-grams} \\ \text{are} \\ \text{normalized} \\ \text{by their} \\ \text{respective} \\ \text{max value} \end{array} \right)$ $\log_2(\text{alphabet_size})$
25,32,39	n -grams_median ^{a,b}	rational	250	
26,33,40	n -grams_mean ^{a,b}	rational	1	
27,34,41	n -grams_min ^{a,b}	integer	250+1	
28,35,42	n -grams_max ^{a,b}	integer	250	
29,36,43	n -grams_perc_25 ^{a,b}	rational	250	
30,37,44	n -grams_perc_75 ^{a,b}	rational	250	
45	shannon_entropy ^{a,b}	rational	≈ 194	

^aFeature ignores public suffix.

^bFeature ignores dots.

^cDefinition of feature in implementation deviates from original paper.

process begins with a start marker and can be stopped once the end marker is encountered or predicted. At test time a sequence can be sampled from the decoder of a trained Seq2Seq model, i.e., beginning with the start marker, the model iteratively predicts the next character with the currently sampled prefix as prior. This sampling technique is commonly referred to as closed-loop, since the predicted characters are fed back into the model at each step.

IV. A MATHEMATICAL REVIEW

Here we view the plain mathematical definition of the feature extraction process as such and assess the invertability of the process. The goal of inversion is to find a valid function $E^{-1}: F \rightarrow S$. Due to the feature extractor $E: S \rightarrow F$ not being bijective, function E^{-1} can obviously only fulfill $E^{-1}(E(s)) = s$ for samples $s \in R$ of a certain subset $R \subset S$ for which we are concerned that it includes real-world NXDs. Estimating E^{-1} by sampling a complete look-up table would require iterating over all domains in S . In theory, the domain space is of size $|S| = \sum_{i=4}^{253} 39^i \approx 3.55 \cdot 10^{402}$. Size of the feature space F can be estimated based on the multiplication of possible choices for each value in a feature vector (see Table I). For this, we respect the specification's maximum domain length of 253 [23] and choose the minimum length to be four. For the rational-valued features the number of choices is determined by the size of the divisor or if the divisor is another feature, then we view the number of choices as the maximum possibilities for the dividend. Occasionally, the dividend is allowed to be zero, which is denoted by a "+1" in Table I. For the `entropy` feature we estimate the average amount of distinct values it can accommodate. For features that are dependent on others we view the amount of

choices as fixed, or "1" in Table I. Finally, we approximate that $|F| \approx 2.71 \cdot 10^{65}$. Consequently, the feature extraction process performs a reduction of order of magnitude $1.311 \cdot 10^{337}$, i.e., there may on average exist 10^{337} pre-images for each feature vector. In the best case, where all pre-images are distributed equally among all images, inversion would thus be impossible.

A. Inference of new Information

It is possible to infer new information about the original domain sample $s \in S$ via the combination of existing features in $f = E(s)$ and thereby more accurately capture the number of possible pre-images. The information is new in the sense that it is not previously held directly as a value in f . Examples of inferable information are listed in Table II. Further, Shannon entropy is computed as weighted sum of character frequencies $H = \sum_{i=1}^n \text{freq}(x_i) \cdot \log(\text{freq}(x_i))$ with restrictions $\sum_{i=1}^n x_i = p_1$ and $x_i \in \mathbb{N}^{\geq 1} \forall i \in \{1, \dots, n\}$. The underlying character frequency distribution $\text{freq}(\cdot)$ over $n = \text{alphabet_size}$ unknown characters is uniquely determined based on p_1 and entropy.

B. Limitations

Finding the unique solution for the discrete frequency distribution that matches the entropy feature may require enumerating all solution candidates. Due to the way the entropy is calculated, the number of solution candidates is given by the binomial coefficient $\binom{n-1}{k-1}$. It is further possible to estimate the amount of unique digits (u_d), vowels (u_v) and other characters (u_o) by iterating over all valid allocations of bins in the frequency distribution to one of the three groups such that the sum of frequencies for each group's bins matches with the previously inferred total count (i.e., p_4 , p_5 and p_6). There is not necessarily a unique solution to this.

With help of combinatorics we can specify a tighter bound on the number of possible pre-images per feature vector f . First, `dsf-struct(f)` captures the possibilities to structure the DSF in (1): Given the inferred total occurrences of digits (p_4), vowels (p_5) and others (p_6), one can virtually choose p_4 character slots in the DSF of length p_1 and similarly p_5 slots of the remaining $p_1 - p_4$. The final slots are for the third group. Then, there are $\binom{p_1-1}{p_2-1}$ possibilities to split the DSF into p_2 sub-domains by inserting the $p_2 - 1$ separating dots.

TABLE II
NEW INFORMATION INFERABLE FROM OTHER FANCI FEATURES

ID	New Information	Inference Rule
p_1	Length of the DSF	$p_1 = \frac{\text{alphabet_size}}{\text{char_diversity}}$
p_2	Amount of sub-domains	$p_2 = \frac{p_1}{\text{subdomain_lengths_mean}}$
p_3	Length of the public suffix	$p_3 = \text{length} - (p_1 + p_2)$
p_4	Total digit occurrences	$p_4 = \text{digit_ratio} \cdot p_1$
p_5	Total vowel occurrences	$p_5 = \text{vowel_ratio} \cdot (p_1 - p_4)$
p_6	Occurrences other chars	$p_6 = p_1 - (p_4 + p_5)$

$$\text{dsf-struct}(f) = \binom{p_1}{p_4} \cdot \binom{p_1 - p_4}{p_5} \cdot \binom{p_1 - 1}{p_2 - 1} \quad (1)$$

Secondly, for a fixed valid setting of unique character counts $u = (u_d, u_v, u_o)$ we can estimate the following: Possibilities for digit occurrences in the DSF is determined by the choices of an u_d -large subset of all digits and all permutations of each subset of length of the total occurrences of digits in the DSF, i.e., p_4 . Same holds for unique count of vowels (u_v & p_5) and others (u_o & p_6). Since there does not have to be a unique solution for u , one needs to sum over the possible choices for u . Similarly, we can thus capture the total amount of possibilities for the DSF's content by $\text{dsf-cont}(f)$ in (2):

$$\text{dsf-cont}(f) = \sum_u \binom{10}{u_d} \cdot \binom{5}{u_v} \cdot \binom{24}{u_o} \cdot u_d^{p_4} \cdot u_v^{p_5} \cdot u_o^{p_6} \quad (2)$$

Finally, the public suffix list used by the feature extractor fixes the amount of choices for the public suffix or TLD with known length to t_f and in total this results in (3), which more reasonably models the number of possible pre-images.

$$\text{pre-images}(f) = t_f \cdot \text{dsf-struct}(f) \cdot \text{dsf-cont}(f) \quad (3)$$

However, the number of pre-images still remains significantly large and improving this manual reconstruction approach ultimately fails due to the lack of more linguistic information: Even if a frequency distribution can be determined, then the allocation of characters to those frequencies remains undetermined as that information is not held in the feature vector itself. Clearly, the function is not bijective and it is impossible to distinguish between equally likely pre-images. Arguing, to which extent more useful information can be extracted or whether a different manual approach would be more beneficial is a complex matter, which is why we attempt to let a DL model learn a reconstruction mapping based on real-world data.

V. METHODOLOGY

In reality, neither the amount of valid pre-images is equally distributed among all possible feature vectors nor are all pre-images for one feature vector equally likely. In fact, real-world examples for benign NXDs and their corresponding feature vectors will only make up a small fraction of the respective domain space S and feature space F : Besides that some subspace of S (and thereby also a subspace of F) is occupied by the malicious samples, benign NXDs that result from typographical errors may still exhibit linguistic characteristics that are of low-entropy. For feature vectors, we argue that there are semantically invalid combinations of features, e.g., `alphabet_size = 1` while both `vowel_ratio > 0` and `digit_ratio > 0`. Subsequently, the feature extractor will only act on a restriction of the mapping $S \rightarrow F$ in reality.

True pre-image distributions and domain-feature relations are best captured by real-world NXD samples, and hence, we leverage such data sets two-fold: (1) To train a DL model that

may learn the distribution of the sample data and (2) as ground truth to assess the reconstruction capability of the trained models. The rest of this section defines the methodology for the experiment, and the evaluation of a DL reconstructor.

A. Attack Model

The context in which the following experiment is conducted is defined by the following aspects: (1) We assume an adversary that is interested in learning the real inputs to the FANCI feature extractor $E : S \rightarrow F$ for a foreign feature set $S' \subset S$ of a target, i.e., for any $E(s) = f \in F$, the adversary aims to find a corresponding $s' \in S$ such that $E(s') = f$ holds and some closeness $s' \approx s$ is satisfied (Note, that finding just any $s' \in S$ with $E(s') = f$ is trivial). (2) The adversary is semi-honest, i.e., he reliably participates in any sharing scenario through which he acquires the foreign feature set. (3) The feature extractor is public knowledge. (4) We only consider the disclosure of benign NXDs as privacy critical. (5) We assume feature sets are shared in the clear, hence no interaction with the target is required. (6) We allow the adversary to be in possession of an arbitrary large data set $S'' \subset S$ of benign NXDs that does not intersect with the target's data, i.e., $S'' \cap S' = \emptyset$. (7) We do not restrict the adversary's computational power that he may apply to his own data. Hence, we allow the adversary to train an ML model.

B. Reconstruction Quantification

We leverage existing members from the family of edit distances on the string space to compare pairs of original and reconstructed samples. Due to the possible encounter of unequal lengths, the only suitable candidates are the *Levenshtein* [24] distance metric and its variant *Damerau-Levenshtein*, which both compute a minimum-change distance via the number of character edit operations (substitutions, insertions or deletions) required to transform the one input string into the other. We use the latter of both metrics which additionally considers the transposition of adjacent characters as a single operation. Further, we compute a normalized version for the metric by dividing the resulting minimum-change-distance by the length of the longest of both input strings. This division operation invalidates none of the metric's axioms. Note, however, that the normalized metric is a ratio of edit-operations to string length and which can be interpreted as a lower bound on the percentage of misplaced characters in the longer input string.

Consequently, a metric value of zero indicates equality, while dissimilarity grows in parallel to larger metric values. Thereby, quantifying the closeness, as previously mentioned in the attack model, becomes possible. Note that the choice of any threshold $\varepsilon > |s - s'|$ indicating attack success is subjective, as the metric does not regard a semantic comparison.

C. Benign Data Sets

In the following, we briefly comment on the nature and origin of the real-life benign NXD data we use, that are sourced locally by distinct institutions in different countries.

1) *University_A*: RWTH Aachen University in Germany provided us with a record comprising approximately 26 million unique NXDs recorded in the month of September 2019 by their central DNS resolver. This resolver handles academic and administrative networks, the university hospital as well as networks of student residences.

2) *University_B*: We obtained another data set comprising 8 million unique samples recorded between mid-May 2020 and mid-June 2020 at Masaryk University that is located in the Czech Republic.

3) *Association*: CESNET is a 27-member association of Czech universities which develops and operates a national e-infrastructure for science, research, and education, including several university networks. We obtained a partial quantity of a one-day recording on June 16th 2020 containing approximately 362k unique samples.

We use the complete record of the Association and draw a random sub-sample in the size of the Association's record from each of the other two institutions' records. Intersections with one another and with malicious samples drawn from the open source intelligence feed of DGArchive [1] (up until September 1st 2020) are removed from all records prior to sub-sampling.

D. Evaluation Setup

In the following experiment, we assess the reconstruction performance of trained DL reconstruction models via the above mentioned distance metrics. More concretely, we train a DL model for each one of the benign data sources and evaluate each of these models against all the data sources including the one on which the individual model was trained on. For each pair of training and evaluation sets we average each metric's scores over all samples. Thereby, we assess the models' capability to reconstruct domains from foreign feature sets.

VI. DATA-DRIVEN RECONSTRUCTION

The following describes the training setup for the Seq2Seq decoder which is trained on the task of domain sample reconstruction (i.e., learning an inverse mapping $E^{-1}|_R: F \rightarrow S$ on subset R) using an attack set of benign NXDs and their corresponding feature vectors $\{(f_i, s_i)\}_{i=0}^n \subset F \times S$. This is a realistic scenario in any sharing use case where a party receiving a feature set may also be in possession of an own data set of benign NXDs. Basically, we assume that the feature extractor is unknown to the model, and we let it learn the inverse mapping without any domain-specific assistance.

A. Model Architecture

To reconstruct a variable-length domain sample from a fixed-length feature vector, the decoder of a Seq2Seq model is utilized. All models share the same architecture whose design follows a related approach [22]. Beginning with two parallel sequences of two dense layers with 200 units each, this leaves opportunity for the model to manipulate the representation of the input feature vector before the two outputs are used as the initial states for the recurrent unit in the decoder. For the

recurrent unit, a single Long Short-Term Memory (LSTM) layer with 200 units is used. Finally, the model ends with a dense layer of size 42 and a softmax activation to output a prediction vector over all relevant characters, which includes the 39 recognized domain characters plus start, end and empty markers used internally for the sequence encoding of domains. In total, the architecture comprises 301,642 trainable weights.

B. Training Setup

For a good balance between training time and model performance, we fix a batch size of 64 for our experiment. Models are trained using the cross-entropy loss to penalize wrong character predictions. Being unaware of any unbalance-bias, a focal loss is used to dynamically down-weight well-classified samples in the cross-entropy loss during training [25].

Training data is prepared as follows: The test set is a random 20% split of the total data. Another random 5% split of the remaining training data is used as validation set. All entries in a FANCI feature vector are in some finite bounded range of the non-negative rationals and are normalized to the range of $[0, 1]$ by dividing each entry by the upper bound of its value range (see last column of Table I). Domain names are encoded to character sequences with start and end markers.

Each model is allowed to train for at most 1000 epochs, while the training data is shuffled after each epoch and training is stopped early whenever 10 epochs without improvement of the validation loss are exceeded. We follow the common methodology to train a Seq2Seq model and thus employ Teacher Forcing to train the decoder [26]. This essentially sets the input of the decoder to the target sequence shifted by one time step (open loop) instead of feeding the decoder's outputs of previous time steps back into the model (closed loop).

VII. RESULTS

For each domain in an evaluation set, we sample a reconstructed domain from the trained reconstructor models using the normalized feature vector of the original domain as initial state input to the model. Averaged closed-loop reconstruction performance for all combinations of trained models and evaluation sets are given on the left side of Table III.

A. Baseline Reconstruction Performance

Rows in Table III with a highlighted evaluation set show the average Damerau-Levenshtein metric score for the case that the evaluation data is equal to all the data used to train, validate and test the model, and is to be interpreted as baseline reconstruction performance of a trained model.

Although the models achieve a small training and test loss, reconstruction performance is mediocre: For *University_A*, *University_B* and the Association we measure that on average respectively 47.85, 15.00 and 13.66 character edit operation separate each original domain and its reconstruction. The normalized version of the metric measures an average score for *University_A* and *University_B* that is just larger than 0.5, i.e., on average at least 50% of characters in each reconstruction are misplaced. For the Association, we measure a slightly smaller

TABLE III
CLOSED LOOP RECONSTRUCTION PERFORMANCE OF SEQ2SEQ RECONSTRUCTOR & FEATURE SPACE OVERLAP

Network Data Source		Averaged Reconstruction Performance		Feature Space Overlap				
Training	Evaluation	Dam-Leven.	norm.	#Unique FV ^a (Training Data)	Training Evaluation \cap	% of Eval Data	#Unique FV ^a (All Data)	% of Total Data
University _A	University _A	47.85	0.51	288118	-	-	3462	10.3
	University _B	23.22	0.72		5789	32.9		
	Association	20.61	0.66		4809	12.1		
University _B	University _A	72.94	0.75	182786	5789	11.5	3462	30.7
	University _B	15.00	0.53		-	-		
	Association	18.61	0.61		12985	41.8		
Association	University _A	62.07	0.67	169921	4809	24.6	3462	22.9
	University _B	20.01	0.65		12985	30.9		
	Association	13.66	0.46		-	-		

^aFV = Feature Vectors.

average score of 0.45. It seems that the models' baseline reconstruction performance is similarly bad on all data sets.

B. Transferability

The remaining lines in Table III demonstrate the trained models' reconstruction performance on data from foreign networks, i.e., exactly the scenario which we describe in our attack model. In all cases the reconstruction error is higher than in the baseline cases (score higher than 0.65) with the worst performance (0.75) in the case where the model trained on data from University_B is evaluated on that of University_A.

VIII. DISCUSSION

After re-consideration of the mathematical review of the feature extractor, it is plausible that the overall reconstruction performance is poor. After all, FANCI's feature extractor considers only very few features and thereby performs a compression of such extent which is tolerable for good classification performance but hinders good reconstruction quality. In the rest of this section we continue to argue about a quantifiable proof that E is not injective when restricted to the subspace of real-world benign NXDs and review the adversary's theoretical information gain for well-reconstructed domains.

A. Feature Space Overlap

We place our experiment in the scenario in which adversary and target NXD data are disjoint. This does, however, not imply that the sets of feature vectors of each respective data set are also disjoint. Therefore, we also quantify the overlap in feature space for the three data sets used in this study on the right side of Table III. First, it is important to note that although every data set contains approximately 362k unique samples, the amount of unique feature vectors is significantly lower which clearly indicates collisions in the feature space. Secondly, for every combination of two distinct NXD data sets we have an intersection of non-trivial size in the feature space, e.g., 11.5% of University_B's data intersects with University_A's and 41.8% of its data with that of the Association.

A large overlap in the feature space most certainly leads to a degraded reconstruction performance, as for the same feature vector the model may learn to reconstruct a domain different

from the one the adversary wants to sample at test time. The worst-performing baseline and transferability reconstructions (training data of University_B) coincides with the largest feature space overlap w.r.t. all data sets (see Table III).

B. Top 10% Reconstructions

The adversary has no clear way of estimating the confidence of a single reconstruction without ground truth unless he conducts an own analysis of which type of domains are reconstructed well using a second data set. Hence, we also discuss what he could potentially learn from good reconstructions by taking a closer look at the best 10% of all reconstructions for the transferability cases: The average reconstruction performance for the top 10% lies at 0.276. Further, approximately 45-55% of the top 10% performers are occupied by IPv4 and IPv6 reverse DNS lookups and 20-35% by spam-related or other DNS-related services, e.g., DNS blacklists.

We argue that the models perform so well in reconstructing these types of domains with (1) these domains' contents being well-structured, (2) sharing a large suffix, and (3) standing out by containing a lot of numerical characters. Hence, they occupy the sparse areas of the feature space around features such as a high `digit_ratio`, low `subdomains_lengths_mean` or a True value for features such as `only_digits_subdomains` `contains_ipv4_addr`. Further, these NXDs do not necessarily originate from user typos but rather from misconfigured software. This would also better explain the high occurrence of these types of domains in the data.

The question remains whether knowledge of reverse lookups and spam-services is privacy-sensitive information and we claim the opposite. After all, these domains do not reveal any information about end-user browsing or sensitive tooling usage in the network from which the data was sourced.

IX. CONCLUSION AND FUTURE WORK

In this study we analyzed the data privacy capabilities of the feature-based DGA detector FANCI. The main goal was to answer whether feature vectors of FANCI disclose any sensitive information about the original domain names. We provide mathematical reasoning for the success likelihood

for any best-case reconstruction attempt and demonstrate that a manual approach of inferring sensitive information from combination of features has its difficulties and most certainly has its limitations: Reconstruction cannot be easily performed on the basis of a single feature vector.

Therefore, we chose to emulate the logical approach a data-rich adversary would take, namely training an ML model to learn a reconstruction mapping. To provide significance to our results, we make use of three large real-world NXD sets fortunately made available to us. Finally, we find reconstruction performance of the trained models to be worse than desired: On average at least half of all character from a reconstructed domain are misplaced in the baseline cases. The models only perform best on foreign network's data for reverse lookups and other not-sensitive NXDs likely originating from misconfigured software. We find this to be the result of these domains sharing a large portion of the higher-level domains and occupying a special niche in the feature space.

Consequently, our experiment suggest that an ML model aiding in the attack cannot reliably reconstruct NXDs from foreign networks' FANCI feature vectors which would be, however, the main use case in an attack.

Due to its universality, our data-driven analysis approach can be used in the future to perform a similar privacy analysis on other feature extractors used for DGA detection. The general concept of the data-driven analysis approach can also be used for a privacy analysis of feature-based classifiers in other ML use cases.

ACKNOWLEDGMENTS

The authors would like to thank Masaryk University, CES-NET and Jens Hektor from the IT Center of RWTH Aachen University for providing NXD data. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 833418. Simulations were performed with computing resources granted by RWTH Aachen University under project rwth0438.

REFERENCES

- [1] D. Plohmann, K. Yakdan, M. Klatt, J. Bader, and E. Gerhards-Padilla, "A comprehensive measurement study of domain generating malware," in *USENIX Security Symposium*. USENIX Association, 2016, pp. 263–278.
- [2] S. Schüppen, D. Teubert, P. Herrmann, and U. Meyer, "FANCI: Feature-based automated nxddomain classification and intelligence," in *USENIX Security Symposium*. USENIX Association, 2018, pp. 1165–1181.
- [3] J. Woodbridge, H. S. Anderson, A. Ahuja, and D. Grant, "Predicting domain generation algorithms with long short-term memory networks," arXiv preprint arXiv:1611.00791, 2016.
- [4] B. Yu, J. Pan, J. Hu, A. Nascimento, and M. De Cock, "Character level based detection of DGA domain names," in *International Joint Conference on Neural Networks*. IEEE, 2018, pp. 1–8.
- [5] J. Saxe and K. Berlin, "eXpose: A character-level convolutional neural network with embeddings for detecting malicious URLs, file paths and registry keys." arXiv preprint arXiv:1702.08568, 2017.
- [6] A. Drichel, U. Meyer, S. Schüppen, and D. Teubert, "Analyzing the real-world applicability of DGA classifiers," in *Conference on Availability, Reliability and Security*. ACM, 2020, pp. 1–11.
- [7] M. Antonakakis *et al.*, "From throw-away traffic to bots: Detecting the rise of DGA-based malware," in *USENIX Security Symposium*. USENIX Association, 2012, pp. 491–506.
- [8] L. Bilge, S. Sen, D. Balzarotti, E. Kirda, and C. Kruegel, "Exposure: A passive DNS analysis service to detect and report malicious domains," in *Transactions on Information and System Security*. ACM, 2014, pp. 1–28.
- [9] M. Grill, I. Nikolaev, V. Valeros, and M. Rehak, "Detecting DGA malware using netflow," in *IFIP/IEEE International Symposium on Integrated Network Management*. IEEE, 2015, pp. 1304–1309.
- [10] S. Yadav and A. L. N. Reddy, "Winning with dns failures: Strategies for faster botnet detection," in *Security and Privacy in Communication Systems*. Springer, 2011, pp. 446–459.
- [11] S. Schiavoni, F. Maggi, L. Cavallaro, and S. Zanero, "Phoenix: DGA-based botnet tracking and intelligence," in *Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2014, pp. 192–211.
- [12] Y. Shi, G. Chen, and J. Li, "Malicious domain name detection based on extreme machine learning," *Neural Processing Letters*, pp. 1347–1357, 2018.
- [13] G. Ateniese *et al.*, "Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers," in *International Journal of Security and Networks*. Inderscience Publishers, 2015, pp. 137–150.
- [14] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Computer and Communications Security*. ACM, 2015, p. 1322–1333.
- [15] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *Symposium on Security and Privacy*. IEEE, 2017, pp. 3–18.
- [16] N. Papernot, P. McDaniel, A. Sinha, and M. P. Wellman, "Sok: Security and privacy in machine learning," in *European Symposium on Security and Privacy*. IEEE, 2018, pp. 399–414.
- [17] M. Al-Rubaie and J. M. Chang, "Privacy-preserving machine learning: Threats and solutions," in *Symposium on Security and Privacy*. IEEE, 2019, pp. 49–58.
- [18] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning," in *Symposium on Security and Privacy*. IEEE, 2019, pp. 739–753.
- [19] M. Fredrikson *et al.*, "Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing," in *USENIX Security Symposium*. USENIX Association, 2014, pp. 17–32.
- [20] Fanci: Feature-based automated nxddomain classification intelligence. [Online, retrieved: September, 2021]. <https://github.com/fanci-dga-detection/fanci/tree/d6c7d08>
- [21] K. Cho *et al.*, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2014, pp. 1724–1734.
- [22] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*. MIT Press, 2014, pp. 3104–3112.
- [23] Rfc 1034: Domain names - concepts and facilities. [Online, retrieved: September, 2021]. <https://datatracker.ietf.org/doc/html/rfc1034>
- [24] V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals," in *Soviet Physics Doklady*, 1966, pp. 707–710.
- [25] T.-Y. Lin, P. Goyal, R. B. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," *Transactions on Pattern Analysis and Machine Intelligence*, pp. 318–327, 2020.
- [26] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural Computation*, pp. 270–280, 1989.