PLOS ONE

RESEARCH ARTICLE

An end-to-end framework for private DGA detection as a service

Ricardo J. M. Maia[1]*, Dustin Ray[2], Sikha Pentyala[2], Rafael Dowsley[3], Martine De Cock[2,4], Anderson C. A. Nascimento[5], Ricardo Jacobi[1]

1 Department of Computer Science, University of Brasilia, Federal District, Brasília, Brazil, **2** School of Engineering, University of Washington Tacoma, Tacoma, Washington, United States of America, **3** Department of Software Systems and Cybersecurity, Monash University, Melbourne, Australia, **4** Department of Applied Mathematics, Computer Science and Statistics, Ghent University, Ghent, Belgium, **5** Visa Research (Work done while at University of Washington Tacoma), Foster City, California, United States of America

* ricardo.menezes@aluno.unb.br

Check for updates

OPEN ACCESS

Citation: Maia RJM, Ray D, Pentyala S, Dowsley R, De Cock M, Nascimento ACA, et al. (2024) An end-to-end framework for private DGA detection as a service. PLoS ONE 19(8): e0304476. https://doi.org/10.1371/journal.pone.0304476

Editor: Muhammad Nasir Khan, Government College University Lahore, PAKISTAN

Received: December 15, 2023

Accepted: May 13, 2024

Published: August 28, 2024

Copyright: © 2024 Maia et al. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper. To access the malign domains data present on the website https://dgarchive.caad.fkie.fraunhofer.de/ it is necessary to contact them first, as they own the data. For negative DGA matches, we have acquired approximately 1,000,000 domains from the last known version of the dataset 'Alexa top 1 million domains' https://en.wikipedia.org/wiki/Alexa_Internet.

Funding: The author(s) received no specific funding for this work.

Abstract

Domain Generation Algorithms (DGAs) are used by malware to generate pseudorandom domain names to establish communication between infected bots and command and control servers. While DGAs can be detected by machine learning (ML) models with great accuracy, offering DGA detection as a service raises privacy concerns when requiring network administrators to disclose their DNS traffic to the service provider. The main scientific contribution of this paper is to propose the first end-to-end framework for privacy-preserving classification as a service of domain names into DGA (malicious) or non-DGA (benign) domains. Our framework achieves these goals by carefully designed protocols that combine two privacy-enhancing technologies (PETs), namely secure multi-party computation (MPC) and differential privacy (DP). Through MPC, our framework enables an enterprise network administrator to outsource the problem of classifying a DNS (Domain Name System) domain as DGA or non-DGA to an external organization without revealing any information about the domain name. Moreover, the service provider's ML model used for DGA detection is never revealed to the network administrator. Furthermore, by using DP, we also ensure that the classification result cannot be used to learn information about individual entries of the training data. Finally, we leverage post-training float16 quantization of deep learning models in MPC to achieve efficient, secure DGA detection. We demonstrate that by using quantization achieves a significant speed-up, resulting in a 23% to 42% reduction in inference runtime without reducing accuracy using a three party secure computation protocol tolerating one corruption. Previous solutions are not end-to-end private, do not provide differential privacy guarantees for the model's outputs, and assume that model embeddings are publicly known. Our best protocol in terms of accuracy runs in about 0.22s.

1 Introduction

Malicious software (malware) is the class of software that infects computers to perform unauthorized actions in the system or gain unauthorized access to information. Malware is a highly

Competing interests: NO authors have competing interests.

significant source of illicit activities with increasing impact [1–4], and producing substantial losses in sectors such as the government, energy, and manufacturing [5]. Examples of malware families include trojan horses, viruses, ransomware, key loggers, worms, spyware, and hidden cryptominers. Some common objectives of these types of malware are information or identity theft, espionage, and service disruption [1, 4].

Botnets, i.e., computer networks infected by malware, are commonly controlled, operated, and updated through communicating with a Command and Control (C&C) server that is under the control of an adversary or botmaster [6]. When the IP address of the C&C server is hard-coded directly into the malware, intrusion detection systems (IDS) or firewalls on Domain Name System (DNS) servers can blacklist the detected malicious domain names and block the connection to the C&C server, effectively rendering the malware useless. Cyber-attackers have, therefore, adopted innovative techniques for obfuscating and concealing the C&C's identity. Among the most prevalent approaches are Domain Generation Algorithms (DGA) [7].

DGAs are algorithms that periodically generate pseudorandom combinations of characters or words to form hundreds or even thousands of new domain names. The key idea is that DGAs can generate the same set of new domain names when executed by two different machines, such as a botmaster and an infected machine. The botmaster registers one or more generated domain names, while the infected machines systematically query the domains from the generated list until one of them is resolved. The domains from the list that the botmaster has not registered will typically result in a non-existent domain response when queried and can be discarded by the infected machine. Once an infected machine queries a registered domain name, communication between the infected bot and the C&C center is established, and malicious activities, as instructed by the C&C center, can be performed by the bot.

The constant changes to the domain name of the C&C server make it much more difficult for IDS and firewalls to detect and contain the attacks. The challenge of mitigating attacks that use DGA techniques lies in identifying malicious domain names. The DNS server must be able to detect and block malicious domain names while keeping normal operations for benign domain names. In short, the ability to identify malicious domains can drastically decrease the harm caused by malware.

Using machine learning models to create classifiers that can identify and separate benign domains from DGA-generated malware domains is a viable approach (see [8] and references therein). Such classifiers (models) can be deployed as automatic malware detection systems in enterprise networks. The state-of-the-art models use deep learning techniques that achieve high accuracy but require large amounts of training data [9]. Due to this data demand, models are usually available with third-party organizations (service providers) as a DGA detection service where the DNS traffic of an enterprise is sent to the service provider who then classifies the incoming traffic into malicious or benign domains and sends the classification result to the enterprise [10]. Such an outsourced DGA detection as a service model presents numerous privacy challenges. The DNS traffic of an enterprise holds sensitive information that can impact the privacy of all the enterprise network users, which raises privacy concerns for the users in the above DGA detection as a service paradigm. A potential solution to address this concern is to make the ML model used for DGA detection available to enterprise network administrators to deploy it locally. This is, again, problematic as the model is proprietary to the service provider. Moreover, the data used for training such models can be private, and releasing the model to the enterprises renders the underlying training data vulnerable to attacks [11, 12]. Simultaneously protecting the sensitive data of the enterprise and the service provider is a significant challenge.

In this paper, we investigate the possibility of simultaneously providing privacy for the machine learning model holder and the party interested in performing the domain classification. Our paper proposes an end-to-end privacy-preserving framework for outsourced DGA classification that preserves enterprise users' and service providers' privacy. Our result allows network administrators to outsource traffic analysis (in this case, DNS traffic analysis) to outside parties so that no information about the DNS traffic is ever leaked. It also protects the intellectual property of the model holder, ensuring the computing parties cannot obtain any information on the model.

Scientific contributions

We propose a novel framework that provides the benefits of automated and outsourced DGA detection while preserving the privacy of enterprise network users' and DGA detection service providers' data. In our framework, neither the DNS traffic is released in the clear to the DGA detection service providers nor is the model made available to the enterprise network administrators (no private data is revealed to any other party). Furthermore, we incorporate techniques that prevent attackers from gaining additional information about the training data or reconstructing the model from the classification results sent to the enterprise network administrators.

Fig 1 illustrates our framework at a high level. *Alice* represents the enterprise and holds the DNS domains and traffic. *Bob* represents the service provider and has the weights (parameters) of the trained machine learning model.*Bob* can choose to train a multilayer perceptron (MLP), a one-dimensional convolutional neural network (1D-CNN), or a long short-term memory (LSTM), and the chosen model shall be trained with differential privacy guarantees [13]. Furthermore, *Alice* wants her DNS traffic classified by *Bob*. We hereafter refer to the *Alice's* DNS data and *Bob's* model weights as private data.

Our proposed framework employs techniques from secure multi-party computation [14] to preserve input privacy. To this end, *Alice* and *Bob* secret share their private data with a set of untrusted computational servers (parties). The MPC servers perform computations on the secret shares to label domain names as benign or malicious in a way such that:

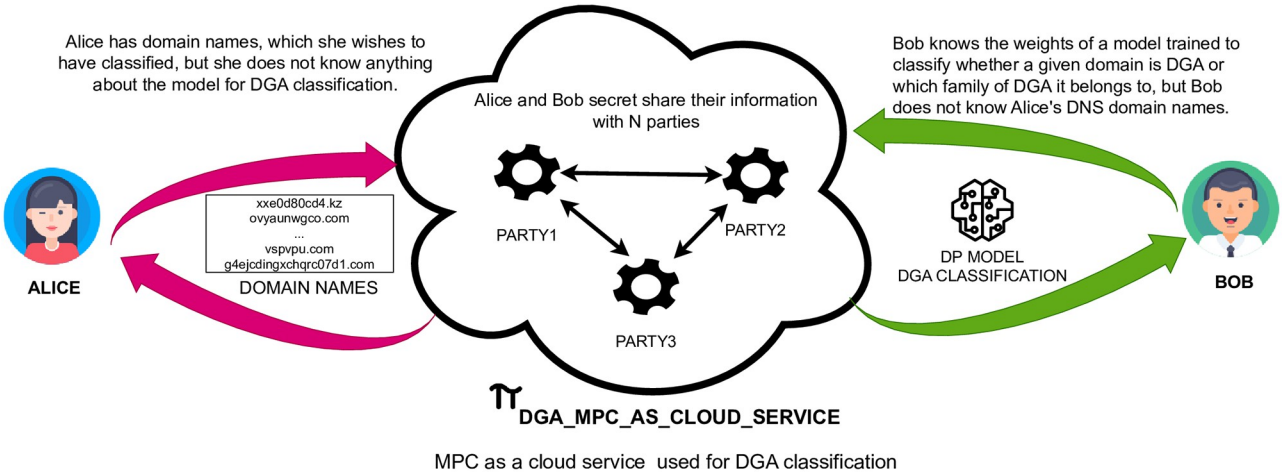

Fig 1. Flow diagram illustrating end-to-end privacy-preserving DGA detection as a service. The DNS domain name classifier of the service provider (Bob) is trained with DP-SGD to provide differential privacy (DP) guarantees (output privacy). New domain names coming from Alice are classified with Bob's model by secure multi-party computation (MPC). Protocols executed by MPC servers in the cloud over encrypted data (input privacy).

https://doi.org/10.1371/journal.pone.0304476.g001

(P1) No individual MPC server should obtain any information about *Alice's* domain names;

(P2) No individual MPC server should learn any information about the weights of *Bob's* machine learning model;

(P3) The result of the classification should be revealed only to *Alice*;

(P4) The result of the classification should reveal no private information about individual entries in *Bob's* training data set to *Alice*.

We note that (P4) provides *output privacy* and is achieved as *Bob* has trained the model with DP guarantees. (P1)–(P3) provide *input privacy* and are achieved through the novel MPC protocols that we propose for inference with a neural network model trained for DGA detection. MPC protocols usually lead to high communication and computation costs, thus impacting the inference runtime and overall performance. To improve the performance of our proposed MPC protocols for secure classification of DGA domains, we leverage the benefits of quantization schemes available in TensorFlow (TFLite). The quantization method reduces the precision of a machine learning model's parameters, typically 32-bit floating point numbers [15]. We propose that *Bob* uses post-training quantization techniques on the DP-trained model before using the model for classification.

The closest related work [16] to ours on privacy-preserving detection of DGAs provides only input privacy through MPC, leaving the inference phase vulnerable to privacy attacks on the training data (see Section 2 for more details). We propose the first end-to-end privacy-preserving framework for DGA classification, ensuring input and output privacy. We summarize our contributions below:

- We propose a novel framework for private classification as a service of domains into DGA/non-DGA, with input privacy (MPC) and output privacy (DP) guarantees.
- Our proposed framework is the first that considers differentially private training of models for DGA classification.
- Our proposed solution works with classifiers based on multilayer perceptrons, 1D convolutional neural networks, and long short-term memory networks. Our MPC protocol for LSTM is novel, efficient, and the first ever implemented in the MP-SPDZ MPC framework [17].
- We evaluate our framework on real datasets—DGArchive and Alexa—for binary and multiclass domain name classification tasks. The binary classification problem distinguishes a domain as benign and malicious. The multiclass classification problem also outputs the DGA family corresponding to the malicious domain.
- We empirically analyze our approach's privacy and utility trade-offs when using MLP, 1D-CNN, and LSTM. We observe that providing output privacy degrades the accuracy by a small amount in our experiments due to the noise introduced to provide DP guarantees. Using our MPC protocols to provide input privacy, on the other hand, does not degrade the utility of the classification model. The fastest model in terms of the runtime is a multilayer perceptron with an inference time equal to 0.07s. The model with the best accuracy is 1D-CNN, achieving an accuracy equal to 93% for an epsilon 5.
- We demonstrate the efficiency of our proposed solution in terms of runtime with 2 or 3 computing parties.
  With quantization, we observe significant improvements in the performance of our MPC protocols, leading to reduced communication rounds and inference time. Our experiments

show a 23% to 42% improvement in inference runtime without affecting accuracy in the 3PC setting (using replicated secret sharing). This demonstrates that we can achieve near real-time secure detection of DGA domains.

We present now how this paper is organized. Section 2 covers related works. Section 3 covers preliminaries related to machine learning, cryptography, and differential privacy. Section 4 describes the proposed framework, security requirements, and threat model. The proposed MPC protocols for embedding, MLP, 1D-CNN, and LSTM and the discussion on post-training float16 quantization are in section 5. Details about the results and experiment are in section 6, which discusses the experiments' Security and Privacy, Utility-Privacy Trade-Off, and runtime metrics.

2 Related works

We now present the related literature and compare our solution to previous ones.

2.1 DGA detection with deep learning

DGA detection methods did not rely on machine learning in the early stages. For instance, Sharifnya et al. [18] developed a technique that identifies hosts with a high volume of failed DNS queries, subsequently adding these hosts to a "suspicious failure matrix." This section lists relevant works regarding DGA detection using deep or machine learning without any privacy guarantees.

Li et al. [19] propose several real-time detection models and frameworks that utilize metadata generated from domains and combine the advantages of a deep neural network model and a lexical features-based model using the ensemble technique. Another work in this line is [20], which uses Helix's architecture that represents DGA as embeddings. The utilization of multilayer perceptron for DGA output detection was also investigated [21, 22].

Huang et al. [23] propose a Helios architecture that uses CNN to detect DGA. Zhou et al. [24] also uses 1D-CNN to detect DGA and enables binary and multiclass analysis of DGA. Berman [25] uses 1D-CNN with a convolutional layer of one-dimensional data to detect DGA. Chen et al. [26] apply 1D-CNN and BiGRU to detect DGA.

Shahzad et al. [27] uses a recurrent neural network to detect DGA outputs on a per-domain basis using the domain name only, with no additional information, which the authors compare to the performance of a DGA classifier based on the following RNN architectures: Unidirectional LSTM network, bidirectional LSTM (Bi-LSTM), and Gated Recurrent Unit (GRU).

Zhang et al. [28] design and implement several DGA classifiers based on machine learning (SVM and RF) and deep learning (CNN, LSTM, and Bi-LSTM) methods. Yang et al. [29] exploits the character-level characteristics of the DGA domain names and proposes a heterogeneous deep neural network framework that includes 1D-CNN and LSTM. LSTM has been used to detect binary DGA and multiclass DGA by the alphanumeric domain name [30, 31]. Tran et al. [32] present a new LSTM algorithm to address the multiclass imbalance problem in DGA-related botnet detection. Strategies with unbalanced datasets are vital for multiclass DGA detection. Balakrishna et al. [33] use an LSTM, which is adapted to predict better if the dataset is unbalanced. Josan et al. [34] use a bidirectional LSTM network for binary DGA and multiclass DGA detection. Other applications that use LSTM to detect DGA are discussed in works [35–38].

Liu et al. [39] combine a convolutional neural network and a bidirectional long short-term memory network to detect DGA. Yun et al. [40] show a method based on natural language

processing and Wasserstein Generative Adversarial Networks and a new way to prevent attackers' evasion of neural network's DGA detection.

Malware detection is relevant for IoT applications (Internet of Things), and DGA detection for this IoT scenario has been investigated [41].

Li et al. [42] make the inference using the Hidden Markov Model (HMM). Koh et al. [43] uses a pre-trained context-sensitive word embedding to classify DGA. Cucchiarelli et al. [44] use the Kullback-Leibner divergence and the Jaccard Index to estimate similarities to detect DGA. Finally, Yilmaz et al. [45] use a method to detect DGA using LSTM and add a GAN (generative adversarial network) to infer previously unknown malicious domains.

Yu et al. [9] comment that simpler architectures are faster in training and inference and are less prone to overfitting. This conclusion is fundamental and was the focus of this work, as we seek lighter architecture toward achieving greater performance in protocols with MPC. Another work from Yu et al. [8] proposes heuristics for automatically labeling domain names monitored in real traffic. This labeled data is essential for improving the accuracy of deep learning models in detecting DGA. Finally, Yu et al. [10] propose a new way to label a large volume of data collected from real traffic as DGA-related and non-DGA-related, allowing models to be trained with large amounts of real traffic.

Sivaguru et al. [46] strengthen DGA detectors against adversarial attacks and evaluate deep learning models and random forests (RFs) to detect DGA using information beyond the domain name.

2.2 Secure multi-party computation for DGA detection

The work of Drichel et al. [16] is the one that is mostly related to our proposal. That work proposes using MPC protocols to classify domains into DGA and non-DGA. They implement their proposals using several different MPC frameworks: PySyft, TF-Encrypted, MP2ML, and SecureQ8 applied in the classifiers Inline [10], NYU [9], ResNet [6], and FANCI. However, despite its pioneering aspect, the work of Drichel et al. [16] still has several deficiencies:

- It uses CNN2D for performing private inference of DNS domains. This adds unnecessary complexity to the classifiers since 1D-CNN is best suited to text classification problems.
- It does not use privacy-preserving embedded layers. In practice, this assumes that *Bob* has to leak the embedding layer to *Alice* for her to perform the embedding operation in the clear. Thus, this solution is not end-to-end private.
- It does not consider LSTM models over MPC. Our work develops the first LSTM implementation available in any existing MPC framework.
- It considers only binary classification. We show how to carry out binary and multiclass DGA predictions.

2.3 Secure multi-party computation for natural language processing

Hao et al. [47] present private inference on transformer BERT based models in a client-server setting. Clients have private inputs, and servers hold proprietary models. One contribution is a customized homomorphic encryption-based method for matrix multiplication.

Adams et al. [48] present the first application of MPC protocols for CNN-based text classification. Their method adapts a CNN2D from the Crypten framework into a 1D-CNN by utilizing two 2D convolutional layers to emulate the behavior of a 1D-CNN. In contrast, our work directly implements a one-dimensional, private embedding layer, resulting in a more efficient

solution. Furthermore, [48] focuses on word-level classifications, while our research emphasizes secure character-level text classification. Lastly, Adams et al. [48] does not address private embeddings. Their approach assumes that *Alice* first converts her text into an embedded vector using a publicly available BERT model before secret-sharing it. This method is infeasible for a model where the embedding is private and part of *Bob's* confidential information. Our work overcomes this limitation by providing protocols and implementation for secure embeddings.

The model SecureNLP [49] has two security protocols for LSTM and RNN in the honest-but-curious model. The difference between SecureNLP and our solution is that we use LSTM for inference with characters, and SecureNLP conducts inference on words. Additionally, SecureNLP does not work with private embedding layers.

Knott et al. [50] offers a comprehensive overview of the Crypten framework, demonstrating its application in text classification, speech recognition, and image classification. In their work, text classification is conducted through a sentiment analysis experiment using a linear layer operating on word embeddings. Our approach differs from theirs in focusing on character-level rather than word-level classification. Moreover, their work does not involve private embeddings or provide an LSTM protocol and implementation.

3 Preliminaries

3.1 Domain generation algorithms

A Domain Generation Algorithm is an algorithm that generates artificial malicious domain names. DGAs play a vital role in malware that relies on network communication between a botmaster and the bots (infected clients) [7, 9, 10]. As illustrated in Fig 2, the key idea is for a botmaster and malware on infected bots to independently run the same DGA with the same seed (e.g. based on the date) to generate the same list of artificial domains. The botmaster subsequently registers one or more of these automatically generated domains, while the malware on the bots attempts to resolve each domain with the DNS.

In Fig 2, for example, the botmaster and the malware on the bots all use the DGA "xxhex" to generate a list of domain names (pseudorandom strings of alphanumerical characters): "xxd80cd4e0.cn", "xxe0d80cd4.kz", etc. Out of these, the botmaster registers "xxe0d80cd4.kz" as the domain related to IP "189.6.29.252". The malware on an infected client will then attempt to resolve each domain with the DNS. In the example in Fig 2, for domain "xxd80cd4e0.cn", the DNS returns a Non-Existent Domain (NXDOMAIN or NXD) error, which indicates the DNS has no listing for "xxd80cd4e0.cn". In contrast, for domain "xxe0d80cd4.kz", the DNS returns IP "189.6.29.252". Finally, the malware connects with IP "189.6.29.252", where the command and control server is registered. Communication then proceeds between the botmaster and the infected machine.

The ability of the malware to dynamically generate and use new domain names prevents ordinary blacklisting from permanently blocking access between the botmaster and infected machines. Indeed, suppose a firewall or intrusion detection service detects and subsequently blocks one of these domains. In that case, the botmaster will use the DGA to create and register a new domain that is not yet blocked, and the malware will then use the same DGA as the botmaster to make a new connection. In this paper, we train and use neural networks that can distinguish such generated domains from real, benign domains. Such machine learning models are used to recognize DGA domains in DNS traffic and mitigate harm [8].

3.2 Neural networks

3.2.1 Multilayer perceptron. A MLP is a machine learning model representing a fully connected neural network architecture. This model has an input layer, an output layer, and

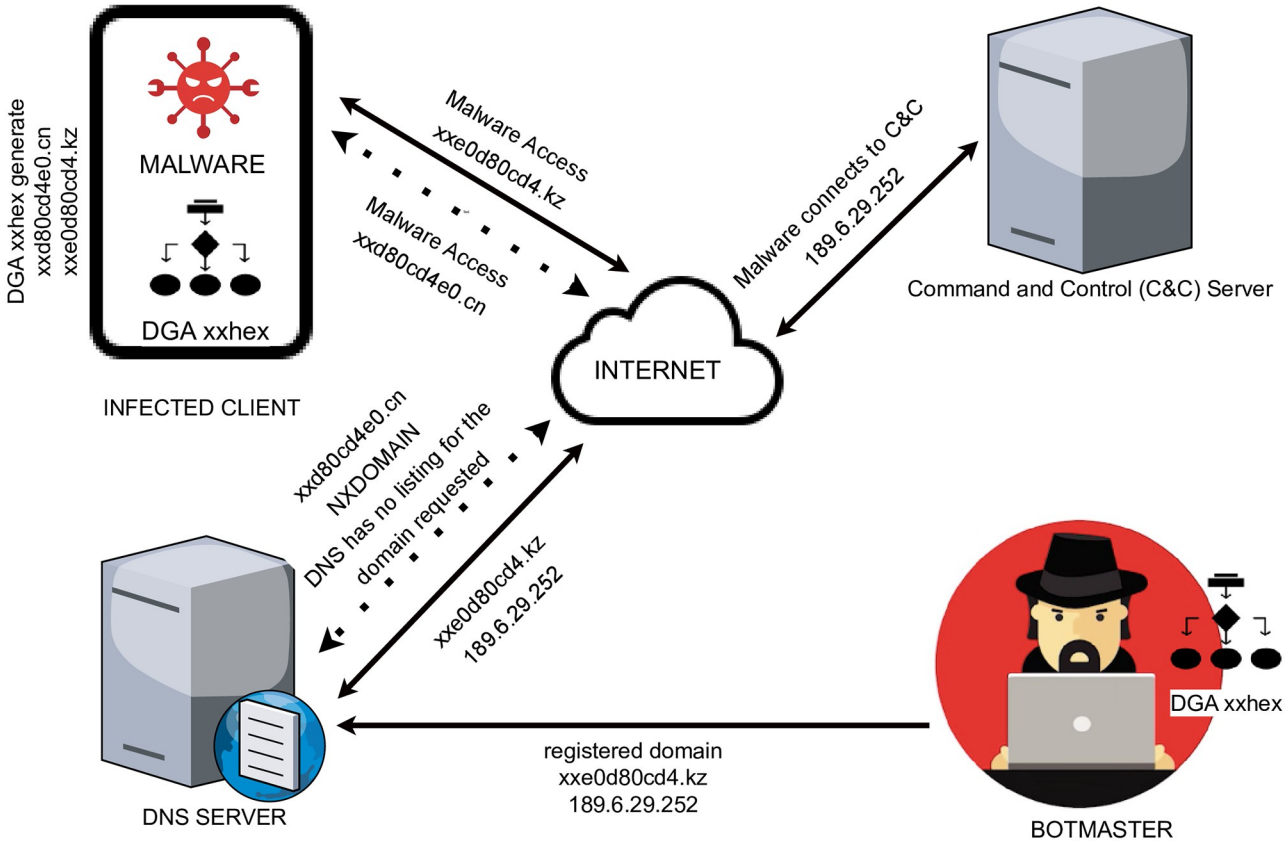

Fig 2. Illustration of the use of a DGA. The botmaster and malware on an infected client generate the same list of domain names. The botmaster registers a domain from the list. The malware attempts to resolve each domain from the list with the DNS until it finds the registered domain and a connection between the infected client and the C&C is successfully established.

https://doi.org/10.1371/journal.pone.0304476.g002

one or more hidden layers. Moreover, the neurons of adjacent layers are connected to each other.

Eq 1 describes an MLP neuron where the output, y, is derived from an activation function, σ. The argument for σ is formed by taking the dot product of the neuron's input vector x with the weight vector w and then adding the bias scalar b.

$$y \leftarrow \sigma(x \cdot w + b) \tag{1}$$

The symbols x and y may be used in different contexts in subsequent sections of this paper. We will explicitly define their meanings each time they are referenced to prevent ambiguity.

3.2.2 Long short-term memory networks. Understanding and abstract reasoning about a given piece of information depends on previous experience. For some tasks, such as reading text, humans can extract knowledge based on the context of previous and recent parts of the text. Standard neural networks do not have memory structures analogous to the human brain. This shortcoming is addressed in recurrent neural networks (RNNs), which closely resemble these memory processes. Therefore, with an RNN, it is possible to train and learn based on a previous element of an input sequence.

With RNNs, it is also possible to model problems where the input data is time-dependent or sequential, i.e., where the next task depends on the previous one. Examples of this can be found in activities such as temperature/weather forecasting, historical air quality trends,

vehicle traffic congestion patterns, etc. Furthermore, in natural language processing, the meaning of texts and language structures humans produce depends on previous texts' content.

Long Short-Term Memory Networks are a kind of RNN designed to resolve problems with vanishing and exploding gradients that occur with long dependencies in the traditional RNN models [51]. LSTMs are well suited for DGA detection [30–32, 34–37, 52, 53], hence we wish to study the effectiveness of LSTMs for the DGA detection problem when privacy is considered.

At their core, LSTM networks revolve around the concept of a cell state, a form of internal memory. LSTM cells have gates that regulate the flow of information into, within, and out of the cell. These gates can add or remove information from the cell state, acting as modification points in the memory system of the network.

We now describe such a cell in more detail and how it is used for inference, i.e., for computing outputs with a trained LSTM. We use the notation and closely follow the explanation presented in [54]. For our specific implementation, the input to the network is a sequence of characters $x_1, \ldots, x_t, \ldots, x_n$. We refer to Section 3.2.4 for describing how each character is converted into a numeric representation. The inputs to the t-th cell consist of the output of the previous cell h_{t-1}, the t-th character x_t, and the state of the previous cell c_{t-1}. The cell outputs h_t, and its state is c_t. We now describe how these quantities are computed. We now describe how these quantities are computed.

We first define how much of the previous state c_{t-1} we will "forget", denoted by the parameter f_t in Eq 2. f_t is computed based on the dot product between the weights w_f and the concatenation of the output of the previous cell h_{t-1} and the cell's input x_t and adding the result to the bias b_f. Weights and biases are learned during training. This result is the input to the *sigmoid* function σ that returns values between 0 and 1. We will see that f_t equal to one means that we keep all of the previous state, while an f_t equal to zero means we forget everything about it when computing the state of the current cell.

$$f_t \leftarrow \sigma(w_f \cdot [h_{t-1}, x_t] + b_f) \tag{2}$$

The state of the t-th cell is computed according to Eq 5. It is a weighted average of the previous cell state c_{t-1}, and a state that depends on the current input x_t, denoted c'_t. The weights are f_t and i_t, a coefficient that tells us how much of the "current" state (c'_t) we want to keep. i_t and c'_t are computed according to Eqs 3 and 4. w_i, w_c are weights and b_i, b_c are biases computed during training. *tanh* is the hyperbolic tangent.

$$i_t \leftarrow \sigma(w_i \cdot [h_{t-1}, x_t] + b_i]) \tag{3}$$

$$c'_t \leftarrow tanh(w_c \cdot [h_{t-1}, x_t] + b_c]) \tag{4}$$

$$c_t \leftarrow f_t \cdot c_{t-1} \cdot i_t \cdot c'_t \tag{5}$$

Finally, the output of the t-th cell h_t is computed according to Eqs 6 and 7. w_o and b_o are weights and biases, respectively.

$$o_t \leftarrow \sigma(w_o \cdot [h_{t-1}, x_t] + b_o]) \tag{6}$$

$$h_t \leftarrow o_t \cdot tanh(c_t) \tag{7}$$

h_t, c_t, and x_{t+1} are then fed into the next cell, and the computations happen similarly for the t + 1-th cell.

3.2.3 Convolutional neural network. A Convolutional Neural Network usually comprises convolution blocks followed by a fully connected network. A standard convolution block consists of a convolution layer, followed by an activation layer, and then by a pooling layer. The standard convolution layer (2D-CNN) takes a 3D input of height h, width w, and depth c and consists of f number of 3D learnable kernels each of size $k \times l \times c$. Each kernel moves along two directions of the input to generate a 3D output. In a 1-dimensional CNN layer, the input is instead a matrix. The kernel moves along only one direction [55].

Let x be the input of a 1D-CNN. Let y be the output of the 1D-CNN layer, and k is the total number of kernels, where the length of y equals $l - k + 1$. The kernel applies a sliding window operation on the input x.

The output of a 1D-CNN can be represented by Eq 8, where $y[i]$ represents the output in the position i. The operation involving x and w is a dot product; b is the bias; w is the weight of the 1D-CNN trained and represents kernels; $w[j]$ is a kernel in the position j.

$$y[i] \leftarrow \sum_{j=0}^{k-1} (x[i+j] \cdot w[j]) + b, \tag{8}$$

3.2.4 Embedding layer. Embedding layers make it possible to represent a text by using a vector of finite precision real numbers. In this work, each letter of a domain name is represented by a vector of finite precision real numbers. Instead of manually specifying the values for the embedding, they are trainable parameters.

We provide the first protocol and implementation for an embedding layer over MPC and implement it in the MP-SPDZ framework. The main idea behind our solution explained in detail in Section 5, is to represent Alice's input to the protocol as a matrix x where each row of x is one hot encoding of a character of the domain name to be classified. So, each row of x consists of a binary vector with Hamming weight equal to 1. A dot product is made between x and the private embedding matrix w, resulting in y, and the result of the embedding layer is the input to 1D-CNN, LSTM, and MLP models.

Our process to get the output of the embedding layer can be described by Eq 9:

$$y \leftarrow x \cdot w, \tag{9}$$

Where the matrix x of dimension $l \times c$ represents one-hot-encoded inputs, the output of the embedding layer is the matrix y of dimension $l \times d$ representing an embedding of the input domain name to be classified, and matrix w of dimension $c \times d$ represents an embedding matrix, where l is the length of input domain, c is the length of the character set, and d is the size of the dimensional vector space.

Drichel et al. [16] proposed a solution for MPC-based private inference of DGA models. Their approach utilized several publicly available MPC frameworks for implementation. At the time, no protocol or implementation for private embedding existed, so the authors assumed that the embedding layer was publicly accessible and implemented in the clear by the party responsible for classifying the domain. However, this assumption could lead to information leakage about the model, especially if the embedding was trained on private data.

3.3 Privacy-enhancing technologies

3.3.1 Differential privacy. Informally, a differentially private algorithm produces a given output with approximately the same probability, regardless of whether a single entry is present or absent in a dataset used to compute the algorithm output. This means that the output is

negligibly affected by the participation of a single user, thereby offering privacy through plausible deniability.

We recall the definition of differential privacy:

Definition 1 *(ϵ, δ)—Differential Privacy* [13]: *A randomized algorithm $\mathcal{M}$ with domain $\mathbb{D}$ is (ϵ, δ)—differentially private if for all $S \subseteq Range(\mathcal{M})$ and for all $x, y \in \mathbb{D}$ such that $||x - y||_1 \leq 1$:*

$$Pr[\mathcal{M}(x) \in S] \leq exp(\epsilon)Pr[\mathcal{M}(y) \in S] + \delta \tag{10}$$

Definition 1 implies that queries on datasets x and y differing on a single entry ($||x - y||_1 \leq 1$) should produce different results with probability bounded by a quantity depending on the parameters ϵ and δ. The constant ϵ is the *privacy budget*. The smaller the value of ϵ, the more privacy the randomized algorithm $\mathcal{M}$ offers. The constant δ captures a small probability of violating the privacy guarantee. When δ equals zero, we say we have pure DP. When $\delta > 0$, we say we have approximate DP. δ is usually heuristically chosen to be smaller or equal to the reciprocal of the dataset size.

Deep learning models often leak information about their training dataset. Moreover, that leak is possible even when only black-box access is available, i.e., when an adversary only observes the output of the deep neural network. Differential privacy is used as a way to prevent such leaks. One can train deep learning models with DP guarantees using DP-SGD (differentially private stochastic gradient descent) [56].

The two essential steps in DP-SGD compared to traditional SGD are gradient clipping and noise addition. Gradient clipping is a technique that limits the magnitude of gradients to a predetermined threshold. Gradient clipping means the model's gradients computed for each data point are scaled down if their magnitude exceeds a certain threshold. This step prevents individual data points from disproportionately impacting the model's learning process, thus reducing the model's sensitivity to any single data point. Once the gradients have been clipped, noise is added to provide privacy. The noise is typically sampled from a Gaussian distribution, with the standard deviation or scale parameter proportional to the chosen privacy budget ϵ. Smaller values of ϵ provide higher privacy but require more significant amounts of noise to be added to the model, which reduces the model's accuracy.

To the best of our knowledge, we are the first to study the trade-off between privacy and the utility of DP-SGD in the context of DGA classification.

3.3.2 Secure multi-party computation. Secure Multi-Party Computation protocols allow mutually distrustful parties to engage in a computation so that, at the end of the protocol, all the honest parties have received the correct output of the computation. No collusion of corrupted parties can learn any information other than what can be inferred from the inputs and outputs of the corrupted parties in the computation. We use a variant of the traditional MPC scenario, where computing servers offer MPC as a service, and the inputs can come from outside parties. These parties do not engage in the MPC protocol. We refer the reader to one of the many available introductions to MPC in the literature [14, 57–59].

The main idea behind all available MPC protocols is to decompose the function to be privately computed into a circuit consisting of addition and multiplication gates. Then, we execute the underlying MPC protocol for evaluating each addition and multiplication gate sequentially till the result is computed and revealed to a designated receiver. Decomposing the functions to be computed into circuits consisting of addition and multiplication gates is a non-trivial task. Efficient representations can dramatically increase the performance of an MPC computation of a given function.

We implement our solutions using MP-SPDZ [17], a publicly available framework for implementing multiple MPC protocols. MP-SPDZ has a high-level interface in Python for presenting

a circuit to be computed over an MPC protocol. MP-SPDZ also has several circuit representations of machine learning algorithms, including multilayer perceptrons and 2D convolutional neural networks. However, no LSTM implementation in MP-SPDZ is available in the literature. Embedding layers are also not available in MP-SPDZ. This work presents novel circuit representations for computing the inference of LSTM networks and embedding layers.

We work with MPC protocols based on secret sharing. Secret sharing is the computational process of splitting a secret s into multiple secret shares x_i and giving these shares to shareholders. Only authorized families of the set of shareholders can recover the secret. Notably, no shareholder can obtain any information on the secret s. Secret sharing-based MPC protocols work by having the inputs to the computation secret shared among all the computing servers. Computations happen on shares rather than on the inputs themselves [14].

3.4 Quantization

The precision used to represent a machine learning model's parameters impacts the model's accuracy, runtime, and size. The performance impact of such precision is magnified when ML models are run on top of secure multi-party protocols. So, it is desirable to quantize (reduce the precision) used in ML models to make them lightweight. We aggressively use post-training float16 quantization to optimize our models [60]. By reducing the precision of our models' weights to 16 bits, we reduce runtimes by approximately 23% to 42% in the 3PC setting (using replicated secret sharing), with a minimal impact on accuracy.

4 Proposed framework, privacy requirements, and threat model

Overview of the proposed framework

We recall that *Alice* holds a DNS domain to be classified, and *Bob* holds a machine learning model that classifies DNS domains into malicious or benign. Our framework consists of set of m untrusted computing servers (MPC servers) $\mathcal{S} = \{S_1, S_2, \ldots, S_m\}$. While our proposed MPC protocols are general and work for any number of servers, we demonstrate our proposed protocols for $m = 2$ and $m = 3$. We assume pairwise authenticated and private communication channels between the servers. The communication between *Alice*, *Bob*, and any server $S_i \in \mathcal{S}$, is also authenticated and private. Our proposed secure classification works as follows:

- Initially, *Alice* and *Bob* convert their real-valued private inputs (domains in the case of *Alice* and the model parameters in the case of *Bob*) into fixed-point representations. They then secret-share their respective fixed-point inputs with the computing servers.
- The computing servers then engage in MPC-based communications and computations to execute the MPC protocols to classify *Alice's* domain names using *Bob's* model.
- The inference result is secret shared among the computing servers at the end of the MPC protocols. The computing servers send their shares to *Alice*. Finally, *Alice* aggregates the secret shares and retrieves the classification result.

Threat model

We build our MPC protocols on existing MPC primitives [17]. Our proposed protocols can be adapted to any threat model by replacing the underlying primitives available in the literature for the given threat model. We demonstrate how our protocols work for the case of "honest-but-curious" servers, which are participants that follow the protocol instructions but try to

obtain as much information as possible about the secret data. We consider threat models: (1) with two and three computing servers; and (2) that withstand attacks by any adversary that can corrupt one server that he chooses, i.e., the privacy guarantees are maintained even when one of the MPC servers is corrupted.

Privacy requirements

Bob cannot learn anything about *Alice's* input. *Alice* should only learn the result of the classification protected by (ϵ, δ)-differential privacy and cannot learn anything else about *Bob's* model and training data. The MPC servers cannot learn anything about *Alice's* and *Bob's* private inputs.

5 Methodology and proposed protocols

5.1 Training DGA classifiers using differential privacy

In our proposed solution, *Alice* is provided with the classification output of her domains. In the binary case, this is a label DGA or non-DGA. In the multiclass case, the output is either a label non-DGA or a label identifying one specific DGA family. This output should preserve the privacy of the individual entries of *Bob's* training data set. We employ the well-known DP-SGD (differentially private stochastic gradient descent) technique [56] to mitigate privacy concerns and provide DP guarantees for *Bob's* training data set according to Definition 1. We specifically train MLP, 1D-CNN, and LSTM models with DP guarantees.

Post-training quantization is an effective technique to improve the inference performance of a trained model. Post-training quantization transforms the model's weights and activations from floating-point precision (32-bit) to lower-precision representations. We propose to quantize the trained model parameters to float16 [15], which causes minimal loss in accuracy and allows for wider deployment of machine learning models with various hardware specifications (e.g. GPU, CPU with float32 or float16 instruction sets). The DP guarantees follow for the post-training quantized model due to the post-processing property of DP. We note that our framework can be adapted to other quantization schemes as well [60–62]. By combining DP and quantization techniques, we can preserve privacy for *Bob's* training data set and enable faster inference.

5.2 Secure inference of DGA domains

During the inference stage, *Alice* has an instance (raw text input, i.e., the domain name string) that needs to be classified as a DGA or non-DGA domain. *Alice* begins by adjusting the length of the given input text to a publicly known value l by truncating the input text or padding the input text with zeros. The fixed length text is then one hot encoded based on ASCII characters, resulting in a matrix x of dimension $l \times 128$, where ASCII is a set of 128 characters. *Alice* then secret shares matrix x and *Bob* secret shares the model parameters with the computing servers. The architecture of *Bob's* model is publicly known.

The computing servers execute MPC protocols that output the secret shares of the classification result to *Alice*. Our proposed framework is equipped to provide inference using models that use MLP, 1D-CNN, and LSTM architectures. All of these models have an embedding layer as the first layer. Next, we describe the proposed MPC protocols for these models that perform secure inference.

Basic building blocks. We build our proposed protocols upon a few building blocks already available in the MP-SPDZ framework [17]. We use MPC protocols π_{SIGMOID} for the sigmoid function, π_{SOFTMAX} for the softmax function, π_{MUL} for secure dot product, π_{TANH} for the

hyperbolic tangent, π_{RELU} for relu function, and π_{DENSE} for dense layer. See [17] for a detailed description of these primitives.

Privacy-preserving embeddings. The embedding layer transforms the high-level input matrix x from *Alice* to provide a dense vector representation of the characters in the input text using *Bob's* learned embedding weights w. Many previous works for MPC-based privacy-preserving text classification, including DGA detection, require *Alice* to embed the input text, which in turn requires the trained embeddings to be made public and may leak information regarding the training data unless trained with DP guarantees [16, 48]. Moreover, these trained embeddings may be proprietary. To mitigate the above scenarios, we propose a novel MPC protocol for the embedding layer to compute the embeddings of private input text in an oblivious manner. The vectors resulting from the embedding layer of our classification models represent the lexical information of the characters in the given DGA domain (URL) in the ASCII character set. The idea behind $\pi_{\text{EMBEDDING}}$ is simple: we extract the vector representation of each character in the input text (in our case, it is a domain or URL) from the trained embeddings with DP guarantees.

One of the simplest ways to extract such embeddings (Protocol 1) is to represent the input text as one hot encoded matrix x of dimension $l \times 128$ and multiply it with the weights of trained embeddings w of dimension 128×128. The product is a matrix y of dimension $l \times 128$ representing a set of vector representations of each character in the input text. We note that feature extraction done this way requires only multiplication operations for which state-of-the-art MPC primitives are available, which results in optimized performance of the MPC protocol for extracting embeddings of the input. Moreover, our protocol is general enough to work with character sets of arbitrary cardinality c.

Protocol 1: $\pi_{\text{EMBEDDING}}$ for secure inference of embedding layer

Input: Secret shared matrices x of dimension $(l \times c)$ representing one-hot-encoded inputs and w of dimension $(c \times d)$ representing embedding weights, where l is the length of the input text in characters, c is the cardinality of the character set, and d is the dimensionality of the embedding space.

Output: A secret shared embedding matrix y of dimension $(l \times d)$ of the input domain to be classified.

1 $y \leftarrow \pi_{\text{MUL}}(x, w)$

2 **return** y

Privacy-preserving MLP. In our work, we leverage an existing implementation of an MPC protocol for secure inference with an MLP, available within the MP-SPDZ framework [17]. MLPs will be used as a baseline method in our framework.

The input is secret shared by *Alice*, while *Bob* secret-shares the model's weights. Bob's model in the architecture with MLP composes the weights of the embedding and dense layers.

Privacy-preserving 1-D convolution. We now present our solution based on 1D-CNN. Our architecture has an input layer with the protocol $\pi_{\text{EMBEDDING}}$, a layer with the protocol $\pi_{\text{1D-CNN}}$, and a layer with protocol π_{DENSE} representing the dense layer in MPC. The input will be secret shared by *Alice*, while *Bob* secret-shares the model's weights.

We leverage an existing proposal for a 1D-CNN [48] but provide our implementation [63]. The protocol $\pi_{\text{1D-CNN}}$ for secure inference with 1D-CNN is built upon (a) the existing MPC protocols available in the literature for π_{RELU}, and π_{MUL} (b) our proposed MPC protocols for embedding $\pi_{\text{EMBEDDING}}$ and 1-D convolution $\pi_{\text{1D-CNN}}$.

Protocol 2 for secure inference with a 1-D convolutional layer takes as input (1) the secret shared embeddings obtained as output from Protocol $\pi_{\text{EMBEDDING}}$ and (2) the secret shared model parameters, i.e. weights of the kernels (k in total) each of size k for the 1-D convolution layer from *Bob*.

Protocol 2: π_{1D-CNN} for secure inference with 1-D Convolution.

Input: The secret shared matrices x (obtained as the output of the private embedding computed by Protocol $\pi_{EMBEDDING}$, b, and w representing input, bias, and weight. The constants k and l represent the number of kernels and rows of x.

Output: A secret shared y of dimension $l - k + 1$.

for $i \leftarrow 0$ ***to*** $l - k + 1$ **do**
  $y[i] \leftarrow b$
  for $j \leftarrow 0$ ***to*** $k - 1$ **do**
    $y[i] \leftarrow y[i] + \pi_{MUL}(x[i + j], w[j])$
  end
end
return $\pi_{RELU}(y)$

Privacy-preserving LSTM protocol. We now present our solution based on LSTM. Our architecture has an input layer with the protocol $\pi_{EMBEDDING}$, a layer with the protocol π_{LSTM}, and a layer with protocol π_{DENSE} representing the dense layer in MPC. The input will be secret shared by *Alice*, while *Bob* secret-shares the model's weights.

Protocol 3 describes the LSTM layer. The input for the LSTM layer is the secret shared output of the embedding layer, while *Bob* secret shares the kernel weights (w_f, w_i, w_o, w_c) and the biases (b_f, b_i, b_o, b_c). We refer the reader to Section 3.2.2 for an explanation of each one of these terms. The operations involve MPC protocols $\pi_{SIGMOID}$ for sigmoid, π_{MUL} for secure multiplications, and π_{TANH} for the hyperbolic tangent.

Protocol 3: π_{LSTM} for secure LSTM

Input: Secret shared vector x (obtained as output of the private embedding computed by Protocol $\pi_{EMBEDDING}$), and secret shared values of the weights (w_f, w_i, w_o, w_c), and biases (b_f, b_i, b_o, b_c). The input length is publicly known. Let [a, b] denote the concatenation of a and b.

Output: A secret shared vector y containing the output after the LSTM layer of the inference process.

$h_0 \leftarrow 0$
$c_0 \leftarrow 0$
for $t \leftarrow 1$ ***to*** n **do**
  $f_t \leftarrow \pi_{SIGMOID}(\pi_{MUL}(w_f, [h_{t-1}, x[t]]) + b_f)$
  $i_t \leftarrow \pi_{SIGMOID}(\pi_{MUL}(w_i, [h_{t-1}, x[t]]) + b_i)$
  $c'_t \leftarrow \pi_{TANH}(\pi_{MUL}(w_c, [h_{t-1}, x[t]]) + b_c))$
  $c_t \leftarrow \pi_{MUL}(f_t, \pi_{MUL}(c_{t-1}, \pi_{MUL}(i_t, c'_t)))$
  $o_t \leftarrow \pi_{SIGMOID}(\pi_{MUL}(w_o, [h_{t-1}, x[t]]) + b_o)$
  $h_t \leftarrow \pi_{MUL}(o_t, \pi_{TANH}(c_t))$
  $y[t] \leftarrow h_t$
end
return y

6 Results

6.1 Dataset

The DGA dataset was obtained from DGArchive [64]. The dataset from DGArchive was the set containing all collected DGA examples up to 2019. We remove DGA families with low representation from the dataset (less than 30k samples). We end up with 1,000,000 examples from the DGArchive dataset for use as positive matches from the following families: bamital, banjori, bedep, beebone, blackhole, bobax, conficker, corebot, cryptolocker, darkshell, dircrypt, dnsbenchmark, dnschanger, downloader, dyre, ekforward, emotet, feodo, fobber, gameover, gameover_p2p.csv, gozi, gspy, hesperbot, locky, madmax, matsnu, modpack, murofet, murofetweekly, necurs, nymaim, oderoor, padcrypt, proslikefan, pushdo, pushdotid, pykspa,

pykspa2, pykspa2s, qadars, qakbot, ramdo, ramnit, ranbyus, randomloader, redyms, rovnix, shifu, simda, sisron, suppobox, sutra, symmi, szribi, tempedreve, tinba, torpig, tsifiri, urlzone, vawtrak, virut, volatilecedar, xxhex.

For negative DGA matches, we have acquired approximately 1,000,000 domains from the last known version of the dataset "Alexa top 1 million domains" [65], which are used for training the model to recognize legitimate domains.

All data is shuffled and split into 80% for training and 20% for testing for binary and multiclass model architectures.

We convert alphanumeric characters representing domain names to lowercase for use in the model. Then, we convert each character to its corresponding ASCII code, which lies between the values of 0 to 127. Finally, the maximum length of a domain ASCII string is set to 64 characters. For domains whose size is less than 64, we prepend zero padding length to 64.

In this section, we evaluate our experimental results. First, we compare the 1D-CNN, LSTM, and MLP models using secure MPC when applied to binary and multiclass DGA detection with and without differential privacy. We also experimented with quantizing the models after the DP training (in which case, after training with DP, we applied quantization to reduce the weights to 16 bits and thus achieve performance gains in the inference phase with MPC).

6.2 Model architectures and parameters

All trained models have an embedding Layer as the first layer, where the input is a vector of 64 numerical elements resulting from transforming each character into ASCII code. The result of embedding is a 128 by 128 dimension array.

The last layer in all binary models comprises one dense layer with one neuron, activation function sigmoid, loss function binary cross-entropy, and optimizer Adam.

The last layer in all multiclass models comprises one dense layer with 65 neurons representing all families, the activation function softmax, the loss function sparse categorical cross-entropy, and the optimizer Adam with a learning rate of 0.001, batch size of 64, and 30 epochs.

We now provide further details about the other layers for each one of the architectures [63] we used:

- **MLP:** The binary and multiclass MLP models have a flatten layer to transform the data resulting from the embedding layer into one dimension data followed by a dense layer with 100 neurons, ReLU activation function and dropout with rate 0.1. The MLP binary model has 835,785 parameters, while the MLP multiclass model has 842,249 parameters.
- **1D-CNN:** The binary and multiclass 1D-CNN models have: (1) a 1D-CNN with filters = 32, kernels = 2, ReLU activation function and dropout with rate 0.1; (2) followed by a flatten layer to transform the output of the previous layer into one-dimensional data; (3) a dense Layer with 100 neurons, ReLU activation function, dropout with rate 0.1. The 1D-CNN binary model has 226,409 parameters, while the 1D-CNN multiclass model has 232,671 parameters.
- **LSTM:** The binary and multiclass LSTM models have: (1) a LSTM with 32 units with ReLU activation function and dropout with rate 0.1; (2) followed by a flatten layer to transform the output of the previous layer into one-dimensional data; (3) a dense layer with 100 neurons, ReLU activation function and dropout with rate 0.1. The LSTM binary model has 48,713 parameters, while the LSTM multiclass model has 55,177.

The DP-SGD parameters used in our experiments are available in TensorFlow Privacy [66] as follows: delta is 6.188×10^{-7}, the clipping norm is 1, and the number of microbatches is 1.

Table 1. Results on the DGA inference accuracy for different noise levels with and without quantization.

Model	Non-Quantized, ϵ =				Quantized, ϵ =			
	0.1	**2**	**5**	∞	**0.1**	**2**	**5**	∞
1D-CNN binary	90%	93%	93%	99%	90%	93%	93%	99%
1D-CNN multiclass	25%	47%	53%	88%	25%	47%	53%	88%
LSTM binary	88%	91%	92%	97%	88%	91%	92%	97%
LSTM multiclass	23%	46%	51%	88%	23%	46%	51%	88%
MLP binary	90%	93%	93%	96%	90%	93%	93%	96%
MLP multiclass	24%	46%	51%	87%	24%	46%	51%	87%

https://doi.org/10.1371/journal.pone.0304476.t001

Regarding the noise multipliers, for an epsilon of 0.1, it is 2.51; for an epsilon of 2, it is 0.61; and for an epsilon of 5, it is 0.46.

6.3 Experiments

Utility-privacy trade-off. Table 1 provides a comprehensive analysis of the trade-off between utility (as measured by the accuracy of the model) and privacy as represented by privacy budget, denoted as ϵ, of the differential privacy mechanism) for all models.

The privacy budget ϵ is a quantifiable measure of privacy. As a rule of thumb, an ϵ value less than one indicates high privacy, a value between one and two suggests moderate privacy, while an ϵ exceeding two is considered low privacy. ϵ equal to infinity represents a model without any differential privacy guarantee. As expected, the model's accuracy decreases as the privacy budget reduces, thus entering a higher privacy regime. The decrease is more severe for the multiclass case. This observation captures the trade-off between utility and privacy when applying differential privacy. Table 1 also presents the accuracy levels of each model when quantization is applied. The quantization step did not change the accuracy, even though quantized models are significantly faster than their non-quantized counterparts.

Runtimes. We implemented the MPC-based inference both in the scenario of two computing servers (2PC) as well as three computing servers (3PC) connected over a Gigabit Ethernet local area network. In the inference experiments, we used three Azure instances with 32 Cores of Intel(R) Xeon(R) Platinum 8272CL CPU @ 2.60GHz and 64GB of RAM. We used the following underlying MPC protocols available on MP-SPDZ [17] in our experiments: semi2k for 2PC and replicated secret sharing for 3PC. The runtimes are available in Table 2 and include communication and computation delays. The MLP model had the best runtimes, but 1D-CNN had the better accuracy, especially for higher values of ϵ (see Table 1).

We note that runtimes are unaffected if the model is protected by differential privacy. Note also that runtimes strongly depend on the corruption threshold: the three party protocols (i.e., honest majority) are faster than the two-party protocols, in which there is no honest majority.

We also performed the same experiments regarding MPC-based inference but using the models we quantized after the DP training. The runtimes are presented in Table 3.

Note that in the 3PC setting (using replicated secret sharing), the quantization step reduced inference runtime by approximately 23% to 42%. The reduction in execution time for the 2PC configuration (using semi2k) was 2% for MLP Binary, 4% for MLP Multiclass, 1% for 1D-CNN, and 42% for LSTM.

Comparison to previous results. The closest result to ours is the work of Drichel et al. [16]. There, they also propose a framework and protocols for private inference and classification of DNS traffic. Their runtimes and accuracy cannot directly compare to ours because of the weaker privacy guarantees they achieve (as pointed out in the related works section).

Table 2. Inference using MP-SPDZ protocol.

Model	Setting	Inference Time (sec)	Rounds	Data Sent (MB)
MLP Binary	3PC	0.0778787	2773	17.217
1D-CNN Binary	3PC	0.319441	8551	21.6062
LSTM Binary	3PC	10.4153	195131	1485.53
MLP Multiclass	3PC	0.133239	3615	24.8676
1D-CNN Multiclass	3PC	0.359278	9382	29.0157
LSTM Multiclass	3PC	10.5449	197123	1489.92
MLP Binary	2PC	14.2051	42923	3951.16
1D-CNN Binary	2PC	14.2954	54151	3983.37
LSTM Binary	2PC	103.472	441023	26752.6
MLP Multiclass	2PC	14.577	44599	4054.73
1D-CNN Multiclass	2PC	14.6503	55831	4089.34
LSTM Multiclass	2PC	104.077	443895	26820.1

https://doi.org/10.1371/journal.pone.0304476.t002

Drichel et al. [16] leak information about the model to Bob (the embeddings). Moreover, their solution does not offer output privacy since there is no differential privacy training of the model. Despite these weaker guarantees, their solution is not faster than ours. All of their models have inference time higher than 1s, and their experiments do not consider network delays. In comparison, our most accurate classifier (CNN 1D) runs under 0.4 s, including network delays. Regarding communication complexity, all the models presented in [16] exchange over 190MB in messages. Our fastest protocol exchanges 17MB (with quantization), and our most accurate one exchanges 21MB (with quantization).

6.4 Security and privacy

Input privacy: The underlying MPC protocols from MP-SPDZ [17] that we use in our solutions (replicated secret sharing in the case of 3 computing servers with an honest majority; and the semi2k protocol in the case of 2 computing servers without an honest majority) implement a secure arithmetic black-box MPC. They only perform operations over secret shares, and no information is leaked during the computation over the secret shares. Moreover, we use MPC sub-protocols π_{SIGMOID} for the sigmoid function, π_{SOFTMAX} for the softmax function, π_{MUL} for secure dot product, π_{TANH} for the hyperbolic tangent, π_{RELU} for relu function, and π_{DENSE} for dense layer from MP-SPDZ. All these sub-protocols do not leak any information and are

Table 3. Inference using MP-SPDZ protocol with quantization after DP training.

Model	Setting	Inference Time (sec)	Rounds	Data Sent (MB)
MLP Binary	3PC	0.0593449	2469	12.3028
1D-CNN Binary	3PC	0.223521	7371	12.3982
LSTM Binary	3PC	5.9959	137044	545.121
MLP Multiclass	3PC	0.076885	3038	15.0529
1D-CNN Multiclass	3PC	0.260958	7940	15.1483
LSTM Multiclass	3PC	6.15767	138573	546.706
MLP Binary	2PC	13.813	42723	3845.97
1D-CNN Binary	2PC	14.1394	51795	3937.96
LSTM Binary	2PC	60.3208	310659	15352.2
MLP Multiclass	2PC	13.9138	44599	3983.37
1D-CNN Multiclass	2PC	14.4177	54151	4054.73
LSTM Multiclass	2PC	60.247	312951	15388.2

https://doi.org/10.1371/journal.pone.0304476.t003

universally composable (UC) secure. The novel protocols that we propose ($\pi_{\text{EMBEDDING}}$, $\pi_{\text{1D-CNN}}$ and π_{LSTM}) therefore do not leak any information to the computing servers responsible for the MPC operations over the secret shares about *Alice's* or *Bob's* inputs; nor any private information is leaked to Alice or Bob other the result of the classification (which Alice can reconstruct by getting all shares of the output from the computing server). Our protocols UC-securely implement the ideal functionality $\mathcal{F}_{PPCDGA}$ for privacy-preserving classification of domains as DGA or non-DGA that is described below for the case of classification using MLP, 1D-CNN or LSTM models.

Functionality $\mathcal{F}_{PPCDGA}$

$\mathcal{F}_{PPCDGA}$ waits until it receives as input from Alice her domain name x and as input from Bob his model m for DGA classification that he trained with DP-guarantees.

Upon receiving both inputs, $\mathcal{F}_{PPCDGA}$ locally computes the result y of classifying x using model m and sends y as a public delayed output to Alice.

Output privacy: The guarantee that *Alice* does not learn information about individual entries used in the training of *Bob's* model—is provided by the DP guarantees of DP-SGD [56] as utilized by *Bob*.

In a nutshell, DP-SGD guarantees (ϵ, δ) – differential privacy by repeated applications of the Gaussian mechanism. When training the model, Bob executes the following steps:

- Bob initializes the model M with random parameters.
- For each input x_i in a batch B, Bob computes the output of the model $M(x_i)$ and, using the corresponding label y_i associated with x_i, computes the gradient of the loss function, denoted by g_i.
- for each gradient g_i in the batch B, Bob clip the gradient if its L2 norm is larger or equal to a threshold c.
- Bob averages all the clipped gradients in Batch B and adds (to each coordinate of the gradient) a random variable sampled from a Gaussian distribution with mean 0 and a standard deviation $\sigma.c$, where c is the clipping threshold, and *sigma* is an appropriately chosen constant that ensures differential privacy.
- Bob then updates the model parameters using these noisy aggregated gradients times the learning rate.
- This process is repeated for each batch in the dataset.

The post-processing property of differential privacy ensures that the outputs of the model are also differentially private.

7 Conclusions and future work

This work presents the first framework for performing outsourced privacy-preserving DGA detection with input and output privacy. Input privacy means that no information about the domain being classified leaks to the computing servers or the model owner *Bob*, and no information about the model leaks to the domain owner (*Alice*) or the computing servers. Output

privacy means that the output of the computation does not allow the extraction of individual entries in the training dataset of the ML model used in the framework.

We additionally proposed MPC protocols for LSTM and embedding layer in the MP-SPDZ framework, furthering the practical application of private and secure machine learning.

We compared the performance of CNN, LSTM, and MLP trained with DP and secure inference with various MPC protocols. The 1D-CNN presented better cost-benefit accuracy and performance on MPC for two parties, and in the case of the three computing servers, there were trade-offs between CNN and MLP.

We demonstrated that post-training float16 quantization with DP improved our runtimes by approximately 23% to 42% without a significant drop in accuracy when we were in the setting of three parties with an honest majority.

In a future work, we propose to apply the techniques of this paper to general tasks of natural language classification.

Limitations and societal impacts

Our work only considers the actual performance against honest-but-curious adversaries. Considering fully malicious adversaries would increase our inference times, particularly for a dishonest majority setting. Developing protocols that perform well even for the malicious behavior case is left as an open problem.

In our adversarial model, we assume players input correct information into the protocol. Nothing in our solution prevents the players from presenting adversarial inputs as well as out of distribution inputs, in the case of Bob.

In terms of societal impacts, our solutions imply in a bit higher energy consumption, since, in order to achieve privacy, we have multiple servers performing the inference, rather than a single central server.

We have also not checked if differential privacy affects underrepresented types of domains in our data sets in a disproportional way. We leave that as an open problem, too.

Acknowledgments

The authors thank Marcel Keller for making the MP-SPDZ framework available and for his assistance. The authors thank Microsoft for the generous donation of cloud computing credits through the UW Azure Cloud Computing Credits for Research program.

Author Contributions

Conceptualization: Ricardo J. M. Maia, Anderson C. A. Nascimento.

Data curation: Ricardo J. M. Maia, Dustin Ray, Sikha Pentyala, Rafael Dowsley, Martine De Cock, Anderson C. A. Nascimento.

Formal analysis: Ricardo J. M. Maia, Sikha Pentyala, Rafael Dowsley, Martine De Cock, Anderson C. A. Nascimento.

Investigation: Ricardo J. M. Maia, Dustin Ray, Sikha Pentyala, Rafael Dowsley, Martine De Cock, Anderson C. A. Nascimento.

Methodology: Ricardo J. M. Maia, Sikha Pentyala, Rafael Dowsley, Martine De Cock, Anderson C. A. Nascimento, Ricardo Jacobi.

Project administration: Ricardo J. M. Maia, Anderson C. A. Nascimento.

Resources: Ricardo J. M. Maia, Sikha Pentyala, Martine De Cock, Anderson C. A. Nascimento.

Software: Ricardo J. M. Maia, Dustin Ray, Sikha Pentyala, Anderson C. A. Nascimento.

Supervision: Ricardo J. M. Maia, Martine De Cock, Anderson C. A. Nascimento, Ricardo Jacobi.

Validation: Ricardo J. M. Maia, Dustin Ray, Sikha Pentyala, Rafael Dowsley, Martine De Cock, Anderson C. A. Nascimento, Ricardo Jacobi.

Visualization: Ricardo J. M. Maia, Dustin Ray, Sikha Pentyala, Rafael Dowsley, Martine De Cock, Anderson C. A. Nascimento, Ricardo Jacobi.

Writing – original draft: Ricardo J. M. Maia, Dustin Ray, Sikha Pentyala, Rafael Dowsley, Martine De Cock, Anderson C. A. Nascimento, Ricardo Jacobi.

Writing – review & editing: Ricardo J. M. Maia, Dustin Ray, Sikha Pentyala, Rafael Dowsley, Martine De Cock, Anderson C. A. Nascimento, Ricardo Jacobi.

References

1. ENISA Threat Landscape—The year in review—ENISA Available at: https://www.enisa.europa.eu/publications/year-in-review/view/++widget++form.widgets.fullReport/@@download/ETL2020+-+A+year+in+review+A4.pdf. 2020 Accessed on 09/09/2023.
2. McAfee. McAfee Labs Threats Report, April 2021. Available at: https://www.mcafee.com/enterprise/en-us/assets/reports/rp-quarterly-threats-apr-2021.pdf. April 2021. Accessed on 09/09/2023.
3. DBIR—Data Breach Investigations Report. Available at: https://enterprise.verizon.com/resources/reports/2021-data-breach-investigations-report.pdf. 2021. Accessed on 09/09/2023.
4. ENISA. ENISA ETL2020—Malware. Available at: https://www.enisa.europa.eu/topics/threat-risk-management/threats-and-trends/etl-review-folder/etl-2020-malware. 2020. Accessed on 09/09/2023.
5. Patsakis Constantinos and Casino Fran. Exploiting statistical and structural features for the detection of Domain Generation Algorithms. *Journal of Information Security and Applications*, 58:102725, 2021. ISSN 2214-2126. https://doi.org/10.1016/j.jisa.2020.102725. Available at: https://www.sciencedirect.com/science/article/pii/S2214212620308632.
6. Arthur Drichel, Ulrike Meyer, Samuel Schüppen, and Dominik Teubert. Analyzing the Real-World Applicability of DGA Classifiers. In *Proceedings of the 15th International Conference on Availability, Reliability and Security (ARES'20)*, article no. 15, 11 pages, Virtual Event, Ireland, 2020. Association for Computing Machinery, New York, NY, USA. ISBN 9781450388337. https://doi.org/10.1145/3407023.3407030.
7. Mayana Pereira, Shaun Coleman, Bin Yu, Martine DeCock, e Anderson Nascimento. Dictionary Extraction and Detection of Algorithmically Generated Domain Names in Passive DNS Traffic. In *Research in Attacks, Intrusions, and Defenses*, páginas 295–314, Springer International Publishing, Cham, 2018.
8. Yu B., Pan J., Gray D., Hu J., Choudhary C., Nascimento A. C. A., et al. Weakly Supervised Deep Learning for the Detection of Domain Generation Algorithms. In *IEEE Access*, volume 7, pages 51542–51556, 2019. https://doi.org/10.1109/ACCESS.2019.2911522
9. B. Yu, J. Pan, J. Hu, A. Nascimento, and M. De Cock. Character Level based Detection of DGA Domain Names. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2018. https://doi.org/10.1109/IJCNN.2018.8489147.
10. B. Yu, D. L. Gray, J. Pan, M. D. Cock, and A. C. A. Nascimento. Inline DGA Detection with Deep Networks. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 683–692, 2017. https://doi.org/10.1109/ICDMW.2017.96.
11. Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *USENIX 2019*, pages 267–284, 2019.
12. Congzheng Song, Thomas Ristenpart, and Vitaly Shmatikov. Machine learning models that remember too much. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 587–601, 2017. Available at: https://dl.acm.org/doi/pdf/10.1145/3133956.3134077.
13. Dwork Cynthia and Roth Aaron. The algorithmic foundations of Differential Privacy. Foundations and Trends® in Theoretical Computer Science, 9(3-4):211–407, 2013. https://doi.org/10.1561/0400000042

14. Cramer Ronald, Damgård Ivan Bjerre, and Jesper Buus Nielsen. *Secure Multiparty Computation and Secret Sharing*. Cambridge University Press, Cambridge, 2015. https://doi.org/10.1017/CBO9781107337756

15. TensorFlow Authors. Post-training quantization. Available at: https://www.tensorflow.org/lite/performance/post_training_quantization/. Accessed: 2023-07-04.

16. Arthur Drichel, Mehdi Akbari Gurabi, Tim Amelung, and Ulrike Meyer. Towards Privacy-Preserving Classification-as-a-Service for DGA Detection. In *2021 18th International Conference on Privacy, Security and Trust (PST)*, pages 1–10, 2021. https://doi.org/10.1109/PST52912.2021.9647755.

17. Marcel Keller. MP-SPDZ: A Versatile Framework for Multi-Party Computation. Cryptology ePrint Archive, Report 2020/521, 2020. Available at: https://eprint.iacr.org/2020/521.

18. R. Sharifnya and M. Abadi. A novel reputation system to detect DGA-based botnets. In *ICCKE 2013*, pages 417–423, 2013. https://doi.org/10.1109/ICCKE.2013.6682860.

19. Shuaiji Li, Tao Huang, Zhiwei Qin, Fanfang Zhang, and Yinhong Chang. Domain Generation Algorithms Detection through Deep Neural Network and Ensemble. In *Companion Proceedings of The 2019 World Wide Web Conference (WWW'19)*, pages 189–196, San Francisco, USA, 2019. Association for Computing Machinery, New York, NY, USA. ISBN 9781450366755. https://doi.org/10.1145/3308558.3316498.

20. Lior Sidi, Yisroel Mirsky, Asaf Nadler, Yuval Elovici, and Asaf Shabtai. Helix: DGA Domain Embeddings for Tracking and Exploring Botnets. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM'20)*, pages 2741–2748, Virtual Event, Ireland, 2020. Association for Computing Machinery, New York, NY, USA. ISBN 9781450368599. https://doi.org/10.1145/3340531.3416022.

21. M. I. Ashiq, P. Bhowmick, M. S. Hossain, and H. S. Narman. Domain Flux-based DGA Botnet Detection Using Feedforward Neural Network. In *MILCOM 2019—2019 IEEE Military Communications Conference (MILCOM)*, pages 1–6, 2019. https://doi.org/10.1109/MILCOM47813.2019.9020730.

22. Mao Jian, Zhang Jiemin, Tang Zhi, and Gu Zhiling. DNS anti-attack machine learning model for DGA domain name detection. *Physical Communication*, 40:101069, 2020. ISSN 1874-4907. https://doi.org/10.1016/j.phycom.2020.101069. Available at: https://www.sciencedirect.com/science/article/pii/S1874490719309036.

23. J. Huang, P. Wang, T. Zang, Q. Qiang, Y. Wang, and M. Yu. Detecting Domain Generation Algorithms with Convolutional Neural Language Models. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*, pages 1360–1367, 2018. https://doi.org/10.1109/TrustCom/BigDataSE.2018.00188.

24. S. Zhou, L. Lin, J. Yuan, F. Wang, Z. Ling, and J. Cui. CNN-based DGA Detection with High Coverage. In *2019 IEEE International Conference on Intelligence and Security Informatics (ISI)*, pages 62–67, 2019.

25. Berman Daniel S. DGA CapsNet: 1D Application of Capsule Networks to DGA Detection. *Information*, 10(5):157, 2019. ISSN 2078-2489. Available at: https://www.mdpi.com/2078-2489/10/5/157.

26. ChaoQuan Chen, LeiLei Pan, and XiaoLan Xie. DGA Domain Name Detection Based on BiGRU-MCNN. In *Proceedings of the 2019 4th International Conference on Intelligent Information Processing (ICIIP 2019)*, pages 315–319, China, China, 2019. Association for Computing Machinery, New York, NY, USA. ISBN 9781450361910. https://doi.org/10.1145/3378065.3378126.

27. H. Shahzad, A. R. Sattar, and J. Skandaraniyam. DGA Domain Detection using Deep Learning. In *2021 IEEE 5th International Conference on Cryptography, Security and Privacy (CSP)*, pages 139–143, 2021. https://doi.org/10.1109/CSP51677.2021.9357591.

28. Y. Zhang. Automatic Algorithmically Generated Domain Detection with Deep Learning Methods. In *2020 IEEE 3rd International Conference on Automation, Electronics and Electrical Engineering (AUTEEE)*, pages 463–469, 2020.

29. Yang L., Liu G., Dai Y., Wang J., and Zhai J. Detecting Stealthy Domain Generation Algorithms Using Heterogeneous Deep Neural Network Framework. *IEEE Access*, 8:82876–82889, 2020. https://doi.org/10.1109/ACCESS.2020.2988877

30. P. Vij, S. Nikam, and A. Bhatia. Detection of Algorithmically Generated Domain Names using LSTM. In *2020 International Conference on COMmunication Systems NETworkS (COMSNETS)*, pages 1–6, 2020.

31. S. Akarsh, S. Sriram, P. Poornachandran, V. K. Menon, and K. P. Soman. Deep Learning Framework for Domain Generation Algorithms Prediction Using Long Short-term Memory. In *2019 5th International Conference on Advanced Computing Communication Systems (ICACCS)*, pages 666–671, 2019.

32. Tran Duc, Mac Hieu, Tong Van, Tran Hai Anh, and Nguyen Linh Giang. A LSTM based framework for handling multiclass imbalance in DGA botnet detection. Neurocomputing, 275:2401–2413., 2018.

ISSN 0925-2312. https://doi.org/10.1016/j.neucom.2017.11.018. Available at: https://www.sciencedirect.com/science/article/pii/S0925231217317320.

33. Simran K, Prathiksha Balakrishna, Vinayakumar Ravi, and Soman KP. Deep Learning based Frameworks for Handling Imbalance in DGA, Email, and URL Data Analysis. 2020. arXiv preprint arXiv:2004.04812. Available at: https://arxiv.org/abs/2004.04812.

34. Gurpreet Josan and Jagroop Kaur. LSTM Network based Malicious Domain Name Detection. 2019.

35. Qiao Yanchen, Zhang Bin, Zhang Weizhe, Sangaiah Arun Kumar, and Wu Hualong. DGA Domain Name Classification Method Based on Long Short-Term Memory with Attention Mechanism. *Applied Sciences*, 9(20):4205, 2019. ISSN 2076-3417. Available at: https://www.mdpi.com/2076-3417/9/20/4205.

36. Ryan R. Curtin, Andrew B. Gardner, Slawomir Grzonkowski, Alexey Kleymenov, and Alejandro Mosquera. Detecting DGA Domains with Recurrent Neural Networks and Side Information. In *Proceedings of the 14th International Conference on Availability, Reliability and Security (ARES'19)*, article no. 20, 10 pages, Canterbury, CA, United Kingdom, 2019. Association for Computing Machinery, New York, NY, USA. ISBN 9781450371643. https://doi.org/10.1145/3339252.3339258.

37. Jonathan Woodbridge, Hyrum S. Anderson, Anjum Ahuja, and Daniel Grant. Predicting Domain Generation Algorithms with Long Short-Term Memory Networks. 2016. arXiv preprint arXiv:1611.00791. Available at: https://arxiv.org/abs/1611.00791.

38. S. Kumar and A. Bhatia. Detecting Domain Generation Algorithms to prevent DDoS attacks using Deep Learning. In *2019 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, pages 1–4, 2019. https://doi.org/10.1109/ANTS47819.2019.9118156.

39. Liu Zhanghui, Zhang Yudong, Chen Yuzhong, Fan Xinwen, and Dong Chen. Detection of Algorithmically Generated Domain Names Using the Recurrent Convolutional Neural Network with Spatial Pyramid Pooling. *Entropy*, 22(9):1058, 2020. ISSN 1099-4300. Available at: https://www.mdpi.com/1099-4300/22/9/1058.

40. Yun X., Huang J., Wang Y., Zang T., Zhou Y., and Zhang Y. Khaos: An Adversarial Neural Network DGA With High Anti-Detection Ability. *IEEE Transactions on Information Forensics and Security*, 15:2225–2240, 2020. https://doi.org/10.1109/TIFS.2019.2960647

41. Vinayakumar R., Alazab M., Srinivasan S., Pham Q., Padannayil S. K., and Simran K. A Visualized Botnet Detection System Based Deep Learning for the Internet of Things Networks of Smart Cities. *IEEE Transactions on Industry Applications*, 56(4):4436–4456, 2020. https://doi.org/10.1109/TIA.2020.2971952

42. Li Y., Xiong K., Chin T., and Hu C. A Machine Learning Framework for Domain Generation Algorithm-Based Malware Detection. *IEEE Access*, 7:32765–32782, 2019. https://doi.org/10.1109/ACCESS.2019.2891588

43. J. J. Koh and B. Rhodes. Inline Detection of Domain Generation Algorithms with Context-Sensitive Word Embeddings. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 2966–2971, 2018.

44. Cucchiarelli Alessandro, Morbidoni Christian, Spalazzi Luca, and Baldi Marco. Algorithmically generated malicious domain names detection based on n-grams features. *Expert Systems with Applications*, 170:114551, 2021. ISSN 0957-4174. https://doi.org/10.1016/j.eswa.2020.114551. Available at: https://www.sciencedirect.com/science/article/pii/S0957417420311957.

45. I. Yilmaz, A. Siraj, and D. Ulybyshev. Improving DGA-Based Malicious Domain Classifiers for Malware Defense with Adversarial Machine Learning. In *2020 IEEE 4th Conference on Information Communication Technology (CICT)*, pages 1–6, 2020.

46. Sivaguru R., Peck J., Olumofin F., Nascimento A., and De Cock M. Inline Detection of DGA Domains Using Side Information. In *IEEE Access*, volume 8, pages 141910–141922, 2020. https://doi.org/10.1109/ACCESS.2020.3013494

47. Meng Hao, Hongwei Li, Hanxiao Chen, Pengzhi Xing, Guowen Xu, and Tianwei Zhang. Iron: Private Inference on Transformers. In *Advances in Neural Information Processing Systems*, 2022. Available at: https://openreview.net/forum?id=deyqjpcTfsG.

48. Samuel Adams, David Melanson, and Martine De Cock. Private Text Classification with Convolutional Neural Networks. In *Proceedings of the Third Workshop on Privacy in Natural Language Processing*, pages 53–58, 2021.

49. Qi Feng, Debiao He, Zhe Liu, Huaqun Wang, and Kim-Kwang Raymond Choo. SecureNLP: A system for multi-party privacy-preserving natural language processing. *IEEE Transactions on Information Forensics and Security*, 15:3709–3721, 2020. https://doi.org/10.1109/TIFS.2020.2997134

50. Brian Knott, Shobha Venkataraman, Awni Hannun, Shubho Sengupta, Mark Ibrahim, and Laurens van der Maaten. Crypten: Secure multi-party computation meets machine learning. *Advances in Neural Information Processing Systems*, 34, 2021.

51. Andrej Karpathy, Justin Johnson, and Li Fei-Fei. Visualizing and Understanding Recurrent Networks. 2015. arXiv preprint arXiv:1506.02078. Available at: https://arxiv.org/abs/1506.02078.

52. Pierre Lison and Vasileios Mavroeidis. Automatic Detection of Malware-Generated Domains with Recurrent Neural Models. 2017. arXiv preprint arXiv:1709.07102. Available at: https://arxiv.org/abs/1709.07102.

53. Hieu Mac, Duc Tran, Van Tong, Linh Giang Nguyen, and Hai Anh Tran. DGA Botnet Detection Using Supervised Learning Methods. In *Proceedings of the Eighth International Symposium on Information and Communication Technology (SoICT 2017)*, pages 211–218, Nha Trang City, Viet Nam, 2017. Association for Computing Machinery, New York, NY, USA. ISBN 9781450353281. https://doi.org/10.1145/3155133.3155166.

54. Christopher Olah. Understanding LSTM Networks. 2023. Available at: https://colah.github.io/posts/2015-08-Understanding-LSTMs/. Accessed: 2023-07-04.

55. Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28, 2015.

56. Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.

57. David Evans, Vladimir Kolesnikov, and Mike Rosulek. A Pragmatic Introduction to Secure Multi-Party Computation. *Found. Trends Priv. Secur.*, 2(2–3):70–246, Dec 2018.

58. Yehuda Lindell. Secure Multiparty Computation. *Commun. ACM*, 64(1):86–96, January 2021. https://doi.org/10.1145/3387108

59. Ronald Cramer, Ivan Damgård, Dario Catalano, Giovanni Crescenzo, Ivan Darmgård, David Pointcheval, et al. Multiparty Computation, an Introduction. 2006.

60. Marcel Keller and Ke Sun. Secure Quantized Training for Deep Learning. Cryptology ePrint Archive, Paper 2022/933, 2022. Available at: https://eprint.iacr.org/2022/933.

61. Anders Dalskov, Daniel Escudero, and Marcel Keller. Secure Evaluation of Quantized Neural Networks. Proceedings on Privacy Enhancing Technologies, 2020(4):355–375, 2020. Available at: http://dx.doi.org/10.2478/popets-2020-0077.

62. Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, et al. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2704–2713, 2018.

63. Ricardo J. M. Maia Private DGA Detection. Available at: https://github.com/ricardojmmaia/private-dga-detection. Accessed: 2024-05-18.

64. Cyber Analysis & Defense department of Fraunhofer FKIE. DGArchive. Available at: https://dgarchive.caad.fkie.fraunhofer.de/welcome/. Accessed: 2023-07-04.

65. Wikipedia Contributors. Alexa Internet. Available at: https://en.wikipedia.org/wiki/Alexa_Internet. Accessed: 2023-07-04.

66. TensorFlow Authors. DPKerasAdamOptimizer. Available at: https://www.tensorflow.org/responsible_ai/privacy/api_docs/python/tf_privacy/DPKerasAdamOptimizer/. Accessed: 2023-07-04.