

Few-Shot Website Fingerprinting Attack

Mantun Chen, Yongjun Wang, Zhiquan Qin and Xiatian Zhu

Abstract—This work introduces a novel data augmentation method for few-shot website fingerprinting (WF) attack where only a handful of training samples per website are available for deep learning model optimization. Moving beyond earlier WF methods relying on manually-engineered feature representations, more advanced deep learning alternatives demonstrate that learning feature representations automatically from training data is superior. Nonetheless, this advantage is subject to an *unrealistic* assumption that there exist many training samples per website, which otherwise will disappear. To address this, we introduce a model-agnostic, efficient, and *Harmonious Data Augmentation* (HDA) method that can improve deep WF attacking methods significantly. HDA involves both intra-sample and inter-sample data transformations that can be used in harmonious manner to expand a tiny training dataset to an arbitrarily large collection, therefore effectively and explicitly addressing the intrinsic data scarcity problem. We conducted expensive experiments to validate our HDA for boosting state-of-the-art deep learning WF attack models in both closed-world and open-world attacking scenarios, at absence and presence of strong defense. For instance, in the more challenging and realistic evaluation scenario with WTF-PAD based defense, our HDA method surpasses the previous state-of-the-art results by more than 4% in absolute classification accuracy in the 20-shot learning case.

Index Terms—User privacy, Internet anonymity, Data traffic patterns, Website fingerprinting, Deep learning, Neural network, few-shot learning, Data augmentation.

I. INTRODUCTION

FOR privacy protection in accessing Internet, an increasing number of users have turned to anonymous networks and Tor [13] is one of the most popular choices [12]. However, this remains not completely secure due to exposure of data transportation pattern before reaching Tor servers. For instance, a local attacker would eavesdrop the connection between a user and the guard node of Tor networks, with the attacking positions including any devices in the same LAN or wireless network, switch, router, and compromised Tor guard node (see Figure 1). By just analyzing the patterns of data package traffic without observing the content inside, the attacker is likely to reason about which website a target user is visiting. This is often known as *website fingerprinting* (WF) attack [18].

To implement WF attack, the attacker needs to first create a particular digital fingerprint for every individual website, and then learn some intrinsic pattern characteristics of these fingerprints for accomplishing attack. Earlier attacking methods rely on manually designed features based on expert domain knowledge [5], [7], [15], [17], [18], [25], [27], [28], [37], [38]. They are not only inflexible but also susceptible




Fig. 1: Illustration of data flow traffic between a user and target websites with a Tor network in-between. Despite being more secure by anonymity, website fingerprinting attackers are still able to reason about which website a victim user is visiting by analyzing the data traffic characteristics at multiple locations, as specified by red dash lines.

to environmental changes over time. This limitation can now be solved by using more advanced deep learning techniques [23]. This is because other than utilizing manually designed features, deep learning methods can learn automatically feature representations directly from training data, and are hence more scalable provided that up-to-date training data are accessible. A couple of latest state-of-the-art studies, Deep Fingerprinting (DF) [32] and Var-CNN [4], have demonstrated this potential in comparison to manual feature based methods. However, these deep learning solutions are not perfect, as their success is established upon an unrealistic assumption that a sufficiently large number (e.g. hundreds) of training samples per website are available. That is, data hungry. When only a small training dataset is given as typical in practical use, their performances are not necessarily superior to traditional methods [15], [27], [37]. It is always expensive, tedious, or even infeasible to collect a vast training set in reality due to highly frequent and continuous changes in Internet environments. Consequently, WF attack is fundamentally a *few-shot learning* problem, which nevertheless is largely unrecognized in the literature.

The nature of few-shot WF attack is also considered in the recent Triplet Fingerprinting method [33], under a condition that there is a large set of relevant auxiliary training samples for model pre-training. It is essentially a transfer learning setting. This will limit significantly its scalability in practice as *in-the-wild* changes of Internet data traffic conditions would render such assumptions to be invalid at high probabilities. On the contrary, we introduce a realistic, generic few-shot WF attack setting where only a handful of training samples are available for every target website, without making any

M. Chen, Z. Qin and Y. Wang are with College of Computer, National University of Defense Technology, Changsha 410073, China. E-mail: {chenmantun19,tanzhiquan14,wangyongjun}@nudt.edu.cn

Xiatian Zhu is with University of Surrey, Stag Hill, University Campus, Guildford GU2 7XH, UK. E-mail: eddy.zhuxu@gmail.com

domain-specific assumptions. Clearly, Triplet Fingerprinting is not applicable in our setting due to the need of auxiliary training data.

We summarize the **contributions** of this paper as follows:

(I) We introduce a novel, practical *few-shot website fingerprinting attack* problem, in which only a few training samples are available without rich auxiliary data. This respects the intrinsic nature of highly dynamic Internet traffic conditions and high cost of collecting large training data in practice. Highlighting the importance of *few-shot learning* without any auxiliary data assumption for the first time, we hope more future efforts would be dedicated for solving practically important WF attack challenge.

(II) To solve the proposed few-shot learning challenges, we embrace the enormous potentials and advantages of deep learning for WF attack by introducing a new *Harmonious Data Augmentation* (HDA) method to explicitly solve the training data scarcity problem in deep learning. Specifically, we augment the original training data by rotating and masking-out randomly individual samples and mixing (linearly combining) sample pairs in arbitrary proportions. With such intra-sample and inter-sample data transformations, our HDA method can efficiently expand a tiny training dataset at any scales.

(III) We benchmark the performance of few-shot WF attack and demonstrate the efficacy of our data augmentation method using existing state-of-the-art deep learning models. In particular, we consider 5-20 shots per website/class in closed-world and open-world settings, with and without defense. The results show that our method can improve the performances of previous state-of-the-art deep learning solutions [4], [16] significantly.

II. RELATED WORK

A. Objectives, Scenarios and Assumptions

The objective of WF attack is to identify which website a victim user is interacting with among a set of monitored target websites. Conceptually, it is a multi-class classification problem with each website regarded as a unique class. There are several scenarios with different assumptions. The most common scenario is *closed-world* attack, where the user is assumed to only visit a set of known target websites under monitoring. This assumption however is not realistic, therefore discarded in the *open-world* scenario. In this scenario, the victim user is considered to likely visit any websites including those monitored ones, as typically experienced in real-world applications.

A third scenario considers *defense* where the user takes some actions to defend against potential attack. This would

lead to higher attack difficulty. Representative defense techniques include Buflo [14], Tamaraw [6], Walkie-Talkie [40] and WTF-PAD [20]. Among them, WTF-PAD is used as the mainstream method for Tor networks due to low bandwidth overhead and zero delay. We considered WTF-PAD based defense in our evaluations.

In the literature, several common assumptions are made. We briefly discussed three main assumptions. In *user behavior*, it is assumed that all Tor users browsed websites sequentially, only opening a single tab at a time. In *background traffic*, it is assumed that the attacker is able to collect all the clean traces generated by the victim's visits against dynamic background traffic. This is increasingly possible as shown in [39] the multiplexed TLS traffic can be split into individual encrypted connections to each website. In *network condition*, the attacker is assumed to have the same conditions as the victim including traffic condition and settings. To compare with the benchmark results, we follow these general assumptions for fair evaluations.

Instead, we focus on addressing the following assumption. Often, the attacker assumes that the training data fall into a similar distribution as the deployment data. This is a particularly strong and artificial assumption as the network condition is actually changing and evolving frequently. Such a property enforces the attacker to update the training data in order to have a robust attacking model over time. This implies that the attacker is less possible to collect a large set of training data at each time due to high acquiring costs. However, existing WF attack methods often ignore this factor by assuming availability of large training data. In contrast, we study the largely ignored few-shot learning setting in WF attack. Specifically, we approach this problem by explicitly solving the small training data issue via synthesizing new labelled training data.

B. Website Fingerprinting Attack Methods

The first pioneer attack against Tor networks was evaluated by Herrmann [17] in 2009. It achieved an accuracy of 2.96% using around 20 training samples per website in the closed-world scenario. Later on, Wang and Ian [38] proposed to represent the traffic data using more fundamental Tor cells (i.e., direction data) as a unit rather than TCP/IP packets. This representation is rather meaningful and informative as it encodes essential characteristics of Tor data. By training a SVM classifier with distance-based kernel, a ground-breaking performance with 90.9% accuracy was achieved on 100 sites each with 40 training samples. Recently, Panchenko et al. [27] proposed an idea of sampling the features from a cumulative trace representation and achieved 91.38% accuracy with 90 training instances per website. Hayes and Danezis [15] exploited random decision forests to achieve similar results. A common design of these above methods is a two-stage strategy including feature design and classifier learning. This is not only constrained by the limitations of hand features but also lacks interaction between the two stages, making the model performance inferior.

Motivated by the remarkable success of deep learning techniques in computer vision and natural language processing

[10], [22], several deep learning WF attack methods have been introduced which can well solve the aforementioned weakness. This is because deep learning methods by design carry out feature learning and classification optimization from the raw training data end-to-end. For example, using VGG network [31] as the backbone, Sirinam et al. [32] proposed a Deep Fingerprinting attack (DF) model that attains 98.3% accuracy on 95 websites. However, this method needs a large training set (e.g. 1000 training samples per website), otherwise it will suffer from significant performance drop. When using 20 training samples per website, DF can only hit around 19.4% accuracy. To overcome this limitation, Bhat et al. [4] developed the Var-CNN model based on ResNet [16] and dilated causal convolution [11], [42]. When small training sets (e.g. 100 samples per website) are available, it achieves superior performance over DF but at dependence on less realistic time features and less scalable hand-crafted statistical information.

A solution to few-shot learning is a recently proposed triplet fingerprinting (TF) method [33]. The key idea of TF is to pre-train a metric model that can measure pairwise distances on new classes. When the pre-training dataset is similar to the target data in distribution, TF can hit an accuracy of 94.5% on 100 websites using only 20 training samples per website. This is a strong transfer learning scenario. However, considering that the dynamics of network conditions is highly unknown and uncontrollable, such a transfer learning assumption is hardly valid in practice. In light of this observation, in this work we propose a more realistic few-shot learning setting without assuming any auxiliary data with similar data characteristics for model pre-training. Hence it is more scalable and generic for real-world deployments. Under the proposed more challenging few-shot setting, TF is unable to work properly due to insufficient network initialization.

C. Data Augmentation

Data augmentation is an important element in deep learning due to its data-hungry nature [23]. For example, random insertion, random swap, and random deletion for text classification in natural language processing [41], or geometric transformations (e.g., flipping, rotation, translation, cropping, scaling), color space transformations (e.g., color casting, varying brightness, and noise injection), inter-image mixup [43] for image analysis [26], [30], [35]. These previous attempts have shown the significance of different augmenting methods for model performance on the respective tasks. Inspired by these findings, we investigate extensively the effectiveness of training data augmentation by adapting existing operations for deep learning WF attack in few-shot learning settings. To the best of our knowledge, this is the first attempt of its kind. Crucially, we demonstrate that existing state-of-the-art deep WF attack method [4] significantly benefits from using the proposed data augmentation operations in varying evaluation scenarios. This result would be encouraging and influential for future investigation of deep learning WF attack methods in particular.

III. METHOD

A. Problem Definition

In website fingerprinting (WF) attack, the *objective* is to detect which website a target user is visiting. The common observations are data traffic traces \mathbf{x} produced by one visit to a website y . Taking each website as a specific class, this is essentially a multi-class classification problem. For model training, a labelled training set $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ is often provided, where $y_i \in \{1, 2, \dots, K\}$ specifying one of K target websites. Two different settings are often considered in model testing: (1) *Closed-world* attack where any test sample is assumed to belong one of target websites/classes, and (2) *Open-world* attack where the above assumption is eliminated, i.e., a test trace may be produced by a *non-target* (unmonitored) website. The latter is a more realistic setting, yet presenting a more challenging task as identifying if a test sample falls into target classes or not is non-trivial.

Feature representation. For common Tor networks, the raw representation of a specific traffic trace consists of a sequence of temporally successive Tor cells travelling between a target user and a website visited. It is derived from TCP/IP data. Specifically, after those TCP/IP packets retransmitted are discarded, TLS records are first reconstructed, their lengths are then rounded down to the nearest multiple of 512 to form the final sequence data \mathbf{x} . In value, each \mathbf{x} is a sequence of 1 (outgoing cell) and -1 (incoming cell), with a variable length. This raw representation is hence known as direction sample. Besides, temporal information about inter-packet time is another modality of data used, but limited by high reliance on network conditions, i.e., not stable and much more noise. Consequently, we mainly consider the direction data samples in this study, which are more scalable and generic.

B. Deep Learning for Website fingerprinting Attack

Most of existing WF attack methods rely on hand-crafted feature representations [5], [7], [15], [17], [18], [25], [27], [28], [37], [38]. This strategy is not only unscalable but also unsatisfactory in performance due to limited and incomplete domain knowledge. Deep learning methods provide a viable solution via learning directly more effective and expressive representation from training data, as shown in a few recent studies [4], [32]. In this work we advance this new direction further.

We explore 1D convolutional neural networks (CNN) [21] for WF attack as the raw data are temporal sequences. Building on the success of deep learning in computer vision, we adopt the same high-level network designs of standard 2D CNN models [24] whilst translating them into 1D counterparts. This is similar to [4], [32].

A CNN model consists of multiple convolutional layers with non-linear activation functions such as ReLU [2] and fully-connected (FC) layers, characterized by end-to-end feature extraction and classification. With convolutional operations, the filters of each layer transform input sequences using learnable parameters and output new feature sequences. This feature transformation is conducted layer by layer in a

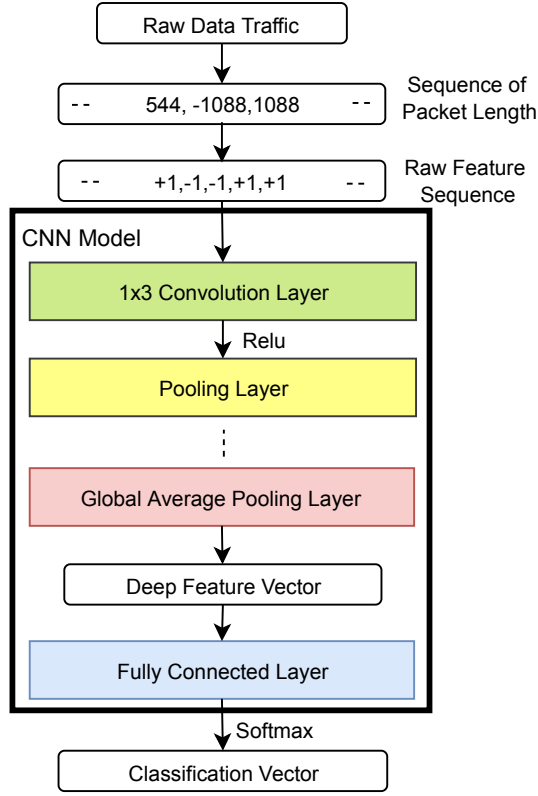


Fig. 2: A deep learning pipeline for website fingerprinting attack that conducts feature representation and website classification end-to-end in a joint learning manner.

hierarchical fashion. The receptive field (kernel) with size 3 is often used in each individual layer to capture local feature patterns. By stacking more layers and pooling operation, the model can perceive the information of larger regions and achieve translational invariance. Another effective method for enlarging receptive field is dilated causal convolutions [11], [42], which has been exploited in [4].

The feature representations \mathbf{f} of WF samples are the output of global average pooling layer on top of the last convolution layer. To obtain the classification probability vector $\hat{\mathbf{y}} = \{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_K\} \in \mathcal{R}^K$ over K target classes, \mathbf{f} is fed into a FC layer and normalized by a softmax function.

For model training, we compute a cross-entropy objective loss function with the classification vector against the ground-truth class label over all N training samples as:

$$\mathcal{L} = \sum_{i=1}^N \sum_{j=1}^K \delta(j = y_i) \log \hat{y}_{i,j} \quad (1)$$

where y_i refers to the ground-truth class label of a training sample \mathbf{x}_i , and $\delta(\cdot)$ is a Dirac function. The objective is to maximize the probability of the ground-truth class in prediction. This loss function is differentiable, with its gradients backpropagated to update all the learnable model parameters.

Once the deep model is trained, we forward a given test sample, obtain a classification probability vector, and take the most likely class as prediction in both closed-world and

open-world settings. For open-world setting, all unmonitored websites are considered to belong to a background class.

Discussion. While deep learning techniques have advanced significantly in the last several years, it is still assumed that a large set of labelled training samples is available. This is not always true, for example, for the WF attack problems. In real-world applications, an attacker is usually faced with highly dynamic network environments. It means that the distribution of raw features is evolving continuously. As such, the training data need to update frequently, which disables collection of large training data with labels in practice due to prohibitively high labelling costs. Consequently, only a small training set is accessible in reality, making deep learning methods ineffective.

C. Harmonious Website Fingerprinting Data Augmentation

To address the above small training data challenge, we propose an intuitive, novel *harmonious data augmentation* (HDA) method. We introduce both *intra-sample* and *inter-sample* augmentation operations that can be applied in a joint and harmonious manner for more effective data expansion.

Intra-sample augmentation. The key idea of intra-sample augmentation is that given an individual training sample, we introduce a certain degree of *random* data perturbation and/or variation whilst keeping the same class labels. Doing so allows us to generate an infinite number of labelled training samples due to the nature of randomness. We consider two perturbation operations: random rotation and random masking.

Random rotation based data augmentation means rotating an original training sample forward or backward by random steps to generate virtual samples (Fig. 3(a)):

$$\text{Rotate}(\mathbf{x}, n_{\text{step}}, \text{dir}) \quad (2)$$

where n_{step} and $\text{dir} \in \{\text{forward}, \text{backward}\}$ specify the steps and the direction to rotate on an input sample \mathbf{x} . The hypothesis behind is that class-sensitive information encoded in a sample is distributed across different sub-sequences and data traffic order is less important than signal patterns. After a sample is rotated, the original class information is largely preserved, i.e., semantically invariant. Hence, the same class can be annotated for the rotated variants. However, this hypothesis is more likely to stand under some certain (unknown) degrees. We therefore introduce an upper bound parameter R_{max} so that the rotation range is limited at most R_{max} steps in both directions, $n_{\text{step}} \leq R_{\text{max}}$.

In contrast, *random masking* introduces localized corruption to an original training sample by setting a random subsequence to zero (Fig. 3(b)). This data augmentation is written as:

$$\text{Mask}(\mathbf{x}, n_{\text{len}}, \text{loc}) \quad (3)$$

where n_{len} and loc denote the length and location of the subsequence that is masked out from an original sample \mathbf{x} . Rather than in form of subsequence, another strategy is to randomly select individual positions to mask. We consider this may introduce more significant corruption to the underlying semantic information.

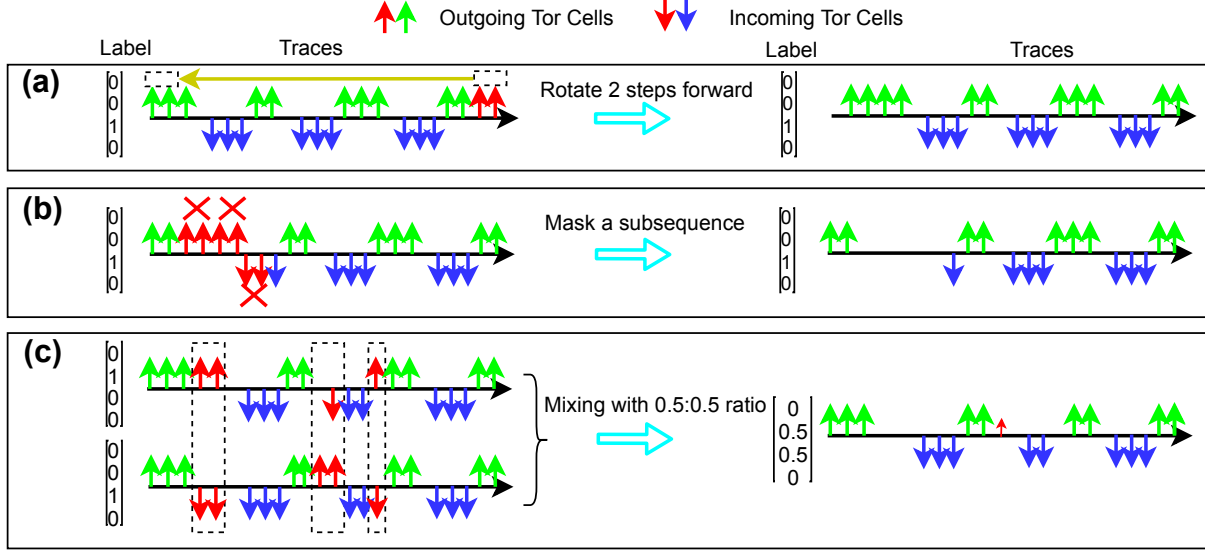


Fig. 3: Illustration of our data augmentation operations for deep learning WF attack, including (a) random rotation, (b) random masking, and (c) random mixing.

Conceptually, random masking simulates varying traffic measurement errors in data transportation. Meanwhile, with the same above hypothesis, such masking would not dramatically change the semantic class information, provided that the masking is subject to some limit, e.g., the length of subsequences masked out M_{len} . It hence offers a complementary data perturbation choice w.r.t. random rotation.

Inter-sample augmentation. Apart from data augmentation on individual samples, we further introduce data perturbation across two different samples to enrich the limited training set.

We propose *random mixing* that generates virtual samples and class labels by linear interpolation between two original samples \mathbf{x}_i and \mathbf{x}_j as:

$$\tilde{\mathbf{x}} = \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_j \quad (4)$$

$$\tilde{\mathbf{y}} = \lambda \mathbf{y}_i + (1 - \lambda) \mathbf{y}_j \quad (5)$$

where $(\mathbf{y}_i, \mathbf{y}_j)$ are the one-hot class labels of \mathbf{x}_i and \mathbf{x}_j . The mixing parameter $\lambda \in [0, 1]$ follows a Beta distribution: $\lambda \sim \beta(\alpha, \alpha)$ with $\alpha > 0$ the parameter that controls the strength of interpolation. This is in a similar spirit of mixup in image understanding domain [43]. Unlike intra-sample augmentation above, random mixing changes the semantic class information, since original samples may be drawn from different classes. It simplifies the data distribution by imposing a linear relationship between classes for complexity minimization. While seemingly counter-intuitive, we will show that such a method brings positive contributions on top of random masking and random rotation.

Combination and compatibility. Different augmentation operations can be applied on the same samples without conflict to each other in a harmony. There is also no particular constraint on the order of applying all the three data augmentation operations in a combination. Given a fixed set of parameters as discussed above, different augmentation orders will result

in different virtual samples. This makes little conceptual difference as the space of sample is just infinite.

Augmentation optimization. In our harmonious data augmentation (HDA), three hyper-parameters $\{R_{max}, M_{len}, \alpha\}$ are introduced. To generate meaningful virtual samples, obtaining their optimal values is necessary otherwise adversarial effects may even be imposed.

Instead of manual tuning, we adopt an automatic Bayesian estimator, called Tree of Parzen Estimators (TPE) [3]. The conventional TPE can take only one parameters alone at a time. So we need to optimize each of the three hyper-parameters independently. This differs from our data augmentation process where the three augmentation operations are typically applied together, making the independently tuned parameters of TPE sub-optimal. This is because, jointly applying three augmentations together makes them inter-dependent.

For solving this problem, we propose a sequential optimization process that takes into account the inter-dependence property of different augmentation operations gradually (see Alg. 1). Specifically, we start with a random, fixed order of applying our random rotation, masking, and mixing operations. Then, we optimize from the first one with TPE, move to next one with all the previous ones optimized and fixed, stop by finishing the last one. Each time, we still optimize a single hyper-parameter whilst keeping all the previous optimized ones fixed. In this way, we expand the inter-dependence among different operations sequentially.

D. Theoretical Foundation and Formulation

The objective of learning a WF attack model is equivalent to derive a function $h \in H$ that fits the latent translation relationship between raw feature vectors $\mathbf{x} \in X$ and corresponding website class labels $\mathbf{y} \in Y$. That is, fitting a joint distribution $P(X, Y)$. To this end, in deep learning we often

Algorithm 1 Data augmentation optimization

Input: A training $\{X_t, Y_t\}$, and validation $\{X_v, Y_v\}$ set.

Output: Data augmentation with optimal parameters B_{aug} .

- 1: Setting $B_{\text{aug}} = \phi$ (empty set);
 - 2: Sequencing data augmentation operations randomly;
 - 3: **while** Enumerating augmentation operations **do**
 - 4: Get the search space S_{aug} of current augmentation A ;
 - 5: Using TPE on S_{aug} to obtain the optimal parameter b_{aug} , with the model trained by B_{aug} and A ;
 - 6: $B_{\text{aug}} = B_{\text{aug}} \cup b_{\text{aug}}$
 - 7: **end while**
 - 8: **return** B_{aug}
-

leverage a loss function L defined to penalize the differences between predictions $h(x)$ and targets y . We minimize the average loss over the joint distribution:

$$R(h) = \int L(h(x), y) dP(x, y) \quad (6)$$

which is known expected risk minimization [36].

However, the joint distribution is often unknown, particularly for WF attack with small training data. Given a limited training data set $D = \{(x_i, y_i)\}_{i=1}^N$, the joint distribution can only be approximated by an empirical distribution as:

$$P_\delta(x, y) = \frac{1}{N} \sum_{i=1}^N \delta(x = x_i, y = y_i) \quad (7)$$

where $\delta(x = x_i, y = y_i)$ is a Dirac mass centered at a sample (x_i, y_i) . Accordingly, the expected risk can now be approximated by an empirical risk:

$$\begin{aligned} R_\delta(h) &= \int L(h(x), y) dP_\delta(x, y) \\ &= \frac{1}{N} \sum_{i=1}^N L(h(x_i), y_i) \end{aligned} \quad (8)$$

The above approximation is in the empirical risk minimization (ERM) principle [36]. The cross-entropy loss (Eq. (1)) is a representative example, which essentially minimizes $R_\delta(h)$ for the classification task.

While ERM is a common strategy, it suffers from high risk of poor generalization due to the tendency of memorization, particularly when a large model is used [34]. To mitigate this issue, we adopt the notion of vicinal distribution [8] which can better approximate the true joint distribution. In particular, the vicinal distribution P_v in the data space is defined as:

$$P_v(\tilde{x}, \tilde{y}) = \frac{1}{n} \sum_{i=1}^n v(\tilde{x}_i, \tilde{y}_i | x_i, y_i) \quad (9)$$

Intuitively, P_v measures the probability of finding a virtual labelled sample (\tilde{x}, \tilde{y}) in the vicinity around an original training sample (x_i, y_i) .

Given such vicinal distributions, we first construct a virtual dataset $D_v := (\tilde{x}_i, \tilde{y}_i)_{i=1}^m$ by sampling P_v randomly, and then minimize an empirical vicinal risk to learn h as:

$$R_v(h) = \frac{1}{m} \sum_{i=1}^m L(h(\tilde{x}_i), \tilde{y}_i) \quad (10)$$

Clearly, at the core of this strategy is performing data augmentation around original training samples. Rather than computing a loss value for every single training sample, it derives a local distribution centered at each individual sample and generates more virtual training samples to reduce the negative memorization effect of deep learning. This is the key rationale of our data augmentation method.

Augmentation formulation. We formulate the proposed harmonious data augmentation operations in the vicinal distribution manner. For intra-sample augmentation (including random rotation and masking), the vicinal distribution is defined as

$$v(\tilde{x}, \tilde{y} | x, y) = T(x) \delta(\tilde{y} = y) \quad (11)$$

where $T()$ is a transformation operator.

For *random rotation*, given any length- n sample $x = \{x_0, \dots, x_i, \dots, x_{n-1}\}$, we first define a circle matrix B for forward rotation as:

$$B(x) = \begin{bmatrix} x_0 & x_1 & \cdots & x_{n-1} \\ x_{n-1} & x_0 & \cdots & x_{n-2} \\ \vdots & \vdots & & \vdots \\ x_{n-1} & x_{n-2} & \cdots & x_0 \end{bmatrix} \quad (12)$$

Then we sample the step size n_{step} *uniformly* from a range of $\{1, \dots, R_{\text{max}}\}$. By one-hot representation of n_{step} , we can obtain a rotation transformation as:

$$T_{\text{rot}}(x) = \text{one-hot}(n_{\text{step}}) B(x) \quad (13)$$

For the backward case, we perform the same process as above but with a backward rotation matrix instead.

For *random masking*, we similarly sample the start position s *uniformly* in the range of $\{1, \dots, n - n_{\text{len}}\}$ where n_{len} is the length of masked subsequence. The masking transformation can be represented by a matrix as:

$$M_{\text{mask}} = \text{diag}\left(\mathbf{1} - \sum_{i=s}^{s+n_{\text{len}}} \text{Row}_i(I)\right) \quad (14)$$

where I is identity matrix, $\mathbf{1}$ is all-one vector, $\text{Row}_i()$ selects the i -th row of a matrix, $\text{diag}()$ transforms a vector to a diagonal matrix. Masking operation is finally conducted by matrix multiplication as:

$$T_{\text{mask}}(x) = x M_{\text{mask}} \quad (15)$$

For inter-sample augmentation, *random mixing* in our case, the vicinal distribution is defined as:

$$\begin{aligned} v(\tilde{x}, \tilde{y} | x_i, y_i, x_j, y_j) &= \\ \delta(\tilde{x} &= \lambda \cdot x_i + (1 - \lambda) \cdot x_j, \tilde{y} = \lambda \cdot y_i + (1 - \lambda) \cdot y_j) \end{aligned} \quad (16)$$

where λ is a random variable drawn from a Beta distribution $\beta(\alpha, \alpha)$ and y is one-hot class label vector. This local vicinity is assumed to respect a linear structure w.r.t. class labels.

IV. EXPERIMENTS

A. Experimental Setup

Datasets. We evaluated our data augmentation method HDA on four standard WF attack datasets as below. (1) **CW₁₀₀** [29]: This dataset provides a total of 100 monitored target websites each with 2,500 raw feature traces. (2) **DF_{95,Nodef}** [32]: This dataset gives 95 monitored websites with each contributing 1,000 feature traces. (3) **ROWUM_{400,000}** [29]: This dataset includes CW₁₀₀ and a large set of samples each was generated by a visit to a page of top 400,000 Alexa websites. (4) **DF_{95,wtf-pad}** [32]: Unlike all the above datasets, this is a more challenging dataset due to presence of WTF-PAD based defense against WF attack. It has the same size as DF_{95,Nodef}, i.e., 95,000 raw feature samples from 95 websites. We considered both closed-world and open-world WF attack scenarios using the above datasets.

Network architectures. We used four different network architectures for testing the generic benefits of the proposed HDA method. (1) Var-CNN [4] is the current state-of-the-art deep learning WF method. (2) Var-CNN* is an improved variant with two more fully-connected layers added to the classifier. (3/4) ResNet-18 and ResNet-34 [16] are two strong and popular networks widely deployed in many different fields such as computer vision.

Implementation details. We conducted our experiments in Keras [9]. In all our experiments, we used the standard training, validation, and test splits for all competitors for fair comparisons. We optimized HDA’s hyper-parameters using Var-CNN [4] as deep learning model on CW₁₀₀ in closed-world setting and applied the same parameter setting for all the other deep learning methods, datasets and settings. This allows testing the generality and scalability of our HDA method. For augmentation optimization, we set the search space as: 1 ~ 20 with step 5 for forward/backward R_{\max} (random rotation) 1 ~ 200 with step 20 for M_{len} (random masking), [0, 1] with step 0.1 for α (random mixing). The optimal parameter values we obtained are $R_{\max} = 20$, $M_{\text{len}} = 180$, and $\alpha = 0.1$. For saving storage, we performed online data augmentation within each mini-batch without any data pre-processing. In each experiment, we trained every deep learning model for 150 epochs and used the checkpoint with best performance on the validation set for model test. We only used the direction feature data, without time sequences and hand-crafted features. We run each experiment by 10 times and reported the mean results and standard deviation as the final performance.

B. Closed-World WF Attack

Setting. We conducted the closed-world attack on CW₁₀₀ and DF_{95,Nodef}. We separated each dataset into training, validation (10 samples per class), and test (70 samples per class) splits. We considered few-shot settings with $n \in \{5, 10, 15, 20\}$ training samples per class. The validation set was used to select the best performing model for test. We used the classification accuracy as performance metric. Besides deep network models, we also compared our method with two conventional

TABLE I: Results of *closed-world* WF attack on CW₁₀₀. Metrics: Accuracy.

Method	5-shot	10-shot	15-shot	20-shot
CUMUL [27]	72.2 ± 1.7	79.7 ± 1.4	83.3 ± 2.0	85.9 ± 0.6
k-FP [15]	79.3 ± 1.0	83.9 ± 1.0	85.9 ± 0.6	87.5 ± 0.8
ResNet-18 [16]	13.9 ± 0.6	23.0 ± 1.1	35.5 ± 2.5	51.3 ± 1.7
ResNet-34 [16]	14.5 ± 0.7	24.3 ± 1.5	40.3 ± 3.1	51.3 ± 86.4
Var-CNN [4]	17.9 ± 1.5	41.4 ± 4.0	65.6 ± 1.9	78.7 ± 1.5
Var-CNN* [4]	34.8 ± 2.6	57.4 ± 3.9	71.9 ± 2.4	80.8 ± 2.4
ResNet-18+HDA	34.3 ± 4.3	61.5 ± 9.4	77.6 ± 5.8	84.6 ± 4.3
ResNet-34+HDA	34.8 ± 6.2	62.3 ± 8.1	78.8 ± 7.1	86.4 ± 2.8
Var-CNN+HDA	59.7 ± 1.5	74.7 ± 2.6	86.4 ± 1.3	90.7 ± 0.8
Var-CNN*+HDA	71.3 ± 4.7	90.2 ± 0.6	93.3 ± 0.2	94.1 ± 0.5

TABLE II: Results of *closed-world* WF attack on DF_{95,Nodef}. Metrics: Accuracy.

Method	5-shot	10-shot	15-shot	20-shot
ResNet-18 [16]	14.3 ± 1.0	26.3 ± 1.6	41.1 ± 1.9	54.3 ± 1.1
ResNet-34 [16]	16.3 ± 1.3	29.5 ± 4.7	41.4 ± 3.2	54.2 ± 2.6
Var-CNN [4]	22.6 ± 3.4	50.0 ± 2.1	75.6 ± 1.5	75.2 ± 1.7
Var-CNN*	19.4 ± 1.6	29.6 ± 2.7	42.5 ± 5.0	53.6 ± 5.2
ResNet-18+HDA	53.9 ± 7.0	77.3 ± 3.6	82.8 ± 2.9	88.3 ± 1.1
ResNet-34+HDA	44.0 ± 6.8	74.3 ± 7.4	84.2 ± 3.1	88.7 ± 1.6
Var-CNN+HDA	71.0 ± 2.3	84.3 ± 1.4	88.6 ± 0.9	91.4 ± 0.5
Var-CNN*+HDA	62.3 ± 4.8	84.3 ± 0.9	87.7 ± 0.7	91.0 ± 0.5

hand-crafted feature based methods: CUMUL [27] and k-FP [15].

Results. The results of different methods are compared in Table I and Table II. We have the following observations: (1) Hand-crafted feature based methods (CUMUL and k-FP) remain competitive, especially very few samples per class are available. The best 5-shot result on CW₁₀₀ is achieved by k-FP, with a moderate edge of 8% over Var-CNN*+HDA. (2) However, deep learning methods (Var-CNN*) becomes clearly stronger when a few more training samples are accessible, suggesting a great deal of potentials. Among previous methods, in 20-shot case Var-CNN* achieves the best result on CW₁₀₀ and DF_{95,Nodef}. (3) With our HDA method for training data augmentation, every deep learning method improves significantly. For example, the 5-shot accuracy of Var-CNN is increased from 17.9% to 59.7% on CW₁₀₀, and 22.6% to 71.0% on DF_{95,Nodef}. Similarly, the accuracy of ResNet-18 is improved from 13.9% to 34.3% on CW₁₀₀, and from 14.3% to 53.9% on DF_{95,Nodef}. Var-CNN and its variant benefit incredibly more, implying a higher demand for larger training data to avoid model overfit. Similar effects are shown for 10/15/20-shot cases. (4) Our HDA can consistently improve different methods on varying datasets, suggesting good generality. (5) The performance deviation of Var-CNN and its variant assisted by our method HDA is the least among all the competitors,

TABLE III: Results of *open-world* WF attack on CW_{100} (target classes) + $ROWWUM_{400,000}$ (non-target classes). Pre: Precision, Rec: Recall. We reported two settings: one is tuned for best precision (top), and one for recall (bottom).

Method	Tuned for Precision							
	5-shot		10-shot		15-shot		20-shot	
	Pre	Rec	Pre	Recall	Pre	Rec	Pre	Rec
ResNet-18 [16]	28.7	0.8	30.0	7.6	49.8	8.7	65.9	18.3
ResNet-34 [16]	32.5	4.7	42.1	6.6	50.4	21.8	61.3	39.6
Var-CNN [4]	39.7	2.7	58.8	9.2	74.2	35.8	78.0	54.4
Var-CNN ⁺	44.7	3.9	58.9	17.9	64.0	49.9	66.1	65.3
ResNet-18+HDA	80.7	0.2	89.6	7.2	91.7	30.7	93.4	46.6
ResNet-34+HDA	72.7	0.8	91.7	8.7	91.5	43.8	92.6	55.1
Var-CNN+HDA	77.4	6.9	91.2	47.2	91.4	64.3	92.9	66.6
Var-CNN*+HDA	92.3	5.4	87.8	59.7	88.7	73.0	90.5	76.1
Method	Tuned for Recall							
	5-shot		10-shot		15-shot		20-shot	
	Pre	Rec	Pre	Recall	Pre	Rec	Pre	Rec
ResNet-18 [16]	10.0	19.8	17.0	34.0	24.8	49.5	32.3	64.3
ResNet-34 [16]	12	25.6	19.2	37.8	30.2	54.5	34.9	68.6
Var-CNN [4]	13.9	27.1	26.7	52.8	37.4	73.9	42.2	82.6
Var-CNN ⁺	16.3	31.9	27.6	54.8	37.0	73.4	41.1	81.4
ResNet-18+HDA	20.0	36.4	36.3	71.5	44.5	87.4	47.2	91.0
ResNet-34+HDA	21.1	37.4	36.2	68.9	45.7	89.2	47.2	91.4
Var-CNN+HDA	28.0	55.8	46.9	88.7	49.3	92.2	49.1	92.5
Var-CNN*+HDA	25.6	48.1	45.7	89.3	47.9	92.8	47.8	94.1

implying strong stability.

C. Open-World WF Attack

Setting. We conducted the open-world attack experiments on the combination of $ROWWUM_{400,000}$ and CW_{100} . We treat the websites of CW_{100} as target (monitored) classes, and those of $ROWWUM_{400,000}$ as non-target (unmonitored) classes. In this test, we selected randomly 8,020 out of 400,000 unmonitored websites, and separated them into three disjoint sets sized at 20/1,000/7,000 for training, validation, and test, respectively. In this scenario, the precision and recall rates were used to evaluate model performance due to the need for detecting non-target classes [19]. We considered the same four deep learning methods (ResNet-18, Resnet-34, Var-CNN [4] and its variant Var-CNN*) for comparisons.

Results. The results of different methods are reported in Table III. We considered two settings, one is tuned for best precision, and one for best recall. Overall, we obtained similar trends as above that our HDA is highly effective for improving both deep learning methods. It is noted that unlike the closed-world scenario, VarCNN*+HDA achieves very top result at most cases under both tuning settings even if it may not be the best one. Similarly, VarCNN*+HDA remains to be more stable and less sensitive to training sample size. Importantly, our HDA method further enhances this strengths by efficient data augmentation, leading to a more robust WF attack solutions.

D. WF Attack against Defense

Setting. In contrast to the two above experiments, we further tested a more challenging WF attack scenario with defense involved. Defense changes the data traffic patterns to be more similar to each other, therefore making the attack more difficult. We considered the most popular defense, WTF-PAD, widely deployed in Tor networks. We used the $DF_{95,wtf-pad}$ dataset in this experiment. We used 100 random samples per website, and separated them into three sets for training (20 samples), validation (10 samples), and test (70 samples), respectively. We reported the classification accuracy as performance metric in closed-world scenario. We compared with previous four deep learning (ResNet-18, Resnet-34, Var-CNN [4] and its variant Var-CNN*) and hand-crafted feature based methods (k-NN [37], SDAE [1], k-FP [15], CUMUL [27], AWF [29]).

Results. We reported the results of closed-world WF attack under WTF-PAD based defense in Table IV. We made the following observations. (1) Some hand-crafted feature based methods (CUMUL, AWF) are superior over recent deep learning methods (ResNet-18, ResNet-34) at the few-shot learning scenarios. This is mainly because the latter suffers from lacking enough training samples, resulting in model overfitting. (2) Using our HDA for training data augmentation, we can directly solve the data scarcity problem and significantly boost the performances of previous deep learning methods. As a result, both Var-CNN+HDA and Var-CNN*+HDA outperform the

TABLE IV: Results of *closed-world* WF attack with *WTF-PAD* based defense on $DF_{95, wtf-pad}$. Metrics: Accuracy.

Method	5-shot	10-shot	15-shot	20-shot
k-NN [37]	–	–	–	16.0
SDAE [1]	–	–	–	36.9
k-FP [15]	–	–	–	57.0
CUMUL [27]	–	–	–	60.3
AWF [29]	–	–	–	60.8
ResNet-18 [16]	7.3 ± 0.3	9.8 ± 0.6	11.4 ± 0.4	14.2 ± 0.6
ResNet-34 [16]	7.4 ± 0.5	9.4 ± 0.7	13.3 ± 1.3	12.3 ± 1.2
Var-CNN [4]	6.6 ± 0.3	9.2 ± 0.7	12.5 ± 0.8	19.2 ± 1.7
Var-CNN*	7.5 ± 0.4	9.8 ± 0.6	11.5 ± 1.0	15.1 ± 1.3
ResNet-18+HDA	12.9 ± 1.2	27.9 ± 2.8	35.2 ± 2.9	40.7 ± 3.4
ResNet-34+HDA	12.3 ± 1.9	28.1 ± 3.7	38.2 ± 6.2	47.7 ± 5.1
Var-CNN+HDA	25.3 ± 2.2	46.9 ± 1.9	48.7 ± 1.4	63.2 ± 1.8
Var-CNN*+HDA	26.0 ± 4.7	48.5 ± 2.6	59.7 ± 1.9	65.4 ± 0.7

other competitors by a large margin, e.g., 4.6% and 2.4% gap over the best competitor CUMUL. (3) ResNet-34 is surpassed by Var-CNN and its variant dramatically. By benefiting more from our data augmentation, Var-CNN finally achieves the best results across all different shot cases. This implies that Var-CNN has higher desire for large training data with higher performance potential, as compared to ResNet-34.

E. Ablation Studies

We carried out a set of component analysis experiments to examine the exact effect of different designs of our method (HDA). We adopted the most common closed-world attack scenario *without* defense on the CW_{100} dataset, following the same setting as Section IV-B. It is noteworthy that this dataset CW_{100} is different from the dataset in Section IV-B because they are different subsets. In this section, we evaluated the 15-shot learning case in particular, using Var-CNN [4] as the deep learning model backbone.

TABLE V: Effect of individual augmentation operations.

Augmentation Operation	Accuracy
<i>None</i>	75.0 ± 3.0
Random Rotation	92.4 ± 0.3
Random Masking	92.4 ± 0.8
Random Mixing	86.7 ± 0.6
Random Rotation + Masking	92.7 ± 0.4
Random Rotation + Mixing	92.6 ± 0.7
Random Masking + Mixing	93.4 ± 0.7
Random Rotation + Masking + Mixing (HDA)	93.5 ± 0.4

Individual augmentation operations. Recall that our data augmentation method (HDA) consists of three different operations (random rotation, masking, and mixing). We have demonstrated their performance advantages of them as a whole

in varying test settings above. For in-depth insights, examining their individual contributions would be informative and necessary, as well as different combinations. We conducted this experiments with an exhaustive set of operation combinations and reported the results in Table V.

It is observed that: (1) Each of the three operations makes significant difference in performance, with rotation and masking the best individual operations that improve the classification accuracy by 17.4%. (2) When jointly using any two augmentation operations, the performance can be further increased. The combination of masking and mixing gives the highest accuracy among them. (3) Combining all the three operations (HDA) achieves the best result with smaller deviation. This suggests that all different operations are complementary and compatible to each other.

Augmentation optimization. For optimal data augmentation, we propose a sequential optimization strategy (see Alg. 1) for capturing the inter-dependence between different augmentation operations applied. To evaluate its effect, we compared with a baseline algorithm that *independently* optimizes each augmentation parameter.

TABLE VI: Effect of augmentation optimization.

Augmentation Optimization	Accuracy
Independent	92.1 ± 0.5
Sequential (Ours)	93.5 ± 0.4

As shown in Table VI, the proposed optimization algorithm (see Alg. 1) is clearly superior, validating our consideration that there exist inter-dependence between different augmentation operations when applied jointly on the same samples. Note that we obtained this performance gain at the same cost as the baseline counterpart. Besides, it is worth noting that even with the simpler optimization, our data augmentation method (HDA) can still greatly improve the previous deep learning model Var-CNN, and achieve new state-of-the-art results (Table VI vs. Table I). This further validates that the proposed augmentation operations are highly compatible with one another and can be applied together well.

V. CONCLUSION

We presented a model-agnostic, simple yet surprisingly effective data augmentation method, called HDA, for few-shot website fingerprinting attack. This is an under-studied and realistically critical problem, as in practice only a handful of training samples per website can be feasibly collected due to the inherent high dynamics of internet networks and expensive label collection cost. Importantly, we focus on deep learning based methods, a line of new research efforts with vast potentials for future investigations. In particular, our HDA method offers three different data augmentation operations, including random rotation, masking, and mixing in intra-sample and inter-sample fashion. They can be applied to the same training samples harmoniously with high complement and compatibility. Moreover, we introduce a sequential augmentation parameter optimization method that captures the inter-dependence

nature between different operations when applied jointly. With recent state-of-the-art deep learning WF attack models, we conducted extensive experiments on four benchmark datasets to validate the efficacy of our HDA method in both closed-world and open-world scenarios, with and without defense. The results show that the proposed data augmentation method makes dramatic differences in performance and enables previous deep learning methods to outperform hand-crafted feature based counterparts in the few-shot learning setting for the first time, often by a large margin. This is achieved without making any artificial assumptions of relevant, large auxiliary training data for model pre-training. With our HDA method, collecting large training data frequently is eliminated, whilst still achieving stronger and more robust WF attack. Finally, we performed detailed component analysis to diagnose the effect of individual model components.

REFERENCES

- [1] K. Abe and S. Goto, "Fingerprinting attack on tor anonymity using deep learning," in *Proceedings of the Asia-Pacific Advanced Network Research Workshop*, 2016, pp. 15–20.
- [2] A. F. Agarap, "Deep learning using rectified linear units (relu)," *arXiv preprint arXiv:1803.08375*, 2018.
- [3] J. Bergstra, R. Bardenet, B. Kégl, and Y. Bengio, "Algorithms for hyperparameter optimization," in *NeurIPS'24: Proceeding of the 24th Neural Information Processing Systems*, 2011, Conference Proceedings.
- [4] S. Bhat, D. Lu, A. Kwon, and S. Devadas, "Var-cnn: A data-efficient website fingerprinting attack based on deep learning," *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 4, pp. 292–310, 2019.
- [5] D. G. Bissias, M. Liberatore, D. Jensen, and B. N. Levine, "Privacy vulnerabilities in encrypted http streams," in *Privacy Enhancing Technologies*, vol. 3856, 2006, pp. 1–11.
- [6] X. Cai, R. Nithyanand, T. Wang, R. Johnson, and I. Goldberg, "A systematic approach to developing and evaluating website fingerprinting defenses," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2014, p. 227–238.
- [7] X. Cai, X. C. Zhang, B. Joshi, and R. Johnson, "Touching from a distance: website fingerprinting attacks and defenses," in *Proceeding of the ACM Conference on Computer and Communications Security*, 2012, pp. 605–616.
- [8] O. Chapelle, J. Weston, L. Bottou, and V. Vapnik, "Vicinal risk minimization," in *Neural Information Processing Systems*, 2000, pp. 416–422.
- [9] F. Chollet et al, "Keras," <https://keras.io>, 2015.
- [10] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. P. Kuksa, "Natural language processing (almost) from scratch," *Journal of Machine Learning Research*, vol. 12, pp. 2493–2537, 2011.
- [11] A. V. Den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [12] T. Developers, "Tor metrics portal," <https://metrics.torproject.org>, 2018.
- [13] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proceedings of the 13th USENIX Security Symposium*, 2004, pp. 303–320.
- [14] K. P. Dyer, S. E. Coull, T. Ristenpart, and T. Shrimpton, "Peek-a-boo, i still see you: Why efficient traffic analysis countermeasures fail," in *Proceedings of the 33rd Annual IEEE Symposium on Security and Privacy*, 2012, pp. 332–346.
- [15] J. Hayes and G. Danezis, "k-fingerprinting: a robust scalable website fingerprinting technique," in *Proceedings of the 25th USENIX Security Symposium*, 2016, pp. 1187–1203.
- [16] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [17] D. Herrmann, R. Wendolsky, and H. Federrath, "Website fingerprinting: Attacking popular privacy enhancing technologies with the multinomial naïve-bayes classifier," in *IEEE International Conference on Cloud Computing Technology and Science*, 2009, pp. 31–42.
- [18] A. Hintz, "Fingerprinting websites using traffic analysis," in *Privacy Enhancing Technologies*, 2003, pp. 171–178.
- [19] M. Juárez, S. Afroz, G. Acar, C. Diaz, and R. Greenstadt, "A critical evaluation of website fingerprinting attacks," in *Computer and Communications Security*, 2014.
- [20] M. Juárez, M. Imani, M. Perry, C. Diaz, and M. Wright, "Toward an efficient website fingerprinting defense," in *Proceeding of the European Symposium on Research in Computer Security*, vol. 9878, 2016, pp. 27–46.
- [21] S. Kiranyaz, T. Ince, R. Hamila, and M. Gabbouj, "Convolutional neural networks for patient-specific ecg classification," vol. 2015, pp. 2608–2611, 2015.
- [22] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," pp. 1097–1105, 2012.
- [23] Y. LeCun, Y. Bengio, and G. E. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, 2015.
- [24] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, pp. 2278 – 2324, 12 1998.
- [25] M. Liberatore and B. N. Levine, "Inferring the source of encrypted http connections," in *Proceedings of the 13th ACM Conference on Computer and Communications Security*, 2006, pp. 255–263.
- [26] F. J. Morenobarea, F. Strazzera, J. M. Jerez, D. Urda, and L. Franco, "Forward noise adjustment scheme for data augmentation," in *IEEE Symposium Series on Computational Intelligence*, 2018, pp. 728–734.
- [27] A. Panchenko, F. Lanze, and M. Henze, "Website fingerprinting at internet scale," in *Proceedings of the 16th Network and Distributed System Security Symposium*, 2016.
- [28] A. Panchenko, L. Niessen, A. Zinnen, and T. Engel, "Website fingerprinting in onion routing based anonymization networks," in *Proceedings of the ACM Workshop on Privacy in the Electronic Society*, 2011, pp. 103–114.
- [29] V. Rimmer, D. Preuveneers, M. Juárez, T. Van Goethem, and W. Joosen, "Automated website fingerprinting through deep learning," in *Proceedings of the Network and Distributed System Security Symposium*, 2018.
- [30] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, p. 60, 2019.
- [31] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014.
- [32] P. Sirinam, M. Imani, M. Juárez, and M. Wright, "Deep fingerprinting: Undermining website fingerprinting defenses with deep learning," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2018, pp. 1928–1943.
- [33] P. Sirinam, N. Mathews, M. S. Rahman, and M. Wright, "Triplet fingerprinting: More practical and portable website fingerprinting with n-shot learning," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2019, pp. 1131–1148.
- [34] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv:1312.6199 [cs.CV]*, 2013.
- [35] L. Taylor and G. Nitschke, "Improving deep learning using generic data augmentation," *arXiv:1708.06020 [cs.LG]*, 2017.
- [36] V. Vapnik, *Statistical learning theory*. Wiley-Interscience, 1998.
- [37] T. Wang, X. Cai, and I. Johnson, Roband Goldberg, "Effective attacks and provable defenses for website fingerprinting," in *Proceedings of the 23rd USENIX Security Symposium*, 2014, pp. 143–157.
- [38] T. Wang and I. Goldberg, "Improved website fingerprinting on tor," in *Proceedings of the ACM Workshop on Privacy in the Electronic Society*, 2013, pp. 201–212.
- [39] —, "On realistically attacking tor with website fingerprinting," in *Proceedings on Privacy Enhancing Technologies (PoPETs)*, 2016, pp. 21–36.
- [40] —, "Walkie-talkie: An effective and efficient defense against website fingerprinting," in *Proceeding of the 26th USENIX Security Symposium*, 2017, pp. 1375–1390.
- [41] J. Wei and K. Zou, "Eda: Easy data augmentation techniques for boosting performance on text classification tasks," in *International Joint Conference on Natural Language Processing*, 2019, pp. 6381–6387.
- [42] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," 2016.
- [43] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *International Conference on Learning Representations*, 2018.