Filtering DDoS Attacks from Unlabeled Network Traffic Data Using Online Deep Learning

Wesley Joon-Wie Tann, Jackie Jin Wei Tan, Joanna Purba, Ee-Chien Chang

National University of Singapore

{wesleyjtann, jackietan, joanna_sara}@u.nus.edu, changec@comp.nus.edu.sg

ABSTRACT

DDoS attacks are simple, effective, and still pose a significant threat even after more than two decades. Given the recent success in machine learning, it is interesting to investigate how we can leverage deep learning to filter out application layer attack requests. There are challenges in adopting deep learning solutions due to the ever-changing profiles, the lack of labeled data, and constraints in the online setting. Offline unsupervised learning methods can sidestep these hurdles by learning an anomaly detector N from the normal-day traffic N . However, anomaly detection does not exploit information acquired during attacks, and their performance typically is not satisfactory. In this paper, we propose two approaches that utilize both the historic N and the mixture M traffic obtained during attacks, consisting of unlabeled requests. First, our proposed approach, inspired by statistical methods, extends an unsupervised anomaly detector N to solve the problem using estimated conditional probability distributions. We adopt transfer learning to apply N on N and M separately and efficiently, combining the results to obtain an online learner. Second, we formulate a specific loss function more suited for deep learning and use iterative training to solve it in the online setting. On publicly available datasets, such as the *CICIDS2017*, our online learners achieve an average of 90.6% accuracy rates compared to the baseline detection method, which achieves around 60.0% accuracy. In the offline setting, our approaches on unlabeled data achieve competitive accuracy compared to classifiers trained on labeled data.

CCS CONCEPTS

• Security and privacy → Denial-of-service attacks.

KEYWORDS

Network security; DDoS mitigation; Online machine learning

ACM Reference Format:

Wesley Joon-Wie Tann, Jackie Jin Wei Tan, Joanna Purba, Ee-Chien Chang. 2021. Filtering DDoS Attacks from Unlabeled Network Traffic Data Using Online Deep Learning. In *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security (ASIA CCS '21)*, June 7–11, 2021, Hong Kong, Hong Kong. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3433210.3453083>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASIA CCS '21, June 7–11, 2021, Hong Kong, Hong Kong

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8287-8/21/06...\$15.00

<https://doi.org/10.1145/3433210.3453083>

1 INTRODUCTION

Distributed Denial-of-Service (DDoS) attacks are well established as a significant threat to our present-day Internet network security, denying legitimate users access to shared and essential resources. The earliest reported DDoS attack in 1999 [2] started a wave of denial-of-service attacks that are distributed in nature. Even though more than two decades have passed, these attacks—simple to set up, difficult to stop, and very effective—are more popular than ever as they remain potent [28]. The allocation of extra resources is a typical protective approach to handle such DDoS traffic. However, catering excess resources goes hand in hand with additional costs. There is a growing urgency to find practical and inexpensive methods, which can effectively abate disruptions by filtering malicious traffic during an attack. We focus on devising such a filtering system.

Existing DDoS filtering methods are generally statistical-based or machine learning-based detection mechanisms. On the one hand, statistical methods typically involve some fast calculation of network traffic properties, such as entropy scoring of network packets [6, 11, 26, 33], applied to IP filtering [19, 25] and IP source traceback [13, 34]. On the other hand, deep learning, which has been successfully applied to the detection domain in recent years, exploits intrinsic traffic properties for better detection. However, there are significant challenges: (1) The attacks are often domain-specific with varying characteristics, making it hard for filtering methods to adapt on time. (2) Time constraint is a consideration in the real-time environment, and online models have to be computationally efficient to respond to requests rapidly during an attack. (3) The lack of labeled requests at the time of the attacks, required for supervised learning algorithms, poses another challenge for learning models.

Most deep learning approaches for DDoS in the literature adopt an offline learning approach. The models are trained on labeled data. The training is conducted much earlier (e.g., at least a few days) before the attack commences [14, 15, 23, 31, 36]; thus, they are not applicable in adaptive online scenarios. There are only a select few works on online learning that performs incremental updates as more attack requests arrive. Çakmakçı et al. gave an up-to-date survey of online DDoS detection [4]. Most online learning methods employed unsupervised clustering (e.g., k-means) on the attack day data [4, 7, 18].

In this paper, we explore how deep learning is leveraged to enhance detection performance, especially in the application-layer traffic, which exhibits intrinsic statistical properties. We utilize data from two operational periods, normal-day traffic N , and the attack-day traffic M that contains a mixture of unlabeled legitimate and malicious traffic. We propose two approaches with two different objectives.

Figure 1: A framework that utilizes both N and M to rank network traffic based on their final n_x/d_x scores. The higher the final score, the more likely the traffic is normal. To achieve computation speedup, the embedding of N is transferred to D .

In the first approach, we employ techniques inspired by the statistical method PacketScore [32], which takes a principled approach to estimate conditional probabilities (the likelihood that a request is an attack given that it is observed in the mixture traffic) of network requests. We perform the learning in a two-step process. (1) We apply unsupervised learning on the normal-day traffic N to obtain a model N that learns the distribution P_n of N . (2) When an attack commences, a similar unsupervised learning model D rapidly learns the distribution P_d of the mixture M . Now, the conditional probability can be estimated using the ratio $P_n(x)/P_d(x)$. We call this approach *N-over-D* nets, as illustrated in Figure 1.

An advantage of the first approach is that existing unsupervised learning models, after customization, can be “re-used” for N and M . In our construction, both models N and D are LSTM recurrent neural networks, designed for application-layer DDoS, where a “request” consists of a sequence of events (e.g., an HTTP GET to the landing page, followed by subsequent HTTP GET requests). To meet the online requirements, we design model D to achieve quick online updates by relying on transfer learning. When an attack occurs, the optimized embedding of N is transferred to D , significantly reducing the training needed for D .

However, there are two main challenges of *N-over-D* nets. First, it is not straightforward to customize an unsupervised learning model to fit into *N-over-D* nets. Typical learning models normalize data values for effective training and smooth gradient flow, and even though the order of values is maintained, the resulting model-calculated probabilities are not exact, leading to the distortion and swapping of the order of ratio $P_n(x)/P_d(x)$. Second, models N and D are separately trained during each operational period. There is no joint training of the models, and we do not fully leverage all available information for each model.

In the second approach, we formulate a machine learning loss function specifically to address these challenges. This approach

requires an estimate of the proportion α of attack requests in M . We formulate a learning objective for a two-class classifier that attempts to classify samples from N as normal, and α proportion of samples in M as attacks. We design an online iterative two-class classifier for this particular loss function that jointly trains on both normal-day N and mixture M traffic. We train the classifier with a similar LSTM model as in *N-over-D* nets. The enhanced iterative classification approach achieves remarkable detection performance. Although the iterative classifier takes a slightly longer training time than *N-over-D* (3–5 times longer), it is nevertheless suitable for attacks with a short timeframe; even used for classifying unlabeled historic M .

Comprehensive experiments demonstrate the ability of our approach to mitigate a range of DDoS attacks. For example, in the *CICIDS2017* dataset, while the N model achieves accuracy and a false-positive rate of (32.5%, 69.97%), both the online *N-over-D* and iterative classifier models accomplish marked improvements of (93.3%, 0.47%) and (93.5%, 3.40%), respectively. Moreover, both the *N-over-D* and the iterative classifier outperform existing online detection methods [4, 5, 18] by a large margin in either accuracy or false-positive rates (see Tables 3 and 5 for experiment details). In the offline setting, even trained on the unlabeled data, the iterative method attains competitive performance compared to two-class classifiers trained on labeled data (see Tables 5 and 6). Although the iterative method requires an estimate of α , our experiments show that the method is robust to inaccurate estimates (see Table 8).

Contribution.

- (1) We propose a learning objective that takes a conditional probability approach inspired by the classical statistical method [32]. Our learning method takes the ratio between two probability distributions of normal-day N traffic P_n and mixture M traffic P_d , ranking each request x by $P_n(x)/P_d(x)$, where a higher score indicates normality.
- (2) We design a specific loss function and propose another learning objective that is more suitable for deep learning. We also introduce an iterative online method for this loss function, which accurately predicts requests from N as normal, and α proportion of requests in M as attacks; the iterative online method achieves better performance.
- (3) To realize our proposed approach, we design an LSTM recurrent neural network to take input requests consisting of sequences of traffic, specifically for application-layer DDoS attacks.

Organization. We organize the rest of the paper by first presenting the background information and challenges of the DDoS problem in Section 2. We introduce the problem in Section 3 and propose an appropriate conditional probability learning approach in Section 4. We then detail the iterative classification approach in Section 5 and the online LSTM architecture of our approaches in Section 6. In Section 7, we describe the datasets and evaluate our approaches on these datasets, comparing them with state-of-the-art methods. We also analyze the performance of our mitigation process and discuss several practical issues. Finally, we summarize related work in Section 8 and conclude the paper in Section 9.

2 BACKGROUND AND CHALLENGES

This section gives an overview of application layer network services and describes the challenges of detecting DDoS attacks in such systems. We also motivate the need for innovative defenses that can perform online network traffic monitoring.

2.1 Application Layer DDoS

A surge in application-layer DDoS attacks, using a network of widely distributed botnets, are increasingly disruptive to current network services [28]. Symantec's telemetry shows that it is often small and medium-sized retailers, selling goods ranging from clothing to gardening equipment to medical supplies, that are on the receiving end of the majority of attacks [29]. It is a serious global issue that could potentially affect any business that operates online, especially the essential services during critical times, such as the COVID-2019 pandemic. Observed in Q1 2020, there is an unexpected spike in the number of attacks; noticeable changes in the distribution of DDoS attacks by type doubled by 80% against Q1 2019 [12].

2.2 Existing Online DDoS Filtering

Several unsupervised learning mitigation strategies [4, 7, 18] have been proposed to counter such attacks in an online manner. These learning-based methods, employing rudimentary machine learning algorithms, e.g., random forest, logistic regression, decision tree (see Sec. 8.2 for details), perform clustering and simple distance measure using a few elementary features. While they can capture some basic network traffic properties, they are unable to learn deeper representations of legitimate and malicious requests to better distinguish between the two types of traffic. Moreover, the optimization goal of these methods is not clear.

2.3 Flooding Attacks

Common application-layer flooding attacks such as HTTP flood attacks are designed to generate an overwhelming amount of traffic, increasing the servers' load to severely impact a target. A popular tactic used by attackers to disrupt legitimate access is to disguise their traffic as flash crowds. A flash crowd is a surge in traffic to a particular Web site that causes the site to be virtually unreachable [35]. Even though attacks and flash crowds originate from different motivations, such attacks' covertness makes them hard to detect. These attacks typically send numerous HTTP packets abruptly and rapidly, flooding a target server with harmless HTTP requests, evading most defense systems. Any defense system designed to detect and mitigate such attacks must differentiate between attack and normal traffic fast, minimize losses, and perform online mitigation.

3 TWO-CLASS LEARNING ON UNLABELED MIXTURE TRAFFIC

In this section, we formalize the learning problem in the DDoS setting. We consider two operational periods, a normal period and an attack period. In practice, the attack period can be readily detected as it is characterised by an unusual high volume of requests. Let $\mathcal{N}$ be the requests logged during the normal period, and $\mathcal{M}$

the requests during the attack period. The mixture $\mathcal{M}$ contains a mixture of both legitimate requests and attack requests. We assume that based on the volume, we have a rough estimate of α , which is the proportion of attack requests in $\mathcal{M}$. While $\mathcal{M}$ is unlabeled, all requests in $\mathcal{N}$ are normal. Given $\mathcal{N}$, $\mathcal{M}$, and α , the goal is to filter out α proportion of attacks from $\mathcal{M}$.

Offline vs Online Setting. The online setting consists of three stages.

- In the *pre-processing* stage, computational intensive processing can be performed on $\mathcal{N}$, giving some intermediate model and representation $\tilde{\mathcal{N}}_0$.
- The *on-line learning* stage is carried out when the attack commences. The time is divided into short intervals of ℓ minutes (e.g., $\ell = 1$), and the incoming mixture $\mathcal{M}$ is also ordered accordingly $(\mathcal{M}_1, \mathcal{M}_2, \dots)$, where $\mathcal{M}_i$ contains the mixture within the i -th interval. During the $(i + 1)$ -th interval, training is performed on $\mathcal{M}_i$ and some intermediate representation $\tilde{\mathcal{N}}_i$, giving the updated $\tilde{\mathcal{N}}_{i+1}$ and a model C_{i+1} . There are stringent computing resource constraints, e.g., the training should be completed within ℓ minutes.
- In the *filtering* stage, the trained model C_{i+1} is deployed to filter the incoming requests. To meet the real-time requirement, instead of applying the model on each request, it is applied to the mixture to obtain an updated blacklist of attack identities (e.g., IP addresses). The blacklist is then activated to filter out the attack requests. For example, C_{i+1} can also be applied on $\mathcal{M}_i$ to obtain the blacklist.

In the offline setting, the training can utilize the full $\mathcal{M}$ and $\mathcal{N}$, and there is no real-time requirement.

DDoS in the application layer. A request in network-layer DDoS such as Syn flood typically consists of a single packet. In contrast, a request in application-layer DDoS, which is the focus of this paper, could be a sequence $\mathbf{x} = (x_1, x_2, \dots, x_T)$ of sub-requests x 's. For instance, a visit to a web-service could consist of an HTTP-GET to the landing page, followed by others. The intrinsic characteristic of such a sequence could be learned to differentiate attacks. In our notation, we view each request $\mathbf{x}$ as a sequence of sub-requests, and a user corresponds to an IP-address. In evaluating performance, instead of measuring the actual byte counts, the reported measures are calculated based on the number of users and IP addresses correctly classified.

4 N-OVER-D APPROACH

We propose the N-over-D approach to extend a one-class unsupervised learning anomaly detection to estimate conditional probabilities in a three-step process. A one-class learning model is trained on a training set containing samples of the same class, giving a predictive model that, on input $\mathbf{x}$, output the probability of having $\mathbf{x}$ in that class.

- (1) During pre-processing, one-class learning is applied to $\mathcal{N}$ to obtain a model N that predicts the distribution P_n .
- (2) During online learning, a similar process is applied to $\mathcal{M}$ to obtain a model D that predicts P_d .

- (3) During filtering, the predicted probabilities, $P_n(\mathbf{x})$ and $P_d(\mathbf{x})$, are combined to give a model that predicts $P_a(\mathbf{x})$, the likelihood that $\mathbf{x}$ is from the attack given that $\mathbf{x}$ is observed in the mixture.

The last step calculates $P_a(\mathbf{x})$ from $P_n(\mathbf{x})$ and $P_d(\mathbf{x})$, following similar techniques in PacketScore [32]. Note that,

$$\begin{aligned} P_a(\mathbf{x}) &= 1 - P(\mathbf{x} \text{ is normal} \mid \text{request } \mathbf{x} \text{ is observed in mixture}) \\ &= 1 - \frac{(1 - \alpha)P_n(\mathbf{x})}{\alpha P_d(\mathbf{x}) + (1 - \alpha)P_n(\mathbf{x})} \\ &= \frac{1}{1 + \left(\frac{1 - \alpha}{\alpha}\right) \frac{P_n(\mathbf{x})}{P_d(\mathbf{x})}} \end{aligned} \quad (1)$$

where α is the proportion of attack requests in the mixture data $\mathcal{M}$. From $P_a(\cdot)$, we can classify $\mathbf{x}$ as an attack when $P_a(\mathbf{x}) > 0.5$.

Moreover, to facilitate fast online learning, we could deploy transfer learning from N to D so that we can train D readily. We design model D to achieve quick online updates. When an attack occurs, the optimized embedding of N is transferred to D , thereby significantly reducing the training needed from a random initialization of D .

Proportion α . Although the proportion α of attack requests is used in Equation 1, the role of α is not significant in the N-over-D approach. The performance does not rely on the accuracy of the estimated α . Typically, during filtering, a threshold is set so that a certain predetermined percentage of requests would be allowed. It is tolerable to accept attack requests as long as the system can handle the load. Such a threshold can be selected by sorting probabilities $P_a(\cdot)$ of the received requests and determining the appropriate cutoff. Since only the ranking matters, we can ignore the constant terms and focus only on the ratio P_n/P_d , which does not depend on α .

Choice of one-class classification. Not all one-class classification is applicable in the N-over-D approach. The candidate should meet the following two requirements. (1) Firstly, Equation (1) assumes that the output score of the model is the predicted probability. Nevertheless, many models transform the probability and output a score that preserves the ranking but distorted the actual value. The ratio of such distorted values from the two models might not preserve the ratio ranking. Hence, it is essential that the score should be the actual probability, or with some additive and multiplicative constants. (2) Secondly, as indicated earlier, we need a mechanism to transfer learning from N to D .

Limitations. The N-over-D approach has a limitation in that it relies on the accuracy of the one-class classification, and the classification is applied to the normal and mixture almost independently (except some information that is transferred from N to D via transfer learning). Intuitively, more accurate results could be attained via joint-learning. For instance, some combination of features values could have few samples from N but a large population from $\mathcal{M}$. When applied independently, N might overgeneralize those features and predict a higher probability. However, with the additional information that there is indeed a high population from $\mathcal{M}$, a more refined and accurate boundary could be learned.

5 ITERATIVE CLASSIFIER APPROACH

In this section, to address the limitations of the N-over-D nets, we propose a method that optimizes a specifically formulated loss function designed for deep learning, which directly trains a two-class classifier on both $\mathcal{N}$ and $\mathcal{M}$.

Guided by the intuition that the two-class classifier should classify samples from $\mathcal{N}$ as normal, and α proportion from $\mathcal{M}$ as attacks, we define the following loss function.

$$\begin{aligned} \mathcal{L} = & -\frac{1}{|\mathcal{N}|} \sum_{y \in \mathcal{N}} \log(1 - p(y)) \\ & -\frac{1}{\alpha|\mathcal{M}|} \sum_{y \in \text{Sel}_\alpha(\mathcal{M})} \log(p(y)) \\ & -\frac{1}{(1 - \alpha)|\mathcal{M}|} \sum_{y \in \mathcal{M}, y \notin \text{Sel}_\alpha(\mathcal{M})} \log(1 - p(y)) \end{aligned} \quad (2)$$

where $p(y)$ is the model-predicted probability of the observed input y being the attack, and $\text{Sel}_\alpha(\mathcal{M})$ contains $\alpha|\mathcal{M}|$ samples from $\mathcal{M}$ with the highest model-predicted probability $p(\cdot)$. The set $\text{Sel}_\alpha(\mathcal{M})$ can be obtained by ranking the samples in $\mathcal{M}$ according to the predicted probabilities and selecting the highest $\alpha|\mathcal{M}|$. In other words, the three summations correspond to the three sets $\mathcal{N}$, $\text{Sel}_\alpha(\mathcal{M})$, and $(\mathcal{M} - \text{Sel}_\alpha(\mathcal{M}))$. The loss function favors the first and third set to be classified as “normal”, and the second set to be classified as “attack”.

We employ an iterative method to train such a binary classification model. Let the trained classifier C_i be the model obtained after the i -th iteration.

- (1) At each interval ℓ during an attack, we assume all traffic is malicious.
- (2) We then train classifier C_i on the newly observed data to update the model C_{i+1} .
- (3) Using classifier C_{i+1} , we predict and rank the observed data, selecting the top predicted attacks to retrain the classifier; by performing this procedure iteratively, the classifier directly minimizes our defined loss function (Equation 2).

The iterative classifier is able to perform fast online learning because it uses one supervised learning model, and it is suitable for attacks with a short timeframe. We assume all requests during the normal-day traffic as legitimate and all the traffic during an attack period, consisting of a mixture of both normal and attack traffic, as malicious and label them accordingly. By making such assumptions, we construct the initial set of inexactly labeled training data.

Next, we train the classifier using these inexact labels. We then predict the same initial training data set with this trained classifier, retain a certain top A percentage of traffic predicted as malicious to form a reduced attack dataset, and construct the next set of training data. Ideally, we would like to retain the α proportion of the predicted attacks, but in the early iterations of training, the predicted values are not accurate. Hence, to increase the confidence of likely attacks, A should be some value smaller than α ¹. By iteratively performing this prediction and reducing attack data, it allows the iterative classifier to repeatedly improve its model performance by

¹We choose $A = 40\%$ in the experiments

increasingly retaining the traffic during an attack that is most likely to be truly malicious.

6 LSTM N-OVER-D

Figure 2: An overview of the LSTM-based N-over-D.

This section realizes the N-over-D approach by using LSTM-based unsupervised learning models. Recall that the N-over-D approach consists of two “partners”—a normal-day model N trained using normal-day dataset $\mathcal{N}$ and a model D trained using the mixture dataset $\mathcal{M}$. Here, we adopt LSTM-based unsupervised learning (see Figure 2) to obtain N and D and provide the details of the input features, model architectures, and design choices below (Sections 6.1 to 6.3).

6.1 Sequence Data Preparation

The data processing procedure for online learning widely differs from that of a static dataset of requests. Hence, we prepare the data into a streaming form; sequences of sub-requests, that are grouped by source and destination IP addresses.

Let τ_0 be the starting interval of the traffic and τ_i are the following ℓ -min intervals, where τ_1 is the first interval, τ_2 the second, and so forth. By extracting the intervals of network traffic into arrays, we prepare the data for sequential online learning, which allows the LSTM to learn the distribution of both normal and DDoS attack traffic. In the rest of this section, we describe in detail each stage of the preparation process (see Appendix for details of the data preprocessing procedure algorithm).

During each interval τ_i , the preparation procedure captures network traffic that is collected for an allocated amount of time ℓ . The request sequences are then constructed using the sub-requests gathered during the interval. We give an example of a sequence of sub-requests (see Table 1). The example sequence consists of T sub-requests with attributes, consisting of both dynamic and static attributes. In our case, we chose eight dynamic attributes that capture the information that varies from request to request and two static attributes that are most common for all requests for a sequence. Sequences longer than T are split to form a new sequence, while sequences shorter than T are zero-padded at the end.

6.2 Model for Normal Traffic

The model N for normal-day traffic is made up of two particular components: (1) a probabilistic sequential learning model that

learns to predict sequences of temporal features, and (2) a histogram frequency distribution approximate representation to model the static features extracted from the network traffic. We model N using a parameterized LSTM network θ_n and a Histogram method r_n .

LSTM. At each interval τ_i , the LSTM θ_n takes as input the sequence of requests $\mathbf{x} = (x_1, x_2, \dots, x_T)$ that are the temporal features in the sequences. We set $T = 200$, balancing between efficiency of the model and memory requirements to learn the data. The LSTM learns from inputs at each step to output the predictions $(\hat{x}_1, \hat{x}_2, \dots, \hat{x}_T)$; the training process is illustrated in Figure 3, and the hyperparameters are detailed below.

Figure 3: The LSTM model architecture.

Input layer: Recall in Section 6.1 that we preprocess the requests into intervals of sequences, simulating a real-time data streaming process. By creating an environment suitable for sequential online learning, it enables the LSTM to leverage its modeling capabilities to learn the correlation between traffic requests in the same sequence. This layer takes as inputs the data sequences $\mathbf{x}$ of each interval τ_i , represented as a 3D tensor $T \times m \times l$, where T is the length of the sequences and each x_i is represented as a $m \times l$ matrix. The matrix contains m samples and l features.

Embedding layer: The inputs $\mathbf{x}$ are then passed to the embedding layer emb_n , which creates a dense matrix that is learned through backpropagation of the training process. This layer reduces all features into a dense vector representation of dimension 16. The embedding, a dense matrix in a linear space, achieves two important functions. First, by using an embedding with a much smaller dimension than the feature size, it reduces the dimension of feature representations to the embedding size and reduce the model complexity. Second, the input feature embedding groups semantically similar features in an Euclidean space, and allow for a better feature representations.

LSTM layer: The outputs from the embedding layer then form the compact inputs to the LSTM layer, which is a powerful machine learning architecture to learn from sequence data. LSTMs are designed to predict the presented observations at each timestep sequentially. By taking a sequence $\{x_1, x_2, \dots, x_T\}$ as input, the LSTM constructs a corresponding sequence of hidden state representations $\{h_1, h_2, \dots, h_T\}$. A single-layer LSTM uses the hidden

Table 1: Details of the eight dynamic and two static features in a sequence sample.

Req #	Absolute Time	Request Len	IP Flags	TCP Len	TCP Ack	TCP Flags	TCP Window Size	Highest Layer (Protocol)	Extra Info	List of Protocols
1	23352.0	66.0	0x00004000	0.0	19.0	0x00000012	14600.0	TCP	443 → 8918 [SYN, ACK] Seq=0 Ack=1 Win=14600 Len=0 MSS=1460 SACK_PERM=1 WS=512	eth:ethertype:ip:tcp
2	23352.1	60.0	0x00004000	0.0	30.0	0x00000010	15872.0	TCP		
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮		
$t-1$	23368.5	85.0	0x00004000	31.0	42.0	0x00000018	16896.0	TLSv1.2		
t	23368.7	60.0	0x00004000	0.0	42.0	0x00000011	16896.0	TCP		
$t+1$	0	0	0	0	0	0	0	0		
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮		
T	0	0	0	0	0	0	0	0		

representations $\{h_1, h_2, \dots, h_T\}$ for estimation and prediction. In our case, we use two LSTM layers, where the second layer uses the hidden states of the previous layer as inputs. Every hidden state in each layer performs memory-based learning to remember relevant features using previous inputs. Previous hidden states and current inputs are transformed into a new hidden state, and it is achieved through a recurrent operator that takes in (h_{t-1}, x_t) , such as:

$$h_t = \tanh(W_h h_{t-1} + W_x x_t + b) \quad (3)$$

where W_h , W_x , and b are parameters of the layer and $\tanh(\cdot)$ represents the standard hyperbolic tangent function.

The LSTM is made up of these cells that are specifically designed for series data to trace the history from previous network requests. In order to overcome both the vanishing gradient and long-term dependency issues, each LSTM cell retains an internal cell state c_t and a hidden state that is the output. Both c_t and h_t are computed via three gate functions to retain both long and short term storage of information. The forget gate f_t , input gate i_t , and output gate o_t control the flow of information. In this work, we use the LSTM without peep-hole connections to handle complex sequences with long-range structure. We design each cell to consist of 32 units, the dimensionality of the output space, which is passed to the prediction layer (see Appendix for details).

Prediction layer: Following the process (see Figure 3), the LSTM takes in x_t at each step t and makes a prediction $\hat{x}_t$ of the request at the next step $t+1$. We set the output function as the softmax function:

$$\text{Softmax}(x_t) = \frac{\exp(x_t)}{\sum_j \exp(x_j)} \quad (4)$$

where j is the total number of possible outputs. The softmax function converts a real vector to a vector of categorical probabilities, where the elements of the output vector are in range $(0, 1)$ and sum to one. We can interpret the results as a probability distribution and take the maximum as the prediction $P(x_{t+1} | x_0, \dots, x_t)$ of the next step in the sequence, where x_0 is the start of sequence symbol. We start all sequences with a this symbol and the probability of observing it at step 0 is one.

Hence, by assuming that the state of each time step depends on previous time steps where the current information has a dependency on previous data, we can use standard conditional probability theory to determine the joint probability of each sequence of requests by calculating:

$$P(\mathbf{x}) = P(x_0) \cdot P(x_1 | x_0) \cdots P(x_T | x_0, \dots, x_{T-1}) \quad (5)$$

$$= P(x_0, x_1, \dots, x_T)$$

and assigning this computed probability as the *LSTM prediction* $\theta_n(\mathbf{x})$.

When training the LSTM network θ_n , our objective is to minimize the difference between the model prediction and observed request at each timestep t . The loss function we minimize is defined as:

$$\mathcal{L} = - \sum_{c=1}^M y_{t,c} \log(\hat{x}_{t,c}) \quad (6)$$

where M is the total number of possible next request, $y_{t,c}$ is the indicator if ground-truth label c is the correct prediction for observation $x_{t,c}$ in the t -th step of the sequence, and $\hat{x}_{t,c}$ is the LSTM predicted probability of the input at time t . As the possibilities M can be exceedingly large, it could cause an explosive increase in model computational complexity whereby drastically increasing the time required to train the model. A large M presents a potential model robustness issue in large datasets. Moreover, there could be other considerations, such as memory limitations and GPU computational constraints. Hence, we employ a feature hashing trick [27] to reduce M and solve the mentioned issues of a large number of request possibilities.

Histogram. For each sequence of unique source and destination IP addresses, the static features are inputs for the histogram model r_n . The Extra Info and List of Protocols attributes are the responses of the server after receiving requests and the list of network protocols used. They generally remain constant throughout the entire sequence. Since the two features are observed to be static, we assume independence between them and the dynamic features.

A histogram model is used to capture the approximate frequency distribution of the static attributes. We count the number of specific Extra Info and List of Protocols pairs, and calculate the histogram probability prediction of each pair:

$$r_n(\mathbf{x}) = \frac{m_i(\mathbf{x})}{\sum_{i=1}^k m_i} \quad (7)$$

where $\mathbf{x}$ is a static feature pair, m_i is a function that counts the number of observations of a particular pair, and k is the total number of static feature pairs.

Model N prediction. To predict how likely a sequence belongs in the normal traffic, we incorporate both the *LSTM prediction* $\theta_n(x)$ and the *Histogram prediction* $r_n(x)$. By taking the product of the two predictions using dynamic features and static features:

$$\begin{aligned} P_n(x) &= N \text{ LSTM}(x) \times N \text{ Hist.}(x) \\ &= \theta_n(x) \times r_n(x) \end{aligned} \quad (8)$$

we obtain the probability of a sequence of network requests, which serves as the numerator of the N-over-D framework.

6.3 Model for Mixture Traffic

The model D , similar to the normal-day model N , also consists of a parameterized LSTM and a Histogram model that learns to predict both sequences of dynamic features and static features of the traffic, respectively. The goal of D is to learn meaningful generalization of the mixture traffic during DDoS to predict the probability of each network request to be malicious. It is designed to perform fast online learning, updating the models on the evolving attack traffic every minute.

The model architecture is designed with a significant difference to promote the online learning design. We initialize the embedding layer emb_d with the previously trained embedding emb_n during normal traffic. Since the embedding is a compact representation of the input network requests, we considerably minimize the training of θ_d by transferring the learned embedding layer rather than initializing emb_d randomly. It significantly reduces the number of model parameters, which decreases the computational cost, to achieve fast learning and inference during attacks. In addition, we employ a similar Histogram model r_d on the static features of network requests during the DDoS attacks.

Model D prediction. To predict how likely a sequence belongs in the malicious traffic, we incorporate both the *D LSTM prediction* $\theta_d(x)$ and the *D Histogram prediction* $r_d(x)$. By taking the product of the two predictions using dynamic features and static features:

$$P_d(x) = \theta_d(x) \times r_d(x) \quad (9)$$

we obtain the probability of a sequence of network requests during DDoS attacks, which serves as the denominator of the N-over-D framework.

6.4 N-over-D Likelihood Measure

N-over-D performs by first learning N to model the normal traffic data distribution and predict the likelihood $P_n(x)$ that the presented network request sequences are legitimate. Next, during DDoS attacks, the online-learner D is trained for accelerated prediction $P_d(x)$ of the incoming DDoS traffic to estimate the probability that each sequence is malicious. Finally, we adopt an integrated approach to decide the normality of a sequence by taking the division:

$$\frac{N}{D}(x) = \log \left(\frac{P_n(x)}{P_d(x)} \right) \quad (10)$$

Since a high $P_n(x)$ indicates high legitimacy and a low $P_d(x)$ suggests a small chance of maliciousness, the higher the joint probabilities $\frac{N}{D}(x)$, the more likely the network traffic is normal. Hence,

we differentiate between normal and attack traffic by ranking the combined probabilities of each sequence in the DDoS traffic of each 1-min interval, and discarding request sequences with relatively low joint probabilities.

7 EXPERIMENTS

In this section, we present the evaluation of our approaches performed on various DDoS attacks. We demonstrate the effectiveness of our models (Figure 4 and Table 3) and compare them with state-of-the-art methods (Tables 5 and 6).

7.1 Datasets

To validate the N-over-D and Iterative Classifier, we use three separate volumetric DDoS attacks from two publicly available network datasets, (1) CICIDS2017 [24] dataset², which contains generated traffic resembling real-world attacks, and (2) CAIDA UCSD “DDoS Attack 2007” [8] dataset³ that contains anonymized traffic traces from a DDoS attack on August 4, 2007. A summary (Table 2) of the test data during attacks, shows the actual proportion α of attack traffic.

Table 2: Test data samples. Both *HULK* and *LOIC* attacks are from the *CICIDS2017* benchmark [24], and the 2007 Attack is from *CAIDA07* real-world [8] dataset.

Dataset	# Total	# Normal	# DDoS	DDoS (%)
<i>HULK</i>	11710	4552	7158	61.1
<i>LOIC</i>	2447	1069	1378	56.3
<i>CAIDA07</i>	16310	3262	13048	80.0

7.1.1 Benchmark Data. We evaluate our approach with one of the latest benchmark data *CICIDS2017* [24]. It is a synthetically generated dataset produced by the Canadian Institute for Cybersecurity of the University of New Brunswick (UNB), Canada. The labeled data consists of the most up-to-date common attacks of network activity, including two volumetric attacks, one on Wednesday and one on Friday. The traffic traces provided in the *pcap* format include full packet payloads, data labels, and traffic flow statistics. The types of attacks and the times they occur are listed below:

- (1) *HULK* (Wed 10:43 a.m. – 11 a.m.)
- (2) Low Orbit Ion Canon (*LOIC*) (Fri 3:56 p.m. – 4:16 p.m.)

The attacks last for 17 and 20 mins, respectively. In order to obtain normal traffic to learn the meaningful generalization of legitimate users, we retrieve traffic data prior to the attacks. For attack (1), we use normal network traffic on Wednesday (9:00 a.m. – 9:47 a.m.), and for attack (2), we take Friday (3:00 p.m. – 3:56 p.m.) data to train the model N . In each attack, after training N on the respective normal traffic, we perform online learning and detection using D when an attack occurs, demonstrating progressive improvement of the detection performance. We leave the last five intervals as the test set, to evaluate our models and compare with state-of-the-art methods.

²www.unb.ca/cic/datasets/ids-2017.html

³www.caida.org/data/passive/ddos-20070804_dataset.xml

Figure 4: The x-axis is the rejection ratio (0–1), and the y-axis is the corresponding false rejection rate. If the rejection ratio is at 0.8, it means that 80% of the requests are rejected, and the graph shows the rate of false positives.

Table 3: Detection performance on the various attacks in *CICIDS2017* and *CAIDA07* datasets.

Dataset	Model	ACC	FPR	Prec.	Rec.	F1
HULK	N	0.5903	0.3837	0.7016	0.5739	0.6313
	N-over-D ($\ell = 1$)	0.8086	0.1030	0.9198	0.7524	0.8277
	N-over-D ($\ell = \infty$)	0.8291	0.0766	0.9403	0.7692	0.8462
	Iterative Classifier ($\ell = 1$)	0.9088	0.0378	0.8831	0.7659	0.8204
	Iterative Classifier ($\ell = \infty$)	0.9120	0.0206	0.9299	0.7313	0.8187
	Full Classifier	0.9350	0.0281	0.9808	0.9116	0.9449
LOIC	N	0.3257	0.6997	0.3889	0.3454	0.3659
	N-over-D ($\ell = 1$)	0.9330	0.0047	0.9959	0.8846	0.9370
	N-over-D ($\ell = \infty$)	0.9330	0.0047	0.9959	0.8846	0.9370
	Iterative Classifier ($\ell = 1$)	0.8805	0.0329	0.8808	0.6494	0.7476
	Iterative Classifier ($\ell = \infty$)	0.9486	0.0220	0.9016	0.8275	0.8630
	Full Classifier	0.9713	0.0037	0.9968	0.9121	0.9526
CAIDA07	N	0.6009	0.2639	0.8958	0.5671	0.6945
	N-over-D ($\ell = 1$)	0.6944	0.0251	0.9900	0.6242	0.7657
	N-over-D ($\ell = \infty$)	0.6880	0.0411	0.9837	0.6202	0.7608
	Iterative Classifier ($\ell = 1$)	0.7081	0.0208	0.9919	0.6403	0.7783
	Iterative Classifier ($\ell = \infty$)	0.7088	0.0141	0.9945	0.6396	0.7785
	Full Classifier	0.7094	0.0120	0.9953	0.6398	0.7789

7.1.2 Real-world Data. The other dataset, *CAIDA07* UCSD “DDoS Attack 2007” [8], is a real-world attack collected from various locations worldwide. It occurred on 04 August, 2007. The data contains approximately one hour (20:50 UTC to 21:56 UTC) of anonymized traffic traces from a DDoS attack that includes types of network traffic such as Web, FTP, and Ping. We use all traffic traces before the attack (20:50 – 21:12) as normal network traffic to train the model N . When the attack starts at 21:13, we perform both online and offline learning of D on the first 4 mins; the next min of the attack is left as the test set to evaluate our models. The summary (Table 2) of the *CAIDA07* test data shows 80% attack traffic.

7.2 Evaluation Results

We present a detailed evaluation of the proposed approaches on both the *CICIDS2017* and *CAIDA07* datasets. Setting the rejection rate at 50% and online interval ℓ at 1-min, we measure the performance of iterative classifier ($\ell = 1$) and N-over-D ($\ell = 1$). We also set the interval ($\ell = \infty$) as all the data during attack for both approaches, and present the Accuracy (ACC), False Positive Rate (FPR), Precision (Prec.), Recall (Rec.), and F1 Score (F1) results. Figure 4 and Table 3 illustrate these results. As demonstrated, both N-over-D and iterative classifier achieve low error rates and high accuracy scores across various attacks. These results validate the robust design of our approach.

For comparison, we design another classification model that trains on the entire training data of each attack using actual ground

truth information. In this setting, we suppose that all normal and attack data is available during model training and that the labels accurately distinguish the attack from the normal traffic. The full classifier uses actual ground truth labels effectively provides us a soft performance upper bound. By taking a classification approach that leverages data label information, we can (1) study how this extra information boosts our approach if labels are available, and (2) obtain an upper bound on the online performance.

In Table 3, the results demonstrate that the iterative classifier, which directly optimizes the loss function (Equation 2), achieves better performance than N-over-D, which estimates conditional probabilities of the normal-day and mixture traffic distributions. The iterative classifier is also more robust in FP rates across various rejection ratios (0–1), demonstrated in Figure 4. However, its empirical runtime (Table 4) is roughly three to five times slower than that of N-over-D.

Table 4: Empirical runtimes per epoch ($\ell = 1$) for N-over-D and Iterative Classifier in seconds.

Dataset	N-over-D (s)	Iterative Classifier (s)
HULK	8	41
LOIC	3	8
CAIDA07	9	43

Moreover, the iterative classifier is sensitive to the estimated proportion of attack α in the loss function, whereas N-over-D is not. In our experiments, we use $\alpha = 0.6$ as a standard, but in other attacks where the attack proportion could be much higher, setting an appropriate α is crucial to the detection performance.

Reported evaluation metrics are Accuracy (ACC), False Positive Rate (FPR), Precision (Prec.), Recall (Rec.), and F1 Score (F1). The evaluation metrics are defined as:

$$\text{ACC} = (TP + TN) / (TP + TN + FP + FN)$$

$$\text{FPR} = FP / (FP + TN)$$

$$\text{Prec.} = TP / (TP + FP)$$

$$\text{Rec.} = TP / (TP + FN)$$

$$\text{F1} = 2[(\text{Prec.} \times \text{Rec.}) / (\text{Prec.} + \text{Rec.})]$$

where TP = True Positives, TN = True Negatives, FP = False Positives, FN = False Negatives. We run all the experiments on NVIDIA Tesla P100 GPUs with 12 GB memory. The models have been implemented in Python v3.7.5 using the Keras v2.2.4 library on top of the Tensorflow v1.14.0 machine learning framework.

7.3 Result Comparison

In addition, we compare our approaches, N-over-D and Iterative Classifier, with state-of-the-art DDoS detection methods on the *CICIDS2017* and *CAIDA07* datasets. We note that the other studied methods leverage on the labels of the data for their classification models. However, as mentioned in the problem formulation, labeled data is a commodity that is often unavailable during the time of an attack.

CICIDS2017. We compare our approach with two online and five offline methods. The two online approaches include a CNN classifier LUCID [5] that uses labeled data, and an unsupervised method E-KOAD [4]. SmartDetect [18] employs decision tree learning algorithms and adopts a different interpretation of the online

setting. SmartDetect uses a decision tree that is not updated during attack-day and thus not classified as online training under our definition. A network flow graph-based detector DeepGFL [30], and four deep-learning classifiers introduced by Cyber Security in IoT Networks [23] are offline methods. The four deep learning models are the MLP, 1D-CNN, LSTM, and 1D-CNN+LSTM models.

We note that the closest comparison, E-KOAD, presented the performance of the attack in the *CICIDS2017* dataset that they achieved the best results. However, their smallest reported interval, ℓ , is 2 mins, while we used 1-min intervals. Hence, we report the results for $\ell = 2$ min too. We present the evaluation results of all the models and contrast it with our proposed models (Table 5).

Table 5: Performance comparison with existing methods on the *CICIDS2017* dataset. The “Labels” column indicates if the method is trained with labeled data.

Method	FPR	Prec.	Rec.	F1	Labels	Online
N-over-D	0.001	0.999	0.887	0.940	No	$\ell = 2$
Iter. Classifier	0.034	0.852	0.810	0.831	No	$\ell = 2$
E-KOAD [4]	0.170	N/A	0.952	N/A	No	$\ell = 2$
LUCID [5]	0.005	0.993	0.999	0.996	Yes	$\ell = 1.7$
SmartDetect [18]	0.002	0.992	N/A	0.992	Yes	No
DeepGFL [30]	N/A	0.756	0.302	0.432	Yes	No
MLP [23]	N/A	0.884	0.862	0.873	Yes	No
1D-CNN [23]	N/A	0.981	0.901	0.939	Yes	No
LSTM [23]	N/A	0.984	0.898	0.895	Yes	No
1D-CNN+LSTM [23]	N/A	0.974	0.991	0.982	Yes	No

The four learning models [23] are constructed by combining LSTM, CNN, and fully connected layers. Out of these architectures, designed for DDoS detection in Internet of Things (IoT) networks, the 1D-CNN+LSTM model performs the best. While it produces good classification results, the other models seem to suffer from low true-positive rates. As for DeepGFL [30], it uses a graph representation of low-order features to perform classification. The results exhibit its weakness in identifying true positives, leading to low recall and F1 scores. Even though LUCID [5], a lightweight deep learning DDoS detection system, presents superior classification results, the model applies a CNN to network traffic flows to detect attacks that heavily rely on the accuracy of each traffic flow’s labels to achieve such performance.

On the other hand, N-over-D and Iterative Classifier achieve one of the highest scores in every evaluation metric while retaining low false-positive rates. It demonstrates competitive performance without utilizing the label information of the data.

CAIDA07. The compared methods include a self-organizing map detection scheme SOM [1], a support vector machine classifier SVM [22], a hybrid model SVM-SOM [20], and an enhanced history-based IP filtering method eHIPF [21]. We present the evaluation results of all the models and contrast it with the N-over-D performance (Table 6).

Methods [1, 20, 22] that adopt the SOM and SVM models produced relatively lower evaluation results, ranging from around 4%–8% false-positive rate and 92%–98% true-positive rate. As these are basic machine learning architectures, the models might not capture the high complexity of the network traffic as well as more complex neural architectures. Even though eHIPF [21] yields outstanding results, it is a system that invokes an intricate process

Table 6: Performance comparison with existing methods on the CAIDA07 real-world dataset. The “Labels” column indicates if the method is trained with labeled data.

Method	FPR	TPR	Rec.	F1	Labels	Online
N-over-D	0.025	0.990	0.624	0.765	No	$\ell = 1$
Iter. Classifier	0.020	0.991	0.640	0.778	No	$\ell = 1$
SOM [1]	0.065	0.934	N/A	N/A	No	No
SVM [22]	0.076	0.923	N/A	N/A	Yes	No
SVM-SOM [20]	0.038	0.981	N/A	N/A	Yes	No
eHIPF [21]	0.006	0.992	N/A	N/A	Yes	No

that requires data labels. The performance of N-over-D, without the need for label information, proves to be very close to that of eHIPF.

7.4 Analysis

While we show in the above results (Table 3) that the N-over-D and Iterative Classifier are effective against a range of real-world and synthetic DDoS attacks, the contrast between traffic distributions during normal and DDoS periods for this effect has not been illustrated. We study the performance of various online parameters ℓ , analyze these distributions from an empirical perspective, and discuss the pros and cons of the proposed approaches.

First, we present the performance (Table 7) for the proposed online approaches based on various intervals ℓ . It clearly shows that the iterative classifier outperforms N-over-D. As interval ℓ increases from 1 to 10 min, it significantly improves the iterative classifier’s detection ability. There is a notable increase in performance from $\ell = 1$ to $\ell = 2$ min. However, increasing intervals ℓ did not affect N-over-D much, and the performance is not sensitive to this online-learning parameter.

Table 7: Online performance comparison for values of interval ℓ on the CICIDS2017 dataset—LOIC attack.

Model	ℓ (min)	ACC	FPR	Prec.	Rec.	F1
Iterative Classifier	1	0.8805	0.0329	0.8808	0.6494	0.7476
	2	0.9355	0.0340	0.8529	0.8104	0.8311
	5	0.9462	0.0165	0.9212	0.7932	0.8524
	10	0.9486	0.0220	0.9016	0.8275	0.8630
N-over-D	1	0.9330	0.0047	0.9959	0.8846	0.9370
	2	0.9362	0.0009	0.9992	0.8875	0.9400
	5	0.9371	0.0001	0.0999	0.8882	0.9408
	10	0.9338	0.0037	0.9967	0.8853	0.9377

Next, in Table 8, we analyze the effects of proportion α on detection performance. While N-over-D is slightly more responsive to changes in α than the iterative classifier, the results show that both models are not sensitive to varying estimates of α . As volumetric DDoS attacks characteristically make up more than half of the network traffic, sometimes even up to 80–90% of traffic, this large amount of attack traffic allows us to set a higher α comfortably and not worry about it adversely affecting performance.

Finally, in Figure 5, we examine the differences in the normalized sequence score values of LOIC attack traffic. There is a clear distinction between the distribution of the N predicted values and the

Table 8: Robustness against errors in the estimate of α on the LOIC attack (actual value of $\alpha = 0.6$).

Model	α	ACC	FPR	Prec.	Rec.	F1
Iterative Classifier	0.5	0.9181	0.0706	0.7503	0.8717	0.8064
	0.55	0.9366	0.0277	0.8742	0.7895	0.8297
	0.6	0.9355	0.0340	0.8529	0.8104	0.8311
	0.65	0.9434	0.0173	0.9167	0.7822	0.8441
	0.7	0.9244	0.0098	0.9418	0.6543	0.7721
N-over-D	0.5	0.9362	0.0009	0.9992	0.8875	0.9400
	0.55	0.9730	0.0171	0.9864	0.9653	0.9757
	0.6	0.9430	0.1096	0.9202	0.9840	0.9510
	0.65	0.8996	0.2158	0.8548	0.9893	0.9172
	0.7	0.8501	0.3288	0.7944	0.9893	0.8812

values expected by D of the same traffic. It can be seen that N assigns lower values for the attack traffic as it predicts the likelihood of the traffic as normal. On the other hand, both online and offline D assigns higher scores to the same attack traffic as it predicts a high likelihood for these sequences to be attacks. Hence, by taking the $\frac{N}{D}$ scores, we obtain a more accurate measure of the probability that a sequence of traffic is normal, with a higher score indicating legitimacy.

8 RELATED WORK

In this section, we examine some of the most relevant works that take statistical and machine learning approaches in DDoS defenses.

8.1 Statistical Approaches to DDoS Defense

There is a large body of work on statistical approaches to DDoS mitigation since the early 2000s [37]. Common methods generally involve some statistical measure of network traffic properties, such as entropy scoring of network packets [6, 11, 26, 33], IP filtering [19, 25], and IP source traceback [13, 34]. These statistical and packet-based defense mechanisms are widely popular in defending against DDoS attacks. Statistical methods leverage variations of distributions in traffic attributes to discriminate between the distribution of traffic during normal usage and attack periods, and identify DDoS attacks.

Entropy-based detection approaches generally score packets based on some statistical metrics. In one of such works [33], it proposes a semi-Markov model for normal network browsing behavior. Through simulating the browsing dynamics of legitimate web browsers, the authors observed that a critical condition for successful attacks: the number of active bots of the botnet must not be lower than the number of active legitimate users. Based on the findings, they defined a new fine correntropy metric to detect DDoS attacks. Another proposed method [6] combines an entropy-based and a packet-score-based method for first characterizing incoming packets before filtering the malicious packets. Results show that the combined scoring method can differentiate between DDoS attacks and normal traffic. Similarly, a statistical solution [11] proposes a joint entropy-based security scheme against DDoS attacks. It generates baseline profiles, determines parameters when attacks are detected, and mitigates the DDoS attacks.

Figure 5: The distinct distributions of the LOIC attack traffic determined by models N and D. This divergence enable N-over-D to identify the DDoS attacks. The x-axis is the normalized score values (0–1), and the y-axis is the number of traffic sequences.

IP filtering, based on IP source address filtering, is another practical way to defend against DDoS attacks. Unlike methods that are based on monitoring the traffic volume, the IP filtering scheme [25] uses a history of legitimate IP addresses to decide if incoming IP packets are anomalous. The authors present several heuristic methods to make the IP address database accurate and robust to increase the effectiveness of the scheme. Peng et al. [19] uses a sequential nonparametric change point detection method to monitor the increase of new IP addresses, and demonstrate its effectiveness for DDoS attacks. IP source traceback methods tackle the problem of identifying the DDoS attack sources. As monitoring the traffic attributes and IP filtering could provide limited information to prevent DDoS attacks, traceback methods [13, 34] are another approach for effective defense against these attacks.

A common obstacle to these statistical techniques is the selection of an appropriate model parameters. Given that each DDoS attack in different networks varies from one to the next, it is especially challenging to identify the most suitable parameters that minimize false detection rates. Himura et al. [9] presented an automated parameter tuning of a statistical network traffic anomaly detection method, which determines a learning period for setting a parameter of the method.

However, many statistical solutions are not appropriate in an online setting, as some traffic statistics cannot be computed during data collection periods. To overcome this limitation, a DDoS defense architecture [32] was proposed to support distributed detection and automated online attack characterization, which assigns a score to each packet that estimates a packet's legitimacy given the attributes.

8.2 Machine Learning for DDoS Defense

Another popular approach to DDoS defense is based on Machine Learning (ML) anomaly detection techniques. Learning algorithms such as the Support Vector Machine (SVM), Self-Organizing Map (SOM), Random Forest, K-Nearest Neighbor (KNN), Decision Tree, and Neural Network have achieved various degrees of success in network intrusion detection [7, 16]. Fundamental methods SOM [1] and SVM [20–22] were first proposed to provide artificial intelligence in DDoS mitigation. Braga et al. [1] presented a SOM detection

scheme with 4 and 6 tuples of attributes as a lightweight detection method based on traffic flow features. Another work [22] applied the SVM classifier for detecting attacks. Some others [20, 21] took a hybrid approach that combined the SVM and SOM to enhance the accuracy in differentiating normal flows from abnormal flows.

A proposed method [16] employs a one hidden layer feedforward neural network to detect abnormal traffic, and it achieves high accuracy and low false detection. In another work [7], the authors evaluate nine ML methods on the source side data in the cloud for DDoS attack detection. The features (e.g., DNS packages ratio, Dillie-Hellman key exchange packages, ICMP package rate) are first extracted from the traffic data. The extracted data is then fed into the ML algorithms to distinguish between benign and malicious network traffic. Even though these ML methods are promising, the reported results depend heavily on the selection of features and the evaluated datasets [10]. Hence, to improve the generalization of traffic data and reduce the significance of feature engineering, deep learning models have been increasingly used in recent years.

The recent advancement of machine learning has given rise to increasing exploration of deep learning (DL) models for DDoS detection. We can broadly categorize these DL models into two groups—those that adopt an offline learning approach [3, 14, 15, 23, 31, 36], where all attack data is observed and available for training the models, and online learning models [4, 5, 7, 18] that only receive data, as more attack traffic is observed with time.

One offline method [14], CNN models with different model depths were used to investigate the relationship between performance and the number of CNN layers. They observe that deeper structures do not improve performance, and the LSTM model, a variant of RNN, performs better. The modeling of network traffic through time in a sequential manner, which retains information of historical patterns, instead of treating the inputs as independent packets. RNN-based intrusion detection systems [3, 15, 31] demonstrate that the sequential learning models have superior modeling capabilities, and they can detect some sophisticated attacks for intrusion detection tasks. RNN-based methods have demonstrated superior detection capabilities on some sophisticated attacks for intrusion detection tasks [3, 15, 31]. In addition, a few deep learning

methods [23, 36] that experiment with RNNs and CNNs hybrid models have slightly increased DDoS detection performance. However, all these methods are not designed for online systems, since the entire training data must be preprocessed before the models start training. They are not suitable for incremental updates when new data is observed.

While there are many offline learning-based DDoS detection methods, few online approaches have been proposed. One online method [5] that uses ground-truth labels for training, extracts traffic flow attributes (e.g., Packet Length, TCP Length, Window Size) as input for a deep CNN model.

Only a few methods do not use data labels for training. He et al. [7] evaluated nine basic ML methods with an added online learning module on the source side data in the cloud for DDoS attack detection. In another proposed online approach [18], basic learning algorithms and simple decision tree methods were employed. Last but not least, Çakmakçı et al. [4] employs a kernel-based learning algorithm in addition to the Mahalanobis distance and a chi-square test. It was shown to be competitive with offline DDoS classification algorithms. However, these existing online approaches are defined by their limited modeling capabilities, in which deep learning methods are capable of overcoming the limitations.

9 CONCLUSION

Adopting deep learning to filter out attack requests during DDoS attacks is not straightforward due to the lack of labeled data for training and the constraints imposed by requirements in the online setting. While there is a lack of labeled data, we observed that two types of datasets are available: a normal-day $\mathcal{N}$ dataset that consists of benign requests and the attack-day $\mathcal{M}$ data that consists of a mixture. By exploiting the availability of these two different datasets, we proposed two approaches guided by some learning objectives. In the N-over-D approach, by assuming that an unsupervised learning model can accurately predict the probability P_n and P_d from $\mathcal{N}$ and $\mathcal{M}$ respectively, we derive the conditional probability (that a request is an attack given that it is in the mixture) using its ratio. However, from a machine learning perspective, there is no joint-learning between the two training processes on $\mathcal{N}$ and $\mathcal{M}$, and thus potentially, there is room for improvement. In the iterative approach, we formulate a loss function more suitable for deep-learning. This formulation is guided by the intuition that the trained two-class classifier should classify requests in $\mathcal{N}$ as normal and α proportion of requests in $\mathcal{M}$ as attacks, where α is a given estimated value. We also proposed an iterative method that incorporates joint-training with the proposed loss function. Empirical studies showed that, even without labels, our proposed approaches attain high accuracy in the online setting.

While our work demonstrates that the attack-day dataset can be utilized to improve accuracy, this naturally leads to the next question of whether attackers can adapt the traffic to fool our online learning approaches into accepting malicious traffic as normal. The effectiveness of such DDoS attacks has been studied earlier on statistical-based filtering [17]. It would be interesting to further investigate the effects of adaptive DDoS attacks in the deep learning paradigm.

ACKNOWLEDGMENTS

This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-100E-2019-0033).

REFERENCES

- [1] R. Braga, E. Mota, and A. Passito. 2010. Lightweight DDoS flooding attack detection using NOX/OpenFlow. In *IEEE Local Computer Network Conference*.
- [2] Paul J Criscuolo. 2000. *Distributed denial of service: Trin00, tribe flood network, tribe flood network 2000, and stacheldraht ciac-2319*. Technical Report. California Univ Livermore Radiation Lab.
- [3] Jianjing Cui, Jun Long, Erxue Min, Qiang Liu, and Qian Li. 2018. Comparative Study of CNN and RNN for Deep Learning Based Intrusion Detection System. In *Cloud Computing and Security*.
- [4] Salva Daneshgadah Çakmakçı, Thomas Kemmerich, Tarek Ahmed, and Nazife Baykal. 2020. Online DDoS attack detection using Mahalanobis distance and Kernel-based learning algorithm. *Journal of Network and Computer Applications* (2020).
- [5] R. Doriguzzi-Corin, S. Millar, S. Scott-Hayward, J. Martinez-del Rincon, and D. Siracusa. 2020. LUCID: A Practical, Lightweight Deep Learning Solution for DDoS Attack Detection. *IEEE Transactions on Network and Service Management* (2020).
- [6] A. Gaurav and A. K. Singh. 2017. Entropy-score: A method to detect DDoS attack and flash crowd. In *2017 2nd IEEE International Conference on Recent Trends in Electronics, Information Communication Technology (RTEICT)*.
- [7] Z. He, T. Zhang, and R. B. Lee. 2017. Machine Learning Based DDoS Attack Detection from Source Side in Cloud. In *2017 IEEE 4th International Conference on Cyber Security and Cloud Computing (CSCloud)*.
- [8] P. Hick, E. Aben, K. Claffy, and J. Polterock. 2007. The CAIDA DDoS attack 2007 dataset.
- [9] Y. Himura, K. Fukuda, K. Cho, and H. Esaki. 2009. An Automatic and Dynamic Parameter Tuning of a Statistics-Based Anomaly Detection Algorithm. In *IEEE International Conference on Communications*.
- [10] K. S. Hoon, K. C. Yeo, S. Azam, B. Shumugam, and F. De Boer. 2018. Critical review of machine learning approaches to apply big data analytics in DDoS forensics. In *2018 International Conference on Computer Communication and Informatics (ICCCI)*.
- [11] K. Kalkan, L. Altay, G. Gür, and F. Alagöz. 2018. JESS: Joint Entropy-Based DDoS Defense Scheme in SDN. *IEEE Journal on Selected Areas in Communications* (2018).
- [12] Oleg Kuprev, Ekaterina Badovskaya, and Alexander Gutnikov. 2020. *DDoS attacks in Q1 2020*. <https://securelist.com/ddos-attacks-in-q1-2020/96837/>
- [13] Vadim Kuznetsov, Helena Sandström, and Andrei Simkin. 2002. An Evaluation of Different IP Traceback Approaches. In *Information and Communications Security*.
- [14] D. Kwon, K. Natarajan, S. C. Suh, H. Kim, and J. Kim. 2018. An Empirical Study on Network Anomaly Detection Using Convolutional Neural Networks. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*.
- [15] Chuanhuang Li, Yan Wu, Xiaoyong Yuan, Zhengjun Sun, Weiming Wang, Xiaolin Li, and Liang Gong. 2018. Detection and defense of DDoS attack-based on deep learning in OpenFlow-based SDN. *International Journal of Communication Systems* (2018).
- [16] J. Li, Y. Liu, and L. Gu. 2010. DDoS attack detection based on neural network. In *2010 2nd International Symposium on Aware Computing*.
- [17] Q. Li, E. Chang, and M. C. Chan. 2005. On the effectiveness of DDoS attacks on statistical filtering. In *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies*.
- [18] Francisco Sales de Lima Filho, Frederico AF Silveira, Agostinho de Medeiros Brito Junior, Genoveva Vargas-Solar, and Luiz F Silveira. 2019. Smart Detection: An Online Approach for DoS/DDoS Attack Detection Using Machine Learning. *Security and Communication Networks* (2019).
- [19] Tao Peng, Christopher Leckie, and Kotagiri Ramamohanarao. 2002. Detecting Distributed Denial of Service Attacks Using Source IP Address Monitoring. In *Proceedings of the Third International IFIP-TC6 Networking Conference*.
- [20] T. V. Phan, N. K. Bao, and M. Park. 2016. A Novel Hybrid Flow-Based Handler with DDoS Attacks in Software-Defined Networking. In *2016 Intl IEEE Conferences on Ubiquitous Intelligence Computing, Advanced and Trusted Computing, Scalable Computing and Communications, Cloud and Big Data Computing, Internet of People, and Smart World Congress (UIC/ATC/ScalCom/CBDCCom/IoP/SmartWorld)*.
- [21] T. V. Phan and M. Park. 2019. Efficient Distributed Denial-of-Service Attack Defense in SDN-Based Cloud. *IEEE Access* (2019).
- [22] T. V. Phan, T. Van Toan, D. Van Tuyen, T. T. Huong, and N. H. Thanh. 2016. OpenFlowSIA: An optimized protection scheme for software-defined networks from flooding attacks. In *2016 IEEE Sixth International Conference on Communications and Electronics (ICCE)*.

- [23] M. Roopak, G. Yun Tian, and J. Chambers. 2019. Deep Learning Models for Cyber Security in IoT Networks. In *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)*.
- [24] Iman Sharafaldin., Arash Habibi Lashkari., and Ali A. Ghorbani. 2018. Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy - Volume 1: ICISPP*.
- [25] Tao Peng, C. Leckie, and K. Ramamohanarao. 2003. Protection from distributed denial of service attacks using history-based IP filtering. In *IEEE International Conference on Communications. ICC*.
- [26] R. Wang, Z. Jia, and L. Ju. 2015. An Entropy-Based Distributed DDoS Detection Mechanism in Software-Defined Networking. In *2015 IEEE Trust-com/BigDataSE/ISPA*.
- [27] Kilian Weinberger, Anirban Dasgupta, John Langford, Alex Smola, and Josh Attenberg. 2009. Feature Hashing for Large Scale Multitask Learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*.
- [28] Paul Wood, Ben Nahorney, Kavitha Chandrasekar, Scott Wallace, and Kevin Haley. 2016. *Internet Security Threat Report*. Technical Report. Symantec Corporation.
- [29] Paul Wood, Ben Nahorney, Kavitha Chandrasekar, Scott Wallace, and Kevin Haley. 2019. *Internet Security Threat Report*. Technical Report. Symantec Corporation.
- [30] Yepeng Yao, Liya Su, and Zhigang Lu. 2018. DeepGFL: Deep feature learning via graph for attack detection on flow-based network traffic. In *MILCOM 2018-2018 IEEE Military Communications Conference (MILCOM)*.
- [31] C. Yin, Y. Zhu, J. Fei, and X. He. 2017. A Deep Learning Approach for Intrusion Detection Using Recurrent Neural Networks. *IEEE Access* (2017).
- [32] Yoohwan Kim, Wing Cheong Lau, Mooi Choo Chuah, and H. J. Chao. 2004. Packetscore: statistics-based overload control against distributed denial-of-service attacks. In *IEEE INFOCOM*.
- [33] S. Yu, S. Guo, and I. Stojmenovic. 2015. Fool Me If You Can: Mimicking Attacks and Anti-Attacks in Cyberspace. *IEEE Trans. Comput.* (2015).
- [34] S. Yu, W. Zhou, R. Doss, and W. Jia. 2011. Traceback of DDoS Attacks Using Entropy Variations. *IEEE Transactions on Parallel and Distributed Systems* (2011).
- [35] S. Yu, W. Zhou, W. Jia, S. Guo, Y. Xiang, and F. Tang. 2012. Discriminating DDoS Attacks from Flash Crowds Using Flow Correlation Coefficient. *IEEE Transactions on Parallel and Distributed Systems* (2012).
- [36] Xiaoyong Yuan, Chuanhuang Li, and Xiaolin Li. 2017. DeepDefense: Identifying DDoS Attack via Deep Learning. *2017 IEEE International Conference on Smart Computing (SMARTCOMP)* (2017).
- [37] S. T. Zargar, J. Joshi, and D. Tipper. 2013. A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks. *IEEE Communications Surveys Tutorials* (2013).

A IMPLEMENTATION DETAILS

We provide the LSTM architecture operations, describe the network traffic data preprocessing procedure, and report the default parameter settings for the experiments implemented in this paper.

A.1 LSTM Architecture Computation

We use the LSTM architecture without peep-hole connections. The operations are computed as follows:

$$\begin{aligned}
 i_t &= \sigma(W_{xi}x_t + W_{hi}h_{t-1} + b_i) \\
 f_t &= \sigma(W_{xf}x_t + W_{hf}h_{t-1} + b_f) \\
 o_t &= \sigma(W_{xo}x_t + W_{ho}h_{t-1} + b_o) \\
 g_t &= \tanh(W_{xg}x_t + W_{hg}h_{t-1} + b_g) \\
 c_t &= f_t \odot c_{t-1} + i_t \odot g_t \\
 h_t &= o_t \odot \tanh(c_t)
 \end{aligned}$$

where x_t is the input at time t , $\sigma(\cdot)$ is the sigmoid function, $\tanh(\cdot)$ is the hyperbolic tangent function, and $\odot$ denotes element-wise product. Each gate function has its own weight matrix and a bias vector. We denote the parameters with subscripts f for the forget gate function, i for the input gate function, and o for the output gate function respectively (e.g., W_{xf} , W_{hf} , and b_f are parameters of the forget gate function).

A.2 Data Preprocessing Algorithm

Using the data processing procedure below, we prepare network traffic data into a streaming form, grouped by source and destination IP addresses, for online learning.

Algorithm 1 Network data preprocessing algorithm.

```

1: for each interval  $\tau_i$  do
2:   Sort packets  $\mathbf{x} = \{x_1, \dots, x_n\}$  chronologically by absolute
   time
3:   for all  $x \in \mathbf{x}$  do
4:     if number of packets  $x^{(ip)}$  from an IP  $< \epsilon$  then
5:        $\mathbf{x} \setminus x^{(ip)}$ 
6:     end if
7:     Extract features  $f$ 
8:   end for
9:   Form sequences  $seq^{(ip)} = (f_1^d, f_2^d, \dots, f_T^d)$  with dynamic
   features of length  $seq\_len = T$ 
10:  if  $seq\_len < T$  then
11:     $zero\_pad(seq)$ 
12:  end if
13:   $seq^{(ip)} \cup f^{s(ip)}$   $\triangleright$  Join static features of IP to sequences
14: end for

```

A.3 Parameter Settings

For the numerator model N , we split the training data into an 80/20 training and validation ratio. As for the online model D , every batch of data in a time interval is divided into a 90/20 training and validation ratio to maximize the training data. Appendix A.3 lists the parameter settings of numerator N .

Model Parameters.

Dataset	Parameter	N	D
CAIDA07	epochs	30	10 (1 per interval)
	learning rate	0.005	0.003
	LSTM dimension	300	300
	embedding dimension	512	512
	batch size	512	128
	sequence length	200	200
LOIC	epochs	30	10 (5 per interval)
	learning rate	0.005	0.003
	LSTM dimension	300	300
	embedding dimension	512	512
	batch size	512	128
	sequence length	200	200
HULK	epochs	30	10 (5 per interval)
	learning rate	0.005	0.003
	LSTM dimension	300	300
	embedding dimension	512	512
	batch size	512	128
	sequence length	90	90

B PREDICTION ANALYSIS

We include more experimental analysis below that illustrate the different distributions learned by models N and both the offline and online D models. Appendix B show the divergent prediction distributions for the HULK attack. For the distributions of the predictions for the CAIDA07 attack, refer to Appendix B.

HULK attack. Prediction by N and offline D . **HULK attack. Prediction by N and online D .** **CAIDA07 attack. Prediction by N and offline D .** **CAIDA07 attack. Prediction by N and online D .**

C DATA SAMPLES

We also show some network traffic data samples to note the characteristic differences between normal and attack traffic in each dataset. Attacks include weird GET requests, such as the repeated appearance of 'GET/HTTP/1.0 Continuation ..' requests. These interactions happen almost exclusively between attackers and victims during attack times.

Select features of normal traffic samples (*HULK*).

Req #	Source IP	Dest. IP	TCP Len	Extra Info
4121780	192.168.10.25	204.154.111.110	1176	GET /visit.js?ctx=1828362..
4121781	13.107.4.50	192.168.10.15	2920	[TCP Spurious Retransmission] 80
4121782	13.107.4.50	192.168.10.15	1460	80 → 49672 [ACK] Seq=219..
4121783	192.168.10.15	13.107.4.50	0	49672 → 80 [ACK] Seq=157..
4121784	192.168.10.25	204.154.111.110	1186	GET /visit.js?ctx=182836..

Select features of attack traffic samples (*HULK*).

Req #	Source IP	Dest. IP	TCP Len	Extra Info
6845625	172.16.0.1	192.168.10.50	335	GET /?KRNCJUNAS=OZB HTTP/1.1
6845626	172.16.0.1	192.168.10.50	349	GET /?SNDUV=WKGEQ HTTP/1.1
6845627	192.168.10.50	172.16.0.1	1448	HTTP/1.1 200 OK [TCP segment of ..]
6845628	192.168.10.50	172.16.0.1	5792	80 → 50782 [ACK] Seq=1449 Ack=327 ..
6845629	192.168.10.50	172.16.0.1	0	80 → 50826 [SYN, ACK] Seq=0 Ack=1 ..

Select features of normal traffic samples (*LOIC*).

Req #	Source IP	Dest. IP	TCP Len	Extra Info
7394913	192.168.10.12	72.246.40.98	296	GET /success.txt HTTP/1.1
7394914	72.246.40.98	192.168.10.12	0	80 → 56910 [ACK] Seq=1 ..
7394915	72.246.40.98	192.168.10.12	384	HTTP/1.1 200 OK (text/plain)
7394916	192.168.10.12	72.246.40.98	0	56910 → 80 [ACK] Seq=297 ..
7394917	74.125.22.109	192.168.10.5	85	Application Data

Select features of attack traffic samples (*LOIC*).

Req #	Source IP	Dest. IP	TCP Len	Extra Info
8140894	172.16.0.1	192.168.10.50	20	GET / HTTP/1.0 Continuation ..
8140895	172.16.0.1	192.168.10.50	20	GET / HTTP/1.0 Continuation ..
8140896	192.168.10.50	172.16.0.1	0	80 → 49786 [ACK] Seq=1 ..
8140897	192.168.10.50	172.16.0.1	0	80 → 49787 [ACK] Seq=1 ..
8140898	192.168.10.50	172.16.0.1	5840	HTTP/1.1 200 OK [TCP segment ..]

Select features of normal traffic samples (*CAIDA07*).

Req #	Source IP	Dest. IP	TCP Len	Extra Info
6513	192.95.27.190	71.126.222.64	48	50723 → 35451 [SYN] Seq=0 Win=64512 Len=0 ..
6514	192.120.148.227	71.126.222.64	60	Echo (ping) request id=0xc0200, seq=1280/5 ..
6515	202.1.175.252	71.126.222.64	60	Echo (ping) request id=0xcce1d, seq=1280/5 ..
6516	192.95.27.190	71.126.222.64	60	Echo (ping) request id=0xc495, seq=1280/5 ..
6517	192.95.27.190	71.126.222.64	48	50725 → 35451 [SYN] Seq=0 Win=64512 ..

Select features of attack traffic samples (*CAIDA07*).

Req #	Source IP	Dest. IP	TCP Len	Extra Info
407343	198.8.152.145	71.126.222.64	60	Echo (ping) request id=0xc8001, seq=1280/5 ..
407344	194.176.2.50	71.126.222.64	60	Echo (ping) request id=0xc0200, seq=1280/5 ..
407345	195.110.239.66	71.126.222.64	60	Echo (ping) request id=0xc0200, seq=1280/5 ..
407346	195.110.239.66	71.126.222.64	48	9536 → 44279 [SYN] Seq=0 Win=64512 Len=0 ..
407347	58.201.177.243	71.126.222.64	48	28559 → 53181 [SYN] Seq=0 Win=64512 Len=0 ..