



**QUEEN'S
UNIVERSITY
BELFAST**

GaDoT: GAN-based Adversarial Training for Robust DDoS Attack Detection

Abdelaty, M., Scott-Hayward, S., Doriguzzi-Corin, R., & Siracusa, D. (2022). GaDoT: GAN-based Adversarial Training for Robust DDoS Attack Detection. In *Ninth IEEE Conference on Communications and Network Security (IEEE CNS 2021): Proceedings* Institute of Electrical and Electronics Engineers Inc..
<https://doi.org/10.1109/CNS53000.2021.9705040>

Published in:

Ninth IEEE Conference on Communications and Network Security (IEEE CNS 2021): Proceedings

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

© 2021 IEEE.

This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

GADoT: GAN-based Adversarial Training for Robust DDoS Attack Detection

Maged Abdelaty^{1,3}, Sandra Scott-Hayward², Roberto Doriguzzi-Corin¹, and Domenico Siracusa¹

¹ICT, Fondazione Bruno Kessler - Italy ²CSIT, Queen's University Belfast - Northern Ireland ³University of Trento - Italy
Email: mabdelaty@fbk.eu, s.scott-hayward@qub.ac.uk, rdoriguzzi@fbk.eu, dsiracusa@fbk.eu

Abstract—Machine Learning (ML) has proven to be effective in many application domains. However, ML methods can be vulnerable to adversarial attacks, in which an attacker tries to fool the classification/prediction mechanism by crafting the input data. In the case of ML-based Network Intrusion Detection Systems (NIDSs), the attacker might use their knowledge of the intrusion detection logic to generate malicious traffic that remains undetected. One way to solve this issue is to adopt adversarial training, in which the training set is augmented with adversarial traffic samples. This paper presents an adversarial training approach called GADoT, which leverages a Generative Adversarial Network (GAN) to generate adversarial DDoS samples for training. We show that a state-of-the-art NIDS with high accuracy on popular datasets can experience more than 60% undetected malicious flows under adversarial attacks. We then demonstrate how this score drops to 1.8% or less after adversarial training using GADoT.

Index Terms—Adversarial attacks, Adversarial machine learning, Adversarial training, Distributed Denial of Service, Intrusion detection systems

I. INTRODUCTION

The complexity and volume of Distributed Denial of Service (DDoS) attacks continues to increase, causing significant disruption to critical online services and leading to economic losses. For instance, the number of DDoS attacks is expected to reach 15.4 million globally by 2023 [1]. Meanwhile, a new record of 2.3 Tbps was reached in an attack targeting Amazon Web Services (AWS) in Q1 2020 [2]. To achieve this scale and damage, attackers develop more sophisticated attacks combining multiple vectors (i.e. protocols) and adapting their attacks to bypass NIDSs.

Over recent years, NIDSs have evolved to leverage ML models to cope with the detection of new attacks, which cannot be identified by signature-based detection methods. An ML-based NIDS essentially consists of a feature extractor and a ML-based model. The feature extractor processes the raw network traffic to derive features that are arranged into data samples suitable to input to the model. The model is trained using malicious and benign samples and then deployed to classify new samples extracted from live traffic. Several ML-based NIDSs have been proposed demonstrating detection of DDoS attacks with high accuracy [3]–[5].

However, in the same way that an attacker might adapt their attack to remain just below the threshold of a threshold-based detection system or modify their attack to avoid signature-based detection, ML-based NIDSs are vulnerable to adversarial attack. The adversarial attack aims to evade the classification model

and thus reach the target, achieving the planned DDoS. In a DDoS scenario, an attacker generates perturbed DDoS flows such that the NIDS extracts samples incorporating features slightly different from those of unperturbed flows. The ability to lead the ML model to misclassify adversarial samples and hence reduce the NIDS accuracy has been demonstrated in previous works [6], [7].

Generative Adversarial Networks (GANs) are deep generative models with the capability of learning the distribution of the input data, then producing new similar examples. A GAN consists of two models: a discriminator and a generator that are trained to compete with each other. The generator model captures the training data distribution to produce fake samples resembling those from the training data. Our intuition here is that GAN-based sample generators can be exploited to enrich the training set with adversarial examples similar to those that an attacker would use to fool the NIDS, that is, DDoS examples with the feature distribution of the benign samples.

In this paper, we present GAN-based Adversarial training for robust DDoS attack deTectioN (GADoT), an adversarial training framework designed to produce strong adversarial examples, hence enabling an effective adversarial training of ML-based DDoS classifiers. Our framework employs the generator model of a GAN trained on benign samples to produce fake-benign samples. Using these fake-benign samples, we perturb the DDoS samples by replacing their features with values from those fake samples. The perturbed samples are added to the training dataset, which already contains benign samples and unperturbed DDoS samples, creating an augmented dataset for adversarial training. The trained model is expected to classify DDoS samples correctly, whether they are perturbed or not. To the best of our knowledge, GADoT is the first solution to use GAN-based perturbations that decouples the trained model from the generation process of adversarial samples to prevent the model from influencing the strength of adversarial samples [8].

There are two approaches to implementing adversarial attacks against NIDS; perturbing data samples and perturbing network traffic¹. In the first approach, the adversary introduces perturbations directly on the preprocessed samples before feeding them to the ML model. In the second approach, the adversary generates adversarially perturbed DDoS flows, which

¹Formally, this equates to the feature space and the problem space, respectively, as studied in [9].

are processed by the feature extractor to create data samples to be classified by the ML model. This approach can be more realistic than perturbing extracted samples that require the adversary to bypass the feature extractor to feed such samples to the model.

For a practical evaluation, we adopt the approach of perturbing the network traffic to generate two datasets in which the semantics of DDoS attacks are fully preserved. The first dataset contains six traces with adversarial SYN flood attacks [10], while the other is based on the CICIDS2017 dataset [11] and contains traces with adversarial HTTP GET flood attacks [12]. The evaluation results on both datasets are the basis of our conclusions on the performance of GADoT. To the best of our knowledge, GADoT is the first solution to be evaluated against representative real-world network attack scenarios.

The contributions of this paper are the following:

- We propose GADoT, which is an adversarial training approach to increase the robustness of ML models used for DDoS attack detection. The approach exploits a GAN to generate fake-benign samples in order to perturb DDoS samples. The training dataset is augmented with these perturbed samples in order to train adversarially robust models.
- We provide an analysis of single and multiple (combined) feature perturbations to offer insight into the model's vulnerability before and after adversarial training.
- We offer a practical evaluation and analysis of the robustness of the ML model trained with GADoT by using network traffic traces that capture adversarially perturbed SYN and HTTP DDoS flood attacks.

The rest of the paper is structured as follows: Section II reviews the related work on adversarial attacks against NIDSs and the proposed solutions to increase the robustness of ML-based NIDSs. Section III introduces the threat model. The proposed GADoT approach is explained in Section IV. Section V presents the datasets, details of the specific feature perturbation and experiment design. The results are discussed in Section VI. Finally, Section VII concludes the paper.

II. RELATED WORK

This section reviews the recent advances in the evaluation of NIDSs in an adversarial setting and the approaches for enhancing the robustness of these systems.

A. Adversarial attacks against NIDSs

As highlighted in Section I, the vulnerability of learning-based NIDSs to adversarial perturbation has been explored and implemented by means of two approaches: perturbing data samples and perturbing network traffic.

1) *Perturbing extracted samples*: The key to the success of an adversarial evasion attack is to maintain the characteristic of the attack. For example, for a SYN Flood attack that sends a volume of SYN packets (i.e. SYN flag set), a feature based on the *flags* of the TCP header should not be manipulated. Yan et al. [13] split features in the KDDCup99 dataset into *unchangeable* that

are essential to maintain the malicious function and *changeable* than can be manipulated without affecting this function. Then, the authors employed a Wasserstein GAN (WGAN)-based neural network to manipulate the changeable features of the DDoS attack. The reported results show a 50% drop in the detection accuracy of the NIDS on the KDDCup99 dataset. However, this approach lacks the discussion of how these perturbed samples can be reproduced in live network traffic. This live traffic should be valid according to the different network protocols in order to pass through the network and carry out the intended attack function on the victim side.

2) *Perturbing network traces*: Aiken et al. [7] investigate adversarial attacks targeting an anomaly based NIDS within Software-Defined Networks (SDN). The authors propose perturbing three features: payload size, packet rate and volume of bidirectional traffic, which can be implemented in live traffic without compromising the attack function. The attack is implemented in SYN flood traces that are classified by several ML models, including Support Vector Machine (SVM) and K-Nearest Neighbour (KNN). The results show that by combining perturbed features, a detection accuracy drop from 100% to 0% is observed across some of the classifiers and the KNN classifier is more robust than the others.

B. Defence mechanisms in NIDSs

In the following, we review some of the related work proposed to increase the robustness of NIDSs, which are based on denoising autoencoders [14], adversarial training [15], [16] and combining adversarial training with ensemble voting [17].

Hashemi et al. [14] propose a solution for increasing the robustness through denoising the samples to remove any adversarial perturbations. The authors evaluated their approach on the CICIDS2017 dataset and use the adversarial attacks presented in [18]. The results suggest that RePO is more robust than other NIDSs such as Kitsune [3].

Zhang et al. [17] propose a defence approach that combines three techniques: ensemble voting based on three different ML models, ensemble adversarial training [8], and query detection. A flow is classified as benign only if all the models classify it as benign, forcing the adversary to craft samples that can evade all models simultaneously. Moreover, they use ensemble adversarial training [8] to augment the training dataset with adversarial samples, improving the robustness of individual models. The defence approach is evaluated on the CICIDS2018 dataset [11] and was effectively able to reduce the success rate of the adversarial attacks.

Despite the good performance of these defence approaches, they are only evaluated on perturbations that are implemented on extracted samples, and there is no discussion of how to convert such perturbed samples into real network traffic. Furthermore, a common limitation of the adversarial training approaches [15], [16], [19] is the use of adversarial samples that are generated based on the target/victim models. In such approaches, the victim model learns to generate weak adversarial samples, instead of learning to defend against strong perturbations and hence remains vulnerable to adversarial attacks [8].

TABLE I
ADVERSARY MODEL

Threat Model Characteristic	Type	Adversary View
Adversary Goals	Compromising integrity (Evasion)	✓
	Compromising availability	✗
	Compromising confidentiality	✗
Adversary Knowledge	Feature set	✓
	Satisfy domain constraints	✓
Adversary Capabilities	Manipulate training data	✗
	Manipulate test data	✓
	Manipulate model	✗

In this work, we propose an adversarial training approach called GADoT in which the target model is decoupled from the generation of the adversarial samples used for training. We generate the adversarial samples by perturbing DDoS samples with values obtained from a GAN. The trained model thus learns to defend against strong perturbations and can be resistant to adversarial attacks. Furthermore, we evaluate the proposed approach using realistic network traces that are adversarially perturbed. These traces capture perturbations that are suitable for volumetric DDoS attacks.

III. THREAT MODEL

To characterise the adversary attacking our ML-based NIDS, we identify their goals, capabilities and knowledge, as per the threat modeling framework presented by Biggio et al. in [20]. Our adversary model is summarised in Table I.

A. Attacker’s goal

The adversary aims to evade the ML model through manipulating the DDoS traffic to be misclassified as benign. If successful, this operation would allow the attack traffic to remain undetected and to reach the victim machine or network. This maps to the online evasion attack as defined in the recently published MITRE adversarial threat matrix [21].

B. Attacker’s knowledge

The attackers know that the network traffic is monitored by an ML-based NIDS. However, they have no direct access to the NIDS itself or the underlying ML model. As a result, they cannot compromise the NIDS by bypassing some of its components (such as the feature extractor) in order to feed manipulated samples directly to the model. Thus, state-of-the-art attacks [15], [16], [22] based on techniques such as Fast Gradient Sign Method (FGSM), which manipulate the data samples, are not viable in such an attack scenario. However, since DDoS attacks are well researched in the literature, the attackers can exploit their field knowledge to know the basic traffic features employed for attack detection. Based on that, they manipulate these features to evade the classification model. At the same time, the manipulated traffic must satisfy network domain constraints such that the objective of the attack remains




Figure 1. Illustration of the GADoT approach for adversarial training

valid. For example, as previously noted, a SYN Flood must be comprised of SYN packets. Similarly, the adversary should ensure that a manipulated packet will not be discarded for reasons such as invalid header fields (e.g., the checksum). These assumptions present a realistic threat model for the NIDS problem space.

C. Attacker’s capability

The attackers can perturb the DDoS flows at test time in order to evade the classification model, allowing the DDoS attack to pass undetected. They can also build upon the published adversarial attacks, such as [7], to select the traffic parameters that can be tuned to craft an evasion attack.

IV. THE GADoT APPROACH

To produce a robust deep learning-based DDoS detector, we propose an adversarial training approach, GADoT, in which an ML model is trained with benign samples, DDoS samples, and adversarial DDoS samples. The goal is to increase the robustness against the perturbed DDoS samples, in anticipation of an attacker who can craft packets and perturb malicious flows to mimic those characteristics of benign flows. The approach employs a generator neural network that generates fake-benign samples used to create adversarial samples and a perturbation module that produces the augmented dataset T_{adv} used for adversarial training of the target model, as depicted in Figure 1.

In the following, we briefly introduce the target model used in this work, Lightweight, Usable CNN in DDoS Detection (LUCID)² [4], and how it prepares the data samples before classification, then we explain the modules employed by GADoT. Finally, we describe two baseline approaches against which GADoT is compared.

A. LUCID as a target model

LUCID is an efficient DDoS detection system that implements a pre-processing tool for the network traffic and a Convolutional Neural Network (CNN). The source code of both pre-processing tool and neural network are publicly available at [23]. The preprocessing tool, which we refer to as the feature extractor, converts the flows extracted from the network traffic into data samples compatible with the CNN. In Figure 2, we illustrate the feature extractor and the classification model. The

²We selected LUCID as our target model as it achieves the best classification accuracy of the deep learning-based DDoS detection solutions in the literature. Its preprocessing tool also enables us to adopt the practical approach of perturbing the network traffic.




Figure 2. Illustration of the traffic classification process

TABLE II
THE 11 PACKET HEADER ATTRIBUTES USED IN LUCID [4]

Feature	Description
Time	Relative time of each packet with respect to the first packet in the time window
Packet Len	Length of the entire IP packet
Highest Layer	The highest protocol in a packet, e.g. TCP or HTTP
IP Flags	Three bits taking one of three states: all zeros, do not fragment or more fragments
Protocols	Representing the protocols found in the packet
TCP Len	Size of the TCP segment (header + payload)
TCP Ack	Relative acknowledgment number
TCP Flags	Nine 1-bit flags of a TCP segment
TCP Win Size	The value of the field of the same name in a TCP segment
UDP Len	Size of the UDP segment (header + payload)
ICMP Type	The value of the type field of an ICMP message

figure also contains other details such as dataset names that we introduce later, in Section V. The CNN classifies the data samples as either benign or DDoS. Each data sample includes a set of features extracted from each packet of a bi-directional flow in a specific time window.

Using the LUCID feature extractor, we extract 11 features that are described in Table II and use ten seconds as the time window length. A network flow, or sample, is represented as an array of fixed size 10×11 , whose 10 lines are the packets that belong to the flow collected within the time window, while the 11 columns are the features extracted from each packet. The size of the array is fixed and decided at training time, as this is a requirement for the CNN. Flows with more than 10 packets in a time window are truncated, while shorter flows are zero-padded. The number of actual packets in a sample (non zero-padded rows) can be seen as a further feature that we call *Flow length* in the rest of the paper.

B. The generator neural network

The generator neural network is the generator model \mathcal{G} of a Wasserstein GAN with gradient penalty (WGAN-GP) [24] that achieves a state of the art performance in terms of training stability and diversity of the generated samples. In addition to the generator \mathcal{G} , the WGAN-GP includes a discriminator model \mathcal{D} that learns to distinguish between the real samples from the training data and the generator’s output. During the GAN training, \mathcal{G} and \mathcal{D} compete with each other. \mathcal{G} learns to generate better fake samples that fool \mathcal{D} , whereas \mathcal{D} learns to label those samples, minimising the success of \mathcal{G} . An optimal generator causes the discriminator to fail to distinguish between the real and generated samples.

Algorithm 1 Adversarial dataset generation algorithm

Input: T : Training dataset; Y : True labels of T ; F : List of features to perturb;

Output: T_{adv} : Adversarial dataset;

- 1: $T_{adv} = Copy(T)$; ▷ Initialise T_{adv}
- 2: **for** f **in** F **do**
- 3: $z = \mathcal{N}(0, 1)$; ▷ Source of noise
- 4: $T_g = \mathcal{G}(z)$; ▷ Generate fake-benign samples
- 5: $T_c = Copy(T)$;
- 6: **if** $f == flow_length$ **then**
- 7: Fill zero-padded rows in $T_c[Y == 1]$;
- 8: **else**
- 9: $T_c[Y == 1, f] = T_g[:, f]$;
- 10: **end if**
- 11: $T_{adv}.append(T_c)$;
- 12: **end for**
- 13: **return** T_{adv}

Given a training dataset of benign and DDoS samples, we train the generator and discriminator of the WGAN-GP with the benign samples. After training, the discriminator is removed and the trained generator \mathcal{G} of WGAN-GP becomes the generator of GADoT (Figure 1). This generator takes a source of noise z as an input and generates new samples $\hat{x} = \mathcal{G}(z)$ with the feature distribution of the training data. As WGAN-GP is trained with benign samples, the resulting $\mathcal{G}(z)$ produces what we call *fake benign samples*.

We chose to generate benign samples in order to perturb the features of the DDoS samples with values seen in benign traffic. This matches with the idea of DDoS flows disguised as benign through mimicking the characteristics of normal flows. For instance, [6], [7] show how the attacker can evade a wide range of classifiers after manipulating features such as the number of bytes and packets, inter-arrival time and packet rate in malicious network traffic. An advantage of using a GAN over the real benign samples for perturbation is the wide variety of unique values that a GAN produces. A well-trained GAN produces a unique output for each random input, providing a wide range of fake-benign samples.

The architecture of the generator neural network is based on the deep convolutional generator of WGAN-GP [24] with minimal changes to match the shape of LUCID samples. The architecture is modified to have output shape 10×11 , as per the shape of LUCID samples, and 20 as the dimension of the noise source z to keep it smaller than the dimensionality of the output [25]. To train this generator, we utilise the discriminator of the same WGAN-GP and modify its input shape to be 10×11 . The other hyperparameters of both generator and discriminator are unchanged with respect to the original architecture. Finally, the output layer in [24] uses a *tanh* activation function that produces samples with features between $[-1, 1]$. We rescale the values between $[0, 1]$ in accordance with the normalisation scheme adopted by LUCID, where all input features are normalised between 0 and 1 before classification.

C. The perturbation module

The perturbation module takes a dataset with benign and DDoS samples, and fake-benign samples produced by the

generator neural network and outputs the adversarial training dataset. The perturbed DDoS samples in the adversarial training dataset are crafted either by replacing features with values seen in benign traffic or by replacing the zero-padded rows with dummy packets. Both the dummy packets and the substitute values of features are obtained using the generator model.

Algorithm 1 shows the process of generating the adversarial training dataset T_{adv} . The inputs of the algorithm are a dataset T with both benign and unperturbed DDoS samples, the ground truth Y of the dataset and the list of features to perturb F , while the output is the adversarial training dataset T_{adv} . F is a subset of the 12 features of a sample, 11 packet header features plus *Flow length*, as described in Section IV-A. Note that each element in Y is a binary label $\in \{0, 1\}$, where 1 denotes a DDoS sample, and 0 is a benign sample.

In the algorithm, a copy of T initialises T_{adv} (line 1). To perturb a feature $f \in F$, we start with feeding a source of noise into the generator model $\mathcal{G}(z)$ to generate a set of fake-benign samples T_g ; then we make a copy of T named T_c (lines 3–5). For f the *Flow length* feature, the goal is to perturb the number of packets in each DDoS sample of T_c ; and thus the zero-padded rows in each sample are filled with packets from the fake samples of T_g (line 7). For the other features, the value of feature f in each DDoS sample of T_c is replaced by a value from the fake samples of T_g (line 9). Then T_c is appended to T_{adv} (line 11). In this way, T_{adv} is augmented with adversarial DDoS samples whose features can be observed in benign traffic.

The resulting T_{adv} is unbalanced, with more malicious than benign samples. To obtain a balanced dataset, we duplicate some of the benign samples until we reach an equal number of malicious and benign samples. Of course, the ground truth Y is updated with the labels of the newly added adversarial and benign samples. It is important to note that, in contrast to the related work [16], [19], GADoT decouples the trained model from the generation of the adversarial samples used for training. This is important to prevent the trained model from influencing the strength of the generated adversarial samples [8].

D. Adversarial training

We use an adversarial training dataset T_{adv} in order to train the target model to increase its robustness to adversarially perturbed DDoS attacks. As T_{adv} contains benign samples as well as perturbed and unperturbed DDoS samples, the goal of training is to minimise the error between the ground truth and the predictions through optimising the learnable parameters of the model.

Having presented the modules of GADoT, we revisit Figure 1 to summarise the end-to-end approach. The input for the perturbation module is a labelled training dataset T , which contains unperturbed DDoS samples and benign samples. Then, the generator neural network receives a number of noise samples from the Gaussian noise source $z = \mathcal{N}(0, 1)$ to generate an equal number of fake-benign samples T_g to perturb the DDoS samples, as explained in Algorithm 1. This process produces

TABLE III
OVERVIEW OF THE DATASETS

	Dataset	#Samples	#Benign	#DDoS
Training	Custom-SYN	52698	26349	26349
	UNB201X	265902	132746	133156
Evaluation	Scapy-SYN	18830	–	18830
	CICIDS2017	13193	–	13193

the adversarial training dataset T_{adv} , which is then used to train the model, LUCID. The outcome of the GADoT approach is a robust DDoS detection model capable of distinguishing between benign samples and DDoS samples even under adversarial attack.

E. Baseline approaches

GADoT is compared against two baseline approaches that can generate perturbed DDoS samples for adversarial training. In the first approach, which we refer to as Benign Feature Perturbation (BFP), the training DDoS samples are perturbed with features obtained from the benign samples in the training dataset. We use the same procedure explained in Algorithm 1 but with T_g from the training data to save the cost of the GAN. The second approach is based on the renowned FGSM [26] algorithm that uses the gradients of the neural network to create adversarial samples. We use FGSM with ℓ_∞ bound on the perturbation magnitude to create adversarial samples with minimal perturbations to fool the model. For both baselines, the augmented dataset T_{adv} contains benign samples, DDoS samples, and adversarial DDoS samples with perturbed features as well as perturbed padding rows.

V. DESIGN OF EXPERIMENTS

This section presents the datasets, the process of training different ML models, and the design of the experiments used to evaluate the performance of the GADoT approach in normal and adversarial settings.

A. Datasets

For our evaluation, we consider two types of DDoS attacks: TCP SYN flood [10] and HTTP GET flood [12]. Both aim at exhausting the victim’s resources by sending large amounts of requests to the victim server, making its services unavailable to the users. We use four datasets: Custom-SYN and Scapy-SYN contain TCP SYN flood attacks while the CICIDS2017 [11] and the UNB201X [4] contain HTTP GET flood attacks. An overview of these datasets is provided in Table III.

1) **Custom-SYN dataset:** A traffic trace recorded in our testbed. It combines legitimate traffic, e.g. web browsing and ssh, and SYN flood attack traffic generated by using the Hping network tool [27].

2) **UNB201X dataset:** This dataset was introduced in [4] to enable LUCID to better detect DDoS attacks irrespective of the training dataset. UNB201X is a balanced combination

TABLE IV
SUMMARY OF THE METHODOLOGY FOR TRAINING AND EVALUATION OF
THE LUCID MODELS

Model Name	Model Detection Goal	Training	Evaluation
Model-SYN	SYN flood	Custom-SYN training dataset	Traffic traces: Scapy-SYN Perturbed Features: <i>IP Flags</i> ; <i>TCP Len</i> ; <i>Flow length</i> Crafting tool: Scapy
Model-HTTP	HTTP GET flood	UNB201X training dataset	Traffic traces: CICIDS2017 Perturbed Features: <i>Time</i> <i>TCP Len</i> ; <i>Flow length</i> Crafting tool: Scapy

of preprocessed traces from ISCX2012 [28], CICIDS2017 and CSECIC2018 [11]³.

3) **Scapy-SYN dataset**: This dataset contains a SYN flood attack crafted by using the Scapy tool [29]. Scapy is an open source program to send, capture, dissect and forge network packets. It comprises seven different traffic traces: one with unperturbed SYN attack packets and six traces that contain SYN packets with different perturbed features:

- (i) *IP flags*: bit *do not fragment* randomly set to 0 or 1.
- (ii) *TCP length*: SYN packets with random payload content.
- (iii) *Padding replacement*: the SYN packet is followed by a random number of dummy packets with a random delay.
- (iv) *SYN Packet Replication*⁴: the SYN packet is repeated a random number of times with a variable delay.
- (v) *IP flags & TCP length & Padding replacement*: combination of these three perturbations.
- (vi) *IP flags & TCP length & SYN Packet Replication*: combination of these three perturbations.

With regard to *Padding replacement* and *SYN Packet Replication*, these are perturbations to the *Flow length* feature in which the attacker exploits the zero-padded rows in DDoS samples to fool the DDoS detection model. The attacker’s aim is to force the feature extractor to fill them with a random number of additional packets, making the DDoS samples less dissimilar to the benign samples (in terms of number of packets per flow). This can easily be achieved by injecting dummy packets into the network with the same IP address and TCP port of the SYN packet, or by replicating the same SYN packet multiple times within a time window.

4) **CICIDS2017 dataset**: The CICIDS2017 dataset contains a wide range of normal and malicious network activities, including DDoS. We use the timeslot 3.30PM-5.00PM of the CICIDS2017-Fri7PM trace that contains HTTP DDoS traffic generated by using LOIC [30]. We perturb the CICIDS2017 traffic trace with Scapy to introduce two perturbations: random delays between the DDoS packets and the fragmentation of those packets. The delay between packets is reflected in the packet arrival time, which is one of the features used by LUCID (see Table II). Similarly, packet fragmentation decreases the *TCP*

³The UNB201X dataset has kindly been made available by the authors of [4] to enable our evaluation.

⁴The nature of the TCP SYN flood attack with a single SYN packet constituting a flow makes this attack well suited for experimenting with perturbations on the *Flow length* feature.

Len of each packet and thus increases the number of packets in the attack flows. Both perturbations are particularly suitable for the HTTP DDoS traffic, as each attack flow comprises several packets.

B. Training and evaluating the models

We produce two separate models. A model called *Model-SYN* is trained to detect the SYN flood attacks, while a second model called *Model-HTTP* is trained to detect the HTTP GET flood attacks⁵. Both models are based on the LUCID code available at [23], hence implemented in Python v3.6.8 using the Keras API v2.2.4 on top of Tensorflow 1.13.1 [31] and the ADAM optimization algorithm. In the following, we present the details of training and evaluating each model, with and without GADoT. A summary is provided in Table IV.

1) **Model-SYN**: First, this model is trained with the Custom-SYN training dataset in order to detect the SYN flood attacks in normal settings, i.e. to classify unperturbed flood attacks. Using GADoT, the same model is adversarially trained with the dataset T_{adv} generated by inserting the Custom-SYN training set as input to Algorithm 1. As T_{adv} contains adversarial samples perturbed using the GAN neural network, we expect the model to be robust and correctly classify the SYN flood traffic in adversarial settings. This model is evaluated on the six perturbed traces of the Scapy-SYN dataset to assess the robustness after using GADoT.

2) **Model-HTTP**: To detect HTTP GET flood attacks, we train this model with the UNB201X training dataset. The trained model is used for attack detection in normal settings, i.e. to classify unperturbed flood attacks. Using GADoT, the model is retrained with the dataset T_{adv} generated by using the UNB201X training set as input to Algorithm 1. In both cases T_{adv} is generated based on F that includes the 11 features of LUCID samples as well as perturbed padding rows. We evaluate this model on the manipulated CICIDS2017 traffic traces. This provides an evaluation based on the practical adversary capability of perturbing network traffic. The results of these experiments are presented in sections VI-A and VI-B.

In our experiments, the benign samples in the Custom-SYN and UNB201X test datasets are fixed throughout the experiments, while the DDoS samples in both datasets are replaced by a similar number of perturbed DDoS samples for the relevant experiment. In this way, variations in performance are solely attributed to the adversarially perturbed DDoS samples, and not to a change in benign samples.

C. Evaluation metrics

We evaluate the classification accuracy using the F1 score, which combines both the recall and the precision and hence provides an overall measure of the model’s performance. The robustness can also be indicated by the false negative rate (FNR), which is the number of misclassified DDoS samples to the total number of DDoS samples in a test dataset. These metrics are formally defined as follows:

⁵This enables comparison with the original LUCID paper when using the UNB201X dataset.

$$Pr = \frac{TP}{TP + FP} \quad Re = \frac{TP}{TP + FN}$$

$$FNR = \frac{FN}{FN + TP} \quad F1 = 2 \cdot \frac{Pr \cdot Re}{Pr + Re}$$

where Pr =Precision, Re =Recall, $F1$ =F1 Score, FNR =False Negative Rate, TP =True Positives, FP =False Positives, FN =False Negatives.

VI. EXPERIMENTAL RESULTS

This section presents the detection results of the models trained using the baseline approaches and our adversarial training approach. We employ these models to classify network traces with perturbed DDoS traffic. Moreover, it is important to stress that the GAN is only used to generate adversarial samples for training. In the following, we use the difference Δ in the evaluation metric, e.g. F1 score or FNR, to show the effect of using GADoT on the robustness of the trained models. This difference is formally defined as: $\Delta = Result(After) - Result(Before)$.

A. Evaluation in normal settings

First, we want to measure any negative effect of adversarial training with GADoT on LUCID when tested on unperturbed data. Thus, we want to ensure that training with adversarial samples does not cause any degradation to the performance of LUCID. In this regard, Table V shows the detection results of the trained models on unperturbed test datasets before and after using GADoT. The results of the baseline approaches are omitted from the table due to space limitation.

Before using GADoT, the LUCID models are capable of detecting DDoS flows with F1 scores above 99%, which aligns with the results already reported in [4]. After training with each of the baseline approaches, the models also maintain their detection accuracy with F1 scores above 99% in all datasets. When trained with GADoT, the models maintain high accuracy scores, despite a drop in the F1 score of less than 1.5%. Notably, we observe a slight drop in precision on the SYN datasets, and in recall score on the CICIDS2017 dataset. This indicates that a few ambiguous network flows are misclassified, which matches the observations made in [17] after adversarial training and [32] after employing ensemble learning to improve robustness.

B. Evaluation in adversarial settings

In what follows, we test Model-SYN and Model-HTTP on perturbed flood attacks that can be generated according to the attacker’s knowledge as defined in our threat model (Section III). Moreover, we demonstrate the positive impact of using GADoT on the models’ robustness.

1) **Model-SYN**: The F1 scores of Model-SYN before and after using GADoT for training, and tested on perturbed samples, are shown in Table VI.

With no adversarial training, we observe a substantial sensitivity of the model to adversarial SYN flood attacks built with perturbations on *IP flags*, and *Padding replacement*, which leads to a drop in the F1 score to 65.96%, and 78.75%,

TABLE V
EVALUATION OF MODEL-SYN AND MODEL-HTTP AGAINST UNPERTURBED TEST DATASETS BEFORE AND AFTER USING GADoT

Dataset (Model Name)	Scapy-SYN (Model-SYN)			CICIDS2017 (Model-HTTP)		
	Before	After	Δ	Before	After	Δ
Precision	0.9985	0.9739	-0.0246	0.9873	0.9902	0.0029
Recall	1.0000	1.0000	0.0000	0.9990	0.9978	-0.0012
F1 score	0.9992	0.9867	-0.0125	0.9931	0.9940	0.0009
FNR	0.0000	0.0000	0.0000	0.0009	0.0021	0.0012

respectively. The results obtained with *IP flags* confirm the high ranking of this feature in the analysis presented in the original LUCID paper [4]. On the other hand, the impact of manipulation of the *Flow length* feature through *Padding replacement* can be explained by analysing the properties of the SYN flood attack traffic. Indeed, the *Flow length* of the SYN samples is either one or two packets (the SYN packet of the attacker and the SYN-ACK from the victim server, if its connection tables are not full), while benign samples in our datasets are longer on average. This difference becomes less prominent when the *Flow length* feature of DDoS samples are perturbed with *Padding replacement*, leading to a decrease in the classification accuracy.

These results demonstrate that perturbing the *IP flags* and replacing the padding rows with dummy packets can generate a successful evasion attack. We also identify that similar success can be achieved by perturbing a combination of features in the same SYN network flow. For instance, while introducing *SYN Packet Replication* perturbation has decreased the F1 score to 96%, combining it with the *IP flags* and *TCP length* perturbations yields a further 33% F1 score decrease, producing a stronger evasion attack.

We also observe that adversarial training might produce a decrease in the classification accuracy when perturbed samples are present. This is indicated by the negative Δ values in Table VI. Although minimal, the decrease is caused by the change in the distribution of the features due to the adversarial samples added to the training set. This change moves the decision boundary, and hence the accuracy scores, F1 included.

Adversarial training with the BFP baseline approach allows the Model-SYN to restore most of its detection accuracy with the majority of F1 scores $> 97\%$ except for the *IP Flags* perturbation. This indicates that adversarial training with this approach leaves the model vulnerable to perturbation of features not present in the benign samples. The FGSM-based approach leads to a minor improvement in robustness against most of the perturbations. This minor increase in robustness suggests that FGSM or similar methods (e.g., Projected Gradient Descent (PGD)) might not be the best option to generate adversarial examples for the training of NIDS. On the other hand, after retraining using GADoT, we obtain an overall higher classification accuracy ($> 98\%$) and a substantial reduction in the false negative rate, which drops to reach approximately 0% for all perturbations. The results highlight that the adversarial dataset T_{adv} created with GADoT covers the

TABLE VI
EVALUATION OF MODEL-SYN AGAINST PERTURBED TRACES OF THE SCAPY-SYN DATASET

Perturbations	Before		GADoT				BFP				FGSM			
	F1 Score	FNR	F1 Score	Δ	FNR	Δ	F1 Score	Δ	FNR	Δ	F1 Score	Δ	FNR	Δ
IP Flags	0.6596	0.5071	0.9867	0.3271	0.0003	-0.5068	0.6558	-0.0038	0.5071	0.0000	0.6541	-0.0055	0.5071	0.0000
TCP Len	0.9992	0.0000	0.9867	-0.0125	0.0000	0.0000	0.9953	-0.0039	0.0000	0.0000	0.9933	-0.0059	0.0000	0.0000
SYN Packet Replication	0.9547	0.0851	0.9867	0.0320	0.0000	-0.0851	0.9951	0.0404	0.0000	-0.0851	0.9933	0.0386	0.0000	-0.0851
Padding Replacement	0.7875	0.3494	0.9867	0.1992	0.0000	-0.3494	0.9757	0.1882	0.0382	0.3112	0.9127	0.1252	0.1491	-0.2003
IP Flags; TCP Len; SYN Packet Replication	0.6322	0.5368	0.9846	0.3524	0.0041	-0.5327	0.9888	0.3566	0.0130	-0.5238	0.9066	0.2744	0.1596	-0.3772
IP Flags; TCP Len; Padding Replacement	0.5862	0.5843	0.9867	0.4005	0.0000	-0.5843	0.9731	0.3869	0.0431	-0.5412	0.5999	0.0137	0.5657	-0.0186

TABLE VII
EVALUATION OF MODEL-HTTP AGAINST PERTURBED TRACES OF THE CICIDS2017 DATASET

Perturbations	Before		GADoT				BFP				FGSM			
	F1 Score	FNR	F1 Score	Δ	FNR	Δ	F1 Score	Δ	FNR	Δ	F1 Score	Δ	FNR	Δ
Delay	0.5516	0.6150	0.9871	0.4355	0.0177	-0.5973	0.9038	0.3522	0.1659	-0.4491	0.8105	0.2589	0.3124	-0.3026
Packet Fragmentation	0.9406	0.1024	0.9935	0.0529	0.0000	-0.1024	0.9879	0.0473	0.0112	-0.0912	0.8833	-0.0573	0.2019	0.0995

space of adversarial perturbations that attackers might exploit. This also indicates that Model-SYN is more robust to adversarial attacks with adversarial training using GADoT than without.

2) *Model-HTTP*: As noted in Section IV-A, packet arrival time and *TCP Len* are important features in the detection model. We therefore measure the impact of delay perturbation and packet fragmentation on the detection accuracy. Packet fragmentation perturbs the *TCP Len* feature of each packet and increases the number of packets in each flow as well. Indeed, packet arrival time, packet size and the number of packets in a flow are features employed to generate adversarial network flows across the state-of-the-art, e.g., [7], [13], [14], [17], [18]. We apply the perturbations to the DDoS packets of the HTTP-GET flows in the CICIDS2017 dataset. As shown in Table VII, with no adversarial training, the HTTP GET attack with a delay perturbation generates a strong evasion attack, as the FNR increases to 61.5%, confirming the importance of this feature for the detection model. On the contrary, packet fragmentation generates a weaker evasion attack with FNR around 10%. The reason is that the attack flows naturally comprise several packets, and hence packet fragmentation has a less noticeable impact on the attack flows.

With adversarial training using the baseline approaches, it is evident that the trained models remain vulnerable to DDoS attacks with perturbed samples. The better of the two models, which is the BFP approach, is still vulnerable to the delay perturbation as the F1 score and FNR are approximately 90% and 16% respectively. After using GADoT for training, the detection accuracy of the model in the case of delay perturbation experiences more than 43% improvement in the F1 score, reaching 98.71%. The same improvement is also apparent in the FNR that drops to 1.8%, indicating the increased resilience to adversarial perturbations. A similar level of resilience is also noticed in the case of the packet fragmentation as the FNR has dropped to 0%. This highlights that, in comparison with the baseline approaches, GADoT is able to create adversarial

training dataset T_{adv} with a wide coverage of adversarial perturbations. These results, combined with those of Model-SYN, demonstrate the considerable improvement in robustness after using GADoT.

C. Discussion

The results presented in the previous sections confirm that ML-based NIDSs are vulnerable to adversarial attacks. Potential adversaries can exploit their domain expertise to craft perturbed flood attacks without requiring knowledge of the underlying detection model. For example, they can perturb the delay in the network flows of the HTTP GET flood attack, as we have done in the CICIDS2017 dataset, which causes the FNR (misclassified DDoS samples) to reach 61.5%. Building upon this FNR, they can increase the number of bots (botnet DDoS attacks) to exhaust the victim resources and achieve a successful DDoS attack. In contrast, after adopting GADoT, most of the perturbed flood samples are detected, with the FNR reduced to below 1.8%.

The models trained using GADoT achieve high performance with F1 scores above 98% on perturbed DDoS network traces, improving on the 62% average detection rate reported in [14]. We highlight that the models trained using GADoT have detected the SYN and HTTP GET flood attacks, whether they are perturbed or not, allowing for the mitigation of such attacks. Moreover, the architecture of those models remains unchanged, and thus GADoT has no impact on their performance in terms of processed samples/second and memory requirements.

Finally, a consideration for practical deployment of a ML system such as GADoT is prediction interpretability [33]. We undertook a simple study of how our trained model classifies samples after adversarial training by exploring the filter activations of a convolution layer while processing samples. This revealed a change in the filter activation across different features as well as a shift in their ranking, and hence their contribution to the model's decision. For example, our initial

analysis indicated that there is no longer a heavy dependence on a single feature such as *Highest layer*, as reported in [4], which suggests to us that the adversarial training evens the activation across multiple features, which results in the increased robustness to evasion attacks. However, this is not conclusive and requires an in-depth analysis, which we propose for our future work.

VII. CONCLUSION

In this paper, we introduced GADoT, an approach for adversarial training of ML models detecting DDoS flood attacks. We trained two LUCID models, which are state-of-the-art for DDoS detection, to detect two types of flood attacks, namely, SYN flood and HTTP GET flood. However, these models, as is the case for other ML models, are vulnerable to adversarially perturbed DDoS samples and traces that aim to evade such models, reaching the victim machine and achieving the DoS goal. To defend against such adversarial attacks, we trained both models using GADoT, which leverages a GAN to generate adversarial DDoS samples for training. The trained models were evaluated on real perturbed traces generated without requiring knowledge of the victim model. This approach sets GADoT apart from other similar solutions. The results show that the models trained with GADoT can detect with high accuracy both the SYN and HTTP flood attacks, regardless of whether they are perturbed or not. The models achieved F1 score above 98% and the FNR drops to below 1.8% on perturbed DDoS network traces. Future work will include extending GADoT to other DDoS attack types and studying the use of open source tools, such as LOIC, to generate perturbed DDoS attacks.

REFERENCES

- [1] Cisco, "Cisco annual internet report (2018–2023)," <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html>, access: 2020-08-05.
- [2] A10, "Aws hit by largest reported ddos attack of 2.3 tbps," <https://www.a10networks.com/blog/aws-hit-by-largest-reported-ddos-attack-of-2-3-tbps/>, access: 2020-08-05.
- [3] Y. Mirsky, T. Doitshman, Y. Elovici, and A. Shabtai, "Kitsune: an ensemble of autoencoders for online network intrusion detection," *arXiv preprint arXiv:1802.09089*, 2018.
- [4] R. Doriguzzi-Corin, S. Millar, S. Scott-Hayward, J. Martinez-del Rincon, and D. Siracusa, "Lucid: A practical, lightweight deep learning solution for ddos attack detection," *IEEE Transactions on Network and Service Management*, 2020.
- [5] M. S. Elsayed, N.-A. Le-Khac, S. Dev, and A. D. Jurcut, "Ddosnet: A deep-learning model for detecting network attacks," in *2020 IEEE 21st International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. IEEE, 2020, pp. 391–396.
- [6] A. Kuppaa, S. Grzonkowski, M. R. Asghar, and N.-A. Le-Khac, "Black box attacks on deep anomaly detectors," in *14th International Conference on Availability, Reliability and Security*. ACM, 2019, pp. 1–10.
- [7] J. Aiken and S. Scott-Hayward, "Investigating adversarial attacks against network intrusion detection systems in sdns," in *2019 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*. IEEE, 2019, pp. 1–7.
- [8] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, "Ensemble adversarial training: Attacks and defenses," *arXiv preprint arXiv:1705.07204*, 2017.
- [9] F. Pierazzi, F. Pendlebury, J. Cortellazzi, and L. Cavallaro, "Intriguing properties of adversarial ml attacks in the problem space," in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 1332–1349.
- [10] Imperva, "TCP SYN Flood," <https://www.imperva.com/learn/ddos/syn-flood/>, access: 2020-30-11.
- [11] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in *ICISSP*, 2018, pp. 108–116.
- [12] Imperva, "HTTP Flood," <https://www.imperva.com/learn/ddos/http-flood/>, access: 2020-30-11.
- [13] Q. Yan, M. Wang, W. Huang, X. Luo, and F. R. Yu, "Automatically synthesizing dos attack traces using generative adversarial networks," *International Journal of Machine Learning and Cybernetics*, vol. 10, no. 12, pp. 3387–3396, 2019.
- [14] M. J. Hashemi and E. Keller, "Enhancing robustness against adversarial examples in network intrusion detection systems," *arXiv preprint arXiv:2008.03677*, 2020.
- [15] R. A. Khamis and A. Matrawy, "Evaluation of adversarial training on different types of neural networks in deep learning-based ids," *arXiv preprint arXiv:2007.04472*, 2020.
- [16] R. Abou Khamis, M. O. Shafiq, and A. Matrawy, "Investigating resistance of deep learning-based ids against adversaries using min-max optimization," in *ICC 2020-2020 IEEE International Conference on Communications (ICC)*. IEEE, 2020, pp. 1–7.
- [17] C. Zhang, X. Costa-Pérez, and P. Patras, "Tiki-taka: Attacking and defending deep learning-based intrusion detection systems," in *2020 Cloud Computing Security Workshop (CCSW'20)*. ACM, 2020.
- [18] M. J. Hashemi, G. Cusack, and E. Keller, "Towards evaluation of nids in adversarial setting," in *Proceedings of the 3rd ACM CoNEXT Workshop on Big Data, Machine Learning and Artificial Intelligence for Data Communication Networks*, 2019, pp. 14–21.
- [19] M. Usama, M. Asim, S. Latif, J. Qadir *et al.*, "Generative adversarial networks for launching and thwarting adversarial attacks on network intrusion detection systems," in *15th International Wireless Communications & Mobile Computing Conference (IWCMC)*. IEEE, 2019, pp. 78–83.
- [20] B. Biggio and F. Roli, "Wild patterns: Ten years after the rise of adversarial machine learning," *Pattern Recognition*, vol. 84, pp. 317–331, 2018.
- [21] MITRE, "Adversarial ml threat matrix," <https://github.com/mitre/advmthreatmatrix>, access: 2020-10-27.
- [22] K. Yang, J. Liu, C. Zhang, and Y. Fang, "Adversarial examples against the deep learning based network intrusion detection systems," in *2018 IEEE Military Communications Conference (MILCOM)*. IEEE, 2018, pp. 559–564.
- [23] R. Doriguzzi-Corin, "LUCID source code," <https://github.com/doriguzzi/lucid-ddos>, access: 2020-30-11.
- [24] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," in *Advances in neural information processing systems*, 2017, pp. 5767–5777.
- [25] "Generative adversarial networks," <https://developers.google.com/machine-learning/gan/generator>, access: 2020-11-25.
- [26] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [27] S. Sanfilippo, N. Jombart, D. Ducamp, Y. Berthier, and S. Aubert, "Hping," www.hping.org, access: 2020-08-05.
- [28] A. Shiravi, H. Shiravi, M. Tavallae, and A. A. Ghorbani, "Toward developing a systematic approach to generate benchmark datasets for intrusion detection," *computers & security*, vol. 31, no. 3, 2012.
- [29] P. Biondi, "Scapy: explore the net with new eyes," *Technical report, EADS Corporate Research Center*, 2005.
- [30] "Loic," <https://www.imperva.com/learn/ddos/low-orbit-ion-cannon>, access: 2020-08-05.
- [31] M. Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, <http://tensorflow.org/>, access: 2020-12-04.
- [32] G. Apruzzese, M. Andreolini, M. Marchetti, V. G. Colacino, and G. Russo, "Appcon: Mitigating evasion attacks to ml cyber detectors," *Symmetry*, vol. 12, no. 4, p. 653, 2020.
- [33] A. Jacovi, O. S. Shalom, and Y. Goldberg, "Understanding convolutional neural networks for text classification," *arXiv preprint arXiv:1809.08037*, 2018.