

NetSentry: A Deep Learning Approach to Detecting Incipient Large-scale Network Attacks

Haoyu Liu^a, Paul Patras^a

^a*School of Informatics, The University of Edinburgh, Informatics Forum, 10 Crichton St, Newington, Edinburgh, EH8 9AB, United Kingdom*

ARTICLE INFO

Keywords:

Network-based Intrusion Detection System
Deep Learning
Feature Augmentation

ABSTRACT

ML techniques are increasingly adopted to tackle ever-evolving high-profile network attacks, including Distributed Denial of Service (DDoS), botnet, and ransomware, due to their unique ability to extract complex patterns hidden in data streams. These approaches are however routinely validated with data collected in the same environment, and their performance degrades when deployed in different network topologies and/or applied on previously unseen traffic, as we uncover. This suggests malicious/benign behaviors are largely learned superficially and Machine Learning (ML)-based Network Intrusion Detection Systems (NIDS) need revisiting, to be effective in practice. In this paper we dive into the mechanics of large-scale network attacks, with a view to understanding how to use ML for Network Intrusion Detection (NID) in a principled way. We reveal that, although cyberattacks vary significantly in terms of payloads, vectors and targets, their *early stages*, which are critical to successful attack outcomes, *share many similarities and exhibit important temporal correlations*. Therefore, we treat NID as a time-sensitive task and propose NetSentry, perhaps the first of its kind NIDS that builds on Bidirectional Asymmetric LSTM (Bi-ALSTM), an original ensemble of sequential neural models, to detect network threats before they spread. We cross-evaluate NetSentry using two practical datasets, training on one and testing on the other, and demonstrate F1 score gains above 33% over the state-of-the-art, as well as up to 3× higher rates of detecting attacks such as Cross-Site Scripting (XSS) and web brute force. Further, we put forward a novel data augmentation technique that boosts the generalization abilities of a broad range of supervised deep learning algorithms, leading to average F1 score gains above 35%. Lastly, we shed light on the feasibility of deploying NetSentry in operational networks, demonstrating affordable computational overhead and robustness to evasion attacks.

1. Introduction


The volume of illicit network traffic continues to grow dramatically, with the number of high-profile attacks including DDoS, botnet, and ransomware rising by over 45% annually [1], and the losses incurred expected to exceed 6 trillion US dollars in 2021 [2]. Effective countermeasures to thwart ever-evolving cyber threats are therefore urgently needed. Traditional Network Intrusion Detection Systems (NIDS) largely apply finite rules, preset by human experts, to detect anomalies. This approach lacks flexibility and is often prone to subversion [3]. Machine Learning (ML) is increasingly used to detect cyber intrusions, due to its ability to discover complex statistical patterns hidden in data streams, which can aid in discriminating anomalies based on feature differences [4].

ML is a powerful tool, yet adopting it meaningfully for security purposes is not straightforward. ML techniques used in areas including imaging and natural language processing have been directly applied to NID (e.g., [5]), without adequate analysis of their suitability for this task. For instance, reconstruction-based algorithms like autoencoders were originally designed to learn to recreate benign samples that contain similar patterns, e.g., the same object type in images [6]. However, when deployed for intrusion detection, whether an autoencoder is able to reconstruct heterogeneous benign traffic originating from various applications is rarely

discussed [7]. Secondly, widely-used evaluation methodologies involve training and testing NID models on a single dataset, collected in the same controlled environment. This makes it difficult to assess if the trained models can truly generalize to previously unseen traffic mixes [8]. Moreover, detecting high-volume attacks promptly, before a target system becomes overloaded and unable to thwart malicious traffic with potential to cause severe damage following early system compromise, is difficult. This capability is however critical to the availability and revenue of online businesses [9].

In this paper, we address the above challenges and propose **NetSentry**, a novel Deep Learning (DL)-based NIDS that reliably detects a range of malicious traffics with similar patterns, indicative of incipient high-impact network attacks. As such, we make the following **key contributions**:

- (1) We scrutinize several attack chains and identify key temporal inter-relations between illicit traffic occurring in the wild; based on this analysis, we design Bidirectional Asymmetric LSTM (Bi-ALSTM), an original ensemble of sequential neural models that effectively captures the temporal dynamics of malicious traffic and classifies specific threats, including Denial of Service (DoS), Port Scanning, and Brute Forcing;
- (2) Since not every attack type can be distinguished accurately with limited information available at the network layer, we introduce a novel training technique that relies on feature augmentation and abstract labeling. The feature augmentation scheme improves the heterogeneity of cyber attacks that were collected in a

 haoyu.liu@ed.ac.uk (H. Liu); paul.patras@ed.ac.uk (P. Patras)
ORCID(s):

controlled environment, which helps NN models learn a more robust decision boundary. Abstract labeling, on the other hand, prevents overfitting by grouping similar types of attacks into one class;

- (3) We train our Bi-ALSTM on a large dataset published by the Canadian Institute for Cybersecurity, we cross-evaluate our approach with a previously unseen dataset collected in a different network topology, and we compare its performance against that of state-of-the-art benchmarks. Results demonstrate Bi-ALSTM outperforms existing approaches by at least 33% in terms of F1 score;
- (4) We discuss practical aspects of deploying NetSentry in real-life, including computational overhead and robustness to a range of evasion attacks.

To our knowledge, NetSentry is perhaps the first principled DL-based NIDS that tackles cyberthreats by focusing on the early stages that are essential to the success of large-scale and high-impact attacks such as botnet and ransomware.

2. Threat Model & Anatomy of Attacks

We start from the key observation that in practice traffic flows shall not be considered in isolation, either as benign or malicious. There exist important temporal correlations among different cyber attacks, especially those with high-impact, which rarely occur independently. For instance, assume that an adversary has zero knowledge of a potentially vulnerable target. Conducting a successful webshell injection attack has at least two pre-requisites: (i) port scanning against the target, so as to uncover that it runs a web serve; and (ii) web API enumerating, to verify if file upload is allowed. That is, the attacker must follow a certain sequence of actions (each an attack itself), which would create distinct network traces at various stages. We remark that essential correlations among different kinds of network attacks have not been explored before, but are potentially useful to design a reliable NIDS.

Hence, we decompose network attacks from the perspective of an active adversary, and summarize them into different attack chains. These typically start with gathering information about a target and conclude when a specific technical goal was achieved. We consider three key attack chains, namely botnet, web intrusion, and ransomware, revealing that they are supported by a similar methodology. While our attack chain view may appear on the surface related to earlier Botnet infection modeling, where attack stages are fingerprinted [10], our modeling approach and subsequent NIDS design are fundamentally different. This is because *the attack chains we consider aim to reveal the common stages shared by different large-scale cyber attacks*, so as to impede a specific range of cyber intrusions by interrupting any of these early stages.

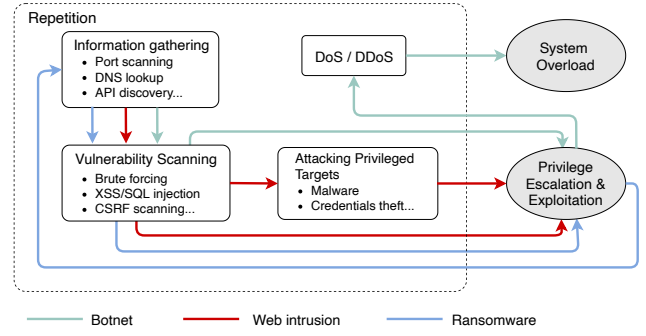


Figure 1: Attack chains employed by large-scale high-impact threats, e.g., botnets, web intrusion, and ransomware. Observe that similar steps are repeated by all, which **NetSentry** exploits for detection.

2.1. Attack Chain Analysis

We particularly focus on two network attack goals, which bring severe damage to target systems. The first is to obtain system privileges temporarily or permanently, by exploiting various security flaws. The second is to overload the system by occupying all its resources. Instead of looking at each attack type individually, we investigate what processes, i.e., attack chains, an adversary must follow to achieve any of these goals, when having zero knowledge about a target. We consider three unique attack chains that are specific to botnet, web intrusion, and respectively ransomware, as shown in Figure 1.¹

Botnet/Mirai. Botnets are collections of Internet-accessible devices hijacked by an attacker, and are usually employed to carry out large-scale, high-impact DDoS attacks. Mirai is one of the most notorious instances in recent years, with ~600,000 devices infected at its peak [12]. Subsequent Mirai variants expand the attack surface to SSH, HTTPS, FTP, etc., but inherit the methodology of the canonical version.

Mirai follows a chain-like methodology that entails *information gathering* → *vulnerability scanning* → *privilege escalation* → *DDoS*. Specifically, (1) TCP SYN packets are sent towards the entire IPv4 address space, on ports 23 and 2323 (Telnet); (2) after identifying potential victims, Mirai attempts to bruteforce the victims' credentials using a dictionary – this process is deemed as vulnerability scanning; (3) upon successful login, a victim's IP and credentials are forwarded to a report server that infects the victim with malicious code; (4) newly infected devices become members of the botnet, and either participate in victim discovery or DDoS attacks [12].

Web intrusion. Web applications integrate a technology stack that includes storage, web engines, Operating Systems (OSs), and communications. Hence, various vulnerabilities are often exploited. Web intrusion can be jointly modeled with the Intrusion Kill Chain [13] and the OWASP Web Penetration Guideline [14], where the former outlines a general intrusion process, while the latter provides specific attack vectors.

¹We note that certain sophisticated attacks may have longer attack chains that expand to application level [11]. However, by detecting their early stages, their later exploitation actions can be prevented.




Figure 2: NetSentry architecture. Feature Augmentor only applied during training. \oplus is fusion operation detailed in Eq. 11.

The attack process entails *information gathering* \rightarrow *vulnerability scanning* \rightarrow *attacking privileged targets* \rightarrow *exploitation*; or *information gathering* \rightarrow *vulnerability scanning* \rightarrow *privilege escalation*. Web intrusion resembles bot-net and ransomware in terms of vulnerability discovering approach, but differs at the later stage, since exploitation is always target-specific, e.g. a XSS attack targets web app users, while SQL injection targets Web APIs. Our focus is on the early stage when the attacker tries to breach a trusted boundary, since later steps occur at system level and are invisible to a NIDS.

Ransomware/WannaCry. Ransomware is a relatively new type of threat that blocks user access to their private data until a ransom is paid to the attacker. For instance, by exploiting EternalBlue [15], WannaCry gains system access via the Server Message Block (SMB) protocol on Windows systems, encrypts user data, and spreads itself to other hosts [16]. The attack chain of WannaCry follows a loop consisting of: *information gathering* \rightarrow *vulnerability scanning* \rightarrow *privilege escalation & exploitation* \rightarrow *information gathering*.

We focus on the procedure ransomware uses to discover and control new victims (instead of the file encryption applied), as the attack is conducted via the network. WannaCry employs repeated TCP scanning on port 445 (serving the SMB protocol) [16]. Targets are further fingerprinted and remote access is achieved by injecting code via a crafted packet, which would be mishandled by SMBv1. WannaCry then encrypts the data on the victim machine and discovers other vulnerable targets.

3. NetSentry Design

In what follows we present **NetSentry**, an original NIDS design that harnesses the unique feature extraction capabilities of recurrent neural models, to detect large-scale, high-impact attacks. NetSentry builds on our observation of correlations between malicious traffics, and handles NID as a time-sensitive task, leveraging an ensembling structure to capture richer contexts and detect intrusions with high efficacy.

3.1. Attack Detection Strategy

Our attack chains analysis revealed that information gathering, vulnerability scanning and DoS are applied across various types of attacks and share the same semantic. Latent network attacks, such as malware downloading, code

injection, Cross-Site Request Forgery (CSRF), and other zero-day attacks, which can obtain system privileges and proceed to exploitation, always follow massive vulnerability scanning, since this is the most efficient way to discover weak entry points. A common argument is that zero-day attacks are heterogeneous, which poses difficulties to any detection logic. However, we argue that *as long as automated activities can be recognized in time, the subsequent zero-day attacks can be blocked*, in order to minimize the chances that attackers may uncover weaknesses and compromise a system. As such, we keep the scope of our detection targets narrow, yet well-directed, as suggested in [8]. With this in mind, we design NetSentry, a NIDS that effectively tackles cyber intrusion by detecting risks at an *early stage*. This also applies to tactics that deviate from the standard attack chains described, as long as they incorporate any common stages to achieve the same end goal.

We maintain that NID, especially of automated attacks, should be treated as a time-sensitive task. Here we consider ‘time-sensitive’ those network intrusions exhibiting temporal correlations among *consecutive traffic flows*, which could potentially exert substantial impact on the decision-making process. This is because a single traffic flow, whose features are extracted as a datum, may not fully reflect the intention of the communications. A straightforward example is a TCP flow encapsulating a complete HTTP request. Assume the flow is terminated quickly after the server responds. Without looking at previous and subsequent traffic, it is impossible to assert whether the flow was initiated by a legitimate user or by DoS tools. Conversely, if we observe a series of statistically similar communications between a pair of hosts, the confidence of classifying them as malicious becomes higher. Network attacks generated with the same tool usually serve the same purpose. Although they would encapsulate different payloads in consecutive flows for obfuscation or fuzzing purposes, this difference is invisible to a NIDS that only has access to protocol headers and timing information. Thus, we leverage sequential neural models in NetSentry to learn such similarity of successive flows generated with automated tools.

3.2. System Architecture

NetSentry is a NIDS that examines the statistical features of network flows and detects illicit traffic via an ensemble of sequential neural models. A traffic flow is built by grouping packets according to a five-tuple (Src IP, Dst IP, Src Port, Dst Port, Protocol). Recall that automated tools tend to initiate

| Feature Type | Direction | Name |
|----------------|--------------|--|
| timing-based | forward | Flow Inter-arrival Time (IAT)*, packets/sec |
| | backward | Flow IAT*, packets/sec |
| | bi-direction | duration, Flow IAT*, packets/sec, bytes/sec, active time*, idle time* |
| protocol-based | forward | # packets, packet length*, PSH counts, URG counts, header length, initial TCP window size, avg segment size, subflow† |
| | backward | # packets, packet length*, PSH counts, URG counts, header length, initial TCP window size, avg segment size, subflow†, |
| | bi-direction | packet length*, flag counts§, down/up ratio, protocol |
| ID-based | None | flow ID, src IP, dst IP, src port, dst port, timestamp |

Table 1

Features used in NetSentry. * means (min, max, avg, std) are computed for a given property. † indicates where (avg packets, avg bytes) are computed. § indicates (FIN, SYN, RST, PSH, ACK, URG, CWE, ECE) are counted in flows.

multiple almost identical flows towards targets during a short period of time. This means that the statistical features of malicious flows share a large degree of similarity. Thus, monitoring the similarities and discrepancies of consecutive flows between pairs of hosts plays an essential role in recognizing anomalies. To learn relevant temporal correlations of the traffic flows and to differentiate malicious patterns, NetSentry incorporates 4 key building blocks, as shown in Figure 2, namely:

- Flow Aggregator & Feature Extractor: groups packets into flows and extracts associated statistical features;
- Sequence Generator: groups flows originating from the same pair of hosts into fixed-length sequences, to be fed as inputs to anomaly detection logic;
- Feature Augmentor: increases the variability of a fraction of malicious traffic features that are non-essential in anomaly detection, but if left unchanged may increase the risk of model becoming trapped in local optima;
- Anomaly Detector (Bi-ALSTM): an ensemble of two asymmetric Long Short-Term Memory (LSTM)-based neural networks operated bidirectionally, taking flow sequences as input to detect malicious traffic.

Next, we explain the inner workings of each component, then detail our bidirectional sequential neural model in §4.

3.2.1. Feature Extraction

NetSentry employs a two-step process to extract numerical or categorical information (features) of the traffic observed, i.e., packet grouping and statistics computation. The former involves aggregating into flows packets generated between same pairs of applications, which can be achieved by monitoring origin, destination, and protocol fields.

Since NetSentry operates at the network layer and is not guaranteed to have access to packet payloads, we confine consideration to features that encompass timing statistics and protocol information. We find that employing popular open-sourced tools for feature extraction, such as CICFlowMeter [17] (which should be able to extract 80+ statistical features), is problematic. Indeed, CICFlowMeter

uses a faulty mechanism to identify the end of TCP flows, which results in benign traffic often being mislabeled as malicious, and vice versa.

With the CICFlowMeter feature extractor, if a new incoming TCP packet has a FIN flag set, the packet is immediately deemed to be the last packet in that *flow*. Obviously, this does not strictly follow the four-way handshake of TCP connection teardown. We show in Figure 3 that the premature assessment of termination leads to mislabeling, which is especially relevant to automated attacks such as DoS.

Assume that A is performing simple HTTP DoS attacks targeting B, quickly reusing a same port 8888. Also assume each time it is B who decides to terminate the TCP connections. Then, applying the mechanism described above on two consecutive flows from (A, 8888) to (B, 80) would generate 4 complete flows and an incomplete one. In this case, any flows from (A, 8888) to (B, 80) should be labeled as DoS. However, Flow #2 and #4 only consist of two packets (ACK and FIN) which cannot reflect any malicious purpose, while Flow #3 that should be labeled as malicious, is marked benign because of its wrongly perceived direction. We confirm that this type of mislabeling occurs for DoS-Hulk attacks in the publicly available CSE-CIC-IDS-2018 dataset [18] (which we use after correct relabeling), but further instances may exist.

We fix this logic error (along with other programming bugs encountered) in CICFlowMeter, by following a complete four-way handshake to terminate TCP flows. The timeout mechanism is preserved for stateless protocols. We also note that the original tool further extracts partial features that are not well-defined. Hence we revise the code and only output 69 features per flow, as summarized in Table 1.²

3.2.2. Sequence Generation

Network anomaly is often nuanced, as a single network flow may be benign on its own, but observing multiple similar instances active at the same time may strongly suggest an automated attack in progress. Therefore, it is necessary to observe consecutive flows between hosts (applications) to further confirm malicious activity.

²We will release our feature extractor's source code upon publication.

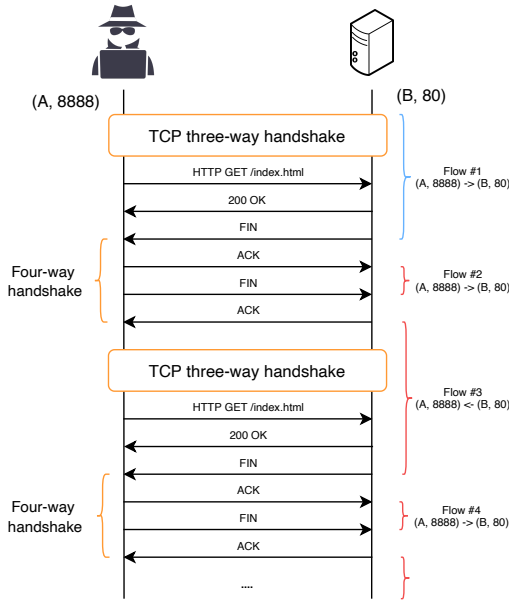


Figure 3: Incorrect flow labeling due to wrong TCP termination rules. Blue curly braces indicate the flow is labeled correctly; red curly braces indicate mislabeling.

Algorithm 1 NetSentry’s sequence generation algorithm incorporating sliding windows and timeout thresholding.

```

1: Inputs:
   Timeout value  $\tau$ , and window size  $\alpha$ 
2: Initialisation:
   conn_table: A connection table storing the
   incoming flows from the feature extractor.
   seq_list: A FIFO list buffering the inputs for the
   feature augmentor or the anomaly detector.
3: while True do
4:   start_time  $\leftarrow$  now()
5:   for flow from the feature extractor do
6:     conn_table[flow.id].add(flow)
7:     if len(conn_table[flow.id]) >  $\alpha$  then
8:       seq_list.add(conn_table.remove(flow.id))
9:     end if
10:    if now() - start_time >  $\tau$  then
11:      seq_list.addAll(conn_table.removeAll())
12:      start_time  $\leftarrow$  now()
13:    end if
14:  end for
15: end while

```

As such, a **sequence** in NetSentry is defined as successive flows using the same protocol between a pair of hosts, that is aggregated by (Src IP, Dst IP, Protocol). This is because with automated attacks, many-to-one (DoS, brute forcing) and many-to-many (port scanning) port attacks between a pair of hosts are common. Therefore, we regard a sequence through grouping not by the tuple used for flow aggregation, but by the origin and destination addresses, along with protocol type.

We adopt a flexible approach to generating sequences, which is a combination of sliding window and timeout




Figure 4: Violin plots of 4 features before and after data augmentation. All values are normalized between 0 and 1.

thresholding techniques, as described by Algorithm 1. NetSentry allows two user-defined parameters for this purpose, namely window size α and timeout value τ , and maintains a connection table with two columns: *ID* and *seq_list*. The *ID* of each flow is a 3-tuple (Src IP, Dst IP, Protocol) and for each *ID*, a FIFO *seq_list* is maintained, storing the flows with the same *ID*. A newly generated flow is added to the *seq_list* determined by its *ID*. Once the length of any *seq_list* is larger than α , the elements in the list are regarded as a sequence to be passed to the neural model. Meanwhile, after every τ seconds, the entire connection table is emptied regardless of the length of the *seq_lists*. Any list whose length is less than α is padded to α for the purpose of alignment. This design is customizable: the larger α is, the more comprehensive the context that the ensemble model obtains, but the higher the memory requirements; the smaller τ is, the more timely classification can be achieved, at the cost of more compute resources.

3.2.3. Feature Augmentation

Since data used for ML training are largely collected in controlled environments, synthetically generated attacks may not offer an accurate view of network threats occurring in real-world [8], which prevents the model from learning a reliable decision boundary. For example, a victim HTTP server was set up to produce the CSE-CIC-IDS2018 dataset [18]; during HTTP DoS generation, during HTTP DoS generation, all flows encapsulated the same backward payloads from victim to attacker, resulting in little variability in payload-related features (see Figure 4, left). In reality, it is hard to predict how the victim would respond, and we show in §5 that such artificially low variability leads to poor generalization abilities for a range of supervised models.

We mitigate this problem by augmenting a collection of payload-related features to emulate a more realistic network environment. Specifically, we set up an HTTP victim server and an attacking client. The client only makes single requests with the Keep-Alive header over one TCP connection, to emulate HTTP DoS attacks. The size of each request and response is sampled from two discrete uniform distributions:

$$http_request\ size \sim \mathcal{U}(100, 400),$$

$$http_response\ size \sim \mathcal{U}(100, 15000).$$

In total, we generate 2,000 flows. A graphical illustration of the augmentation process is depicted in Figure 5. For

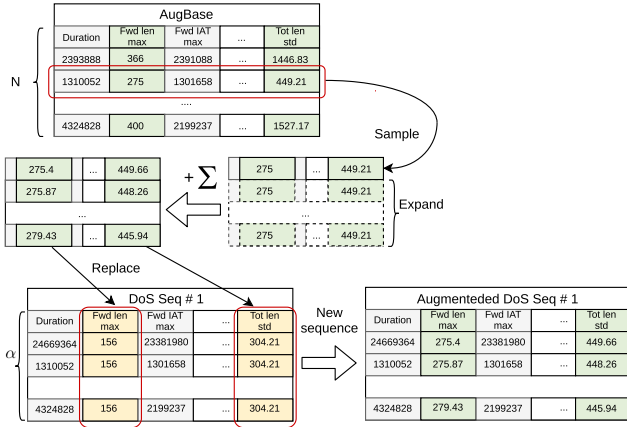


Figure 5: Illustration of the data augmentation process. Yellow cells denote payload features in original training data; green cells represents payload features from AugBase. Σ is a noise matrix with each element sampled from $\sigma \sim \mathcal{N}(0, 5)$.

each sequence that contains single-request HTTP DoS attacks (excluding Slowloris), (1) a flow is randomly sampled from the AugBase set and expanded to sequence length α ; (2) random noise $\sigma \sim \mathcal{N}(0, 5)$ is added to each payload-related cells to mimic minor differences among flows in a sequence; (3) finally, payload-related features in the original sequence are replaced by the new features generated at Step 2. By applying such augmentation, the new payload features of different flows in the same sequence would not differ much, but the features among different sequences would look considerably different. The distributions and the means of a subset of payload features in the augmented set are shown in Figure 4 (right). We use augmented data only for training.

It is also worth noting that the augmented data need not originate from any real traffic, since we only change parts of the features. However, what the model can learn from the augmented set is that (i) payload features possess high variability within some attacks (DoS in our case) whose logic does not rely on specific payloads, thus payload features should not be utilized for decision; and (ii) the rest of features (payload-irrelevant) are more valuable in distinguishing augmented attacks. We choose not to remove payload-related features altogether because they may be important for the model to differentiate other types of attacks, such as Slowloris, which only sends a small amount of payload during a long span. In §5, we demonstrate that augmentation boosts the performance of several supervised models.

3.2.4. Ensemble Network

The sequential ensemble neural network is the critical component of NetSentry and is responsible for detecting malicious traffic based on the inputs provided by the sequence generator or feature augmentor. As explained in §3.1, detecting automated cyber attacks is a time-sensitive task and hinges on temporal correlations between network flows. Let $X := \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_\alpha\}$ and $Y := \{y_1, y_2, \dots, y_\alpha\}$ denote a sequence of inputs and respectively their corresponding correct prediction. Then time-sensitive intrusion detection

can be formalized as

$$\tilde{\theta} = \arg \max_{\theta} P_{\theta} (y_1, y_2, \dots, y_\alpha | \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_\alpha),$$

where the sequential model is parameterized by θ .

NetSentry leverages two LSTM-based models to approximate the probability function above. In the next section, we first introduce the LSTM models we employ, and show both conceptually and empirically that an ensemble of such sequential models, preferably with different architectures, is key to improving the overall NID performance.

4. A Sequential Ensemble for NID

We overview the different blocks that lay foundations for our Bi-ALSTMs, then explain the ensembling and why our approach is essential for high-performance classification.

4.1. Long Short-Term Memory (LSTM)

As a variant of Recurrent Neural Networks (RNN), LSTM [19] incorporates gating functions to simulate the update of memory units along time and has shown excellent ability to model long-term dependencies in sequential data [20, 21]. An LSTM maintains two states: a cell state c_t and a hidden state h_t , which is computed based on inputs up to timestamp t , i.e., $\mathbf{x}_1, \dots, \mathbf{x}_t$ ($\mathbf{x}_i \in \mathbb{R}^d$). To maintain long-term dependencies, the LSTM cell has input (i_t), forget (f_t), and output (o_t) gates controlling the information flowing through at different timestamps. Gates are modeled by single-layer neural networks with parameters $W_{xi}, W_{hi}, W_{xf}, W_{hf}, W_{xo}, W_{ho} \in \mathbb{R}^{h \times d}$ and associated biases, i.e.,

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + b_i), \quad (1)$$

$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + b_f), \quad (2)$$

$$c_t = f_t \circ c_{t-1} + i_t \circ \tanh(W_{xc}x_t + W_{hc}h_{t-1} + b_c), \quad (3)$$

$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + b_o), \quad (4)$$

$$h_t = o_t \circ \tanh(c_t), \quad (5)$$

where σ denotes the sigmoid function and \circ represents element-wise product. When a new input x_t is given to the LSTM, the cell state c_t is updated with information from x_t and the previous cell state c_{t-1} . The proportion of x_t and c_{t-1} in the new cell state is determined by i_t and f_t , as in Eq. 3. h_t can be perceived as a non-linear transformations on c_t , and are always used for downstream tasks, such as classification. In NetSentry, an Multilayer Perceptron (MLP) $g_{\phi}(\cdot)$ is used to approximate the probability of y_t given h_t .

When it comes to intrusion detection against automated network attacks, the order of attack flows in a sequence is less important. Hence, at any timestamp t , both previous inputs and subsequent inputs may be equally highly correlated with x_t . The traditional LSTM can only model temporal information unidirectionally as time evolves. In other words, the hidden representation h_t only comprises the context before or at timestamp t . To generate more comprehensive hidden representations for anomaly detection, a Bidirectional LSTM (Bi-LSTM) which runs two LSTMs separately (one forward, one backward), and whose hidden

states are concatenated before given to $g_\phi(\cdot)$, can be used. The objective of Bi-LSTM is thus:

$$\begin{aligned} & \max_{W, b, \phi} p(y_1, y_2, \dots, y_a | x_1, x_2, \dots, x_a, W, b, \phi) \\ &= \max_{W, b, \phi} \prod_{t=1}^a p(y_t | x_1, \dots, x_a, W, b, \phi) = \max_{W, b, \phi} \prod_{t=1}^a g_\phi(h_t), \end{aligned}$$

where W and b are the weights and biases in LSTM cells and ϕ the parameters of the MLP. Bi-LSTM is one of the benchmarks we consider in evaluating our work.

4.2. ConvLSTM

Convolutional LSTM (ConvLSTM) [22] was first proposed to model spatiotemporal data, such as radar echo maps, whose spatial correlations cannot be extracted by fully connected layers in LSTM. ConvLSTM tailors the convolution operation into LSTM by replacing matrix multiplication operations in (1)–(5) with convolution, as follows:

$$i_t = \sigma(W_{xi} * \mathcal{X}_t + W_{hi} * \mathcal{H}_{t-1} + b_i), \quad (6)$$

$$f_t = \sigma(W_{xf} * \mathcal{X}_t + W_{hf} * \mathcal{H}_{t-1} + b_f), \quad (7)$$

$$C_t = f_t \circ C_{t-1} + i_t \circ \tanh(W_{xc} * \mathcal{X}_t + W_{hc} * \mathcal{H}_{t-1} + b_c) \quad (8)$$

$$o_t = \sigma(W_{xo} * \mathcal{X}_t + W_{ho} * \mathcal{H}_{t-1} + b_o), \quad (9)$$

$$\mathcal{H}_t = o_t \circ \tanh(C_t), \quad (10)$$

in which $*$ denotes the convolution operator, $\mathcal{X}_t \in \mathbb{R}^{d \times 1}$ is the input, and $W_{xi}, W_{hi}, W_{xf}, W_{hf}, W_{xo}, W_{ho} \in \mathbb{R}^{k \times c}$ are the convolution kernels. The above applies to a single ConvLSTM unit, which can be further extend to multiple layers as traditional LSTM does. A pooling layer can be added between each two ConvLSTM layers to reduce computation.

An immediate concern is whether applying convolution-embedded models on network traffic data without obvious spatial information would be effective. In fact, Convolutional Neural Networks (CNN) not only gained success in computer vision [23, 24], but also in areas including web traffic fingerprinting [25] and mobile traffic forecasting [26]. As the sequential data we deal with are one-dimensional, we implement a 1-D ConvLSTM, which takes inputs with channel and length dimensions, where channel by default equals 1.

Differences from CNN-LSTM. The CNN-LSTM, or Long-term Recurrent Convolutional Networks (LRCN), is a combination of CNN and LSTM, first proposed for visual recognition. It differs from ConvLSTM in that in the former a separate CNN handles spatial information before providing input to LSTM. In contrast, the latter has a compact form. While previous studies apply CNN-LSTM to NID [27], our work is the first to leverage ConvLSTM and study the differences between the two structures. Our results in §5 reveal that ConvLSTM consistently outperforms CNN-LSTM.

4.3. Bidirectional Asymmetric LSTM

Bi-LSTM and Bi-ConvLSTM can be perceived as ensemble models with two separate LSTM units. The hidden states from two units are concatenated so that future information is accessible at the current timestamp t , which can potentially benefit the downstream classification task. Since

our targets are sequences with *similar* malicious flows, it is the hidden states at the two ends of the structures that can acquire most information. The hidden representations in the middle of a sequence from two (Conv-)LSTM units yield certain amount of redundant information when the same architecture is used.

Given the fact that different architectures are likely to exploit different facets of input features for classification [29], we propose **Bidirectional Asymmetric LSTM (Bi-ALSTM)**, which consists of two different, asymmetric LSTM units, one for forward and the other for backward processing, to generate intermediate representations that incorporate more comprehensive temporal contexts. The hidden states from two different units are first linearly combined, then fed through an activation function. Precisely, denote $\mathbf{h}_{fc}^t \in \mathbb{R}^{N_1}$ as the hidden state generated by LSTM at timestamp t , and $\mathbf{h}_{conv}^t \in \mathbb{R}^{N_2}$ generated by ConvLSTM (operated backwards). The final hidden states are formed through the following fusion operation:

$$\mathbf{h}^t = \tanh\left(\frac{\mathbf{U}_{conv} \mathbf{h}_{conv}^t}{\|\mathbf{U}_{conv} \mathbf{h}_{conv}^t\|_2} + \frac{\mathbf{U}_{fc} \mathbf{h}_{fc}^t}{\|\mathbf{U}_{fc} \mathbf{h}_{fc}^t\|_2}\right), \quad (11)$$

where $\mathbf{U}_{conv} \in \mathbb{R}^{N \times N_2}$ and $\mathbf{U}_{fc} \in \mathbb{R}^{N \times N_1}$ are learnable parameters. \mathbf{h}_{fc}^t and \mathbf{h}_{conv}^t are projected to the same subspace and L_2 -normalized, so that in the learning process, any single unit would not easily dominate the final results. This design allows two asymmetric LSTMs to produce hidden states with different dimensions, meaning that two different LSTM structures can be tuned separately and flexibly before building the Bi-ALSTM. Finally, a single FC layer $g_\phi(\cdot)$:

Algorithm 2 The training algorithm for Bi-ALSTM

1: Inputs:

$D := \{S_1, \dots, S_n\}; S_i := \{(\mathbf{x}_{i1}, y_{i1}), \dots, (\mathbf{x}_{ia}, y_{ia})\},$
 $\lambda_1 := L_2 \text{ weight decay}, \lambda_2 := \text{learning rate}.$

2: Initialisation:

Denote $h_{W_{fc}}(\cdot)$ and $h_{W_{conv}}(\cdot)$ as the LSTM and ConvLSTM unit parameterized by W_{fc} and W_{conv} , and predictor $g_\phi(\cdot)$ parameterized by ϕ .
 W_{fc}, W_{conv}, ϕ set via Xavier initialization [28].

3: while model has not converged **do**

4: for S_i sampled from D **do**

$\bar{X}_i \leftarrow (\mathbf{x}_{i1}, \dots, \mathbf{x}_{ia}), \bar{y}_i \leftarrow (y_{i1}, \dots, y_{ia})$

$\bar{X}_i \leftarrow \text{reverse}(\bar{X}_i)$

$\bar{H}_{i,fc} \leftarrow h_{W_{fc}}(\bar{X}_i)$

$\bar{H}_{i,conv} \leftarrow \text{reverse}(h_{W_{conv}}(\bar{X}_i))$

$\bar{H}_i \leftarrow \bar{H}_{i,fc} \oplus \bar{H}_{i,conv} \quad \triangleright \text{Eq. 11}$

$\mathcal{L} \leftarrow NLL(\text{softmax}(g_\phi(\bar{H}_i), \bar{y}_i) +$

$\lambda_1(\|W_{conv}\|_2^2 + \|W_{fc}\|_2^2 + \|\phi\|_2^2)$

$\phi, W_{conv}, W_{fc} \leftarrow \text{Adam}(\mathcal{L}, \phi, W_{conv}, W_{fc}; \lambda_2)$

12: end for

13: end while

14: return W_{fc}, W_{conv}, ϕ

$\mathbb{R}^N \rightarrow \mathbb{R}^C$ with softmax function is used to approximate the probability of the samples belonging to each class.

Computational complexity: We first derive the time complexity of a single-step forward pass of LSTM. Given that the time complexity of $W_x x_t$ is $O(hd)$, it is easy to know that the time complexity inside each nonlinearity σ is $O(hd + h^2) = O(h(d + h))$. Assuming σ and \tanh have a constant time complexity, applying σ or \tanh element-wise yields the complexity of h , which can be omitted due to the existence of $O(h^2)$. Therefore, the time complexity of a single-step forward pass equals $O(h(d + h))$. Similarly, given that $W_h * \mathcal{X}_t$ has time complexity $O(dkc)$, the time complexity of a single-step forward pass of ConvLSTM is $O(dkc + dkc^2) = O(dkc^2)$. Assume the input sequence to Bi-ALSTM has the length n . The computational complexity of a single-layer Bi-ALSTM results in $O(n(h(d + h) + dkc^2))$.

We use negative log likelihood with L_2 regularization as the loss function. Bi-ALSTM is stochastically optimized via back propagation through time by the Adam algorithm. The complete training process follows Algorithm 2.

4.4. Why an All-range Multi-class Classifier is Unfeasible

Applying supervised methods for NID commonly involves training a multi-class classifier that seeks to detect as many types of malicious attacks as possible. However, we argue that this approach would lead to an overfitted model because the *ambiguity of the true attack labels* is widely overlooked. To understand this, consider a classifier is trained to differentiate two types of network attacks: SQLmap fuzzing vs. Web password Bruteforcing, both of which repetitively initiate HTTP requests to a given API. Although the two types of attacks would incorporate different payloads (SQL scripts and user/password pairs respectively), this discrepancy appears negligible on a statistical level, because the contents of payloads would not be extracted. Given that both trigger database I/O operation, the extracted timing information would be hardly differentiable.

Thus if a trained model can distinguish between such attacks, it has to be overfitted and learn a decision boundary that is unique to the dataset, rather than truly understand the differences. Attack labels usually indicate the purposes and techniques behind them, but when it comes to network layer, the attack realizations, i.e., traffic flows, would not be clearly dissimilar. In this regards, using sequential models to distinguish as many types of attacks as possible is unrealistic.

4.5. Abstract Labeling

Given that it is hard to correctly predict every type of automated cyber attack with a network-based algorithm, reverting to a binary detector seems sensible. However, since we augment DoS attacks, we decide to use **abstract labeling** in order to evaluate the augmentation technique and to avoid the aforementioned overfitting issue. Specifically, we assign a number of abstract, generic labels, including *benign*, *DoS*, *portscanning* and *bruteforcing & fuzzing*, as ground truth during training. On one hand, this approach can

clearly illustrate the influence of our feature augmentation technique. On the other hand, the model would not put effort in distinguishing the subtle differences between, e.g., DoS HOIC and DoS LOIC, which may not be separable by a network-based algorithm.

5. Experiments

We implement NetSentry in PyTorch and train the model on a GeForce Titan X GPU. To build the Bi-ALSTM, we use an LSTM unit for the forward pass and a ConvLSTM unit for the backward pass. The LSTM unit has the following structure: $\text{dropout}(0.5) \rightarrow \text{lstm}(65, 48) \rightarrow \text{lstm}(48, 48)$, where the arguments in $\text{lstm}(\cdot, \cdot)$ denote the input size and the hidden size. The ConvLSTM unit encompasses $\text{convlstm1D}(1, 3, 3) \rightarrow \text{convlstm1D}(3, 6, 3) \rightarrow \text{maxpool}()$, in which the arguments in $\text{convlstm1D}(\cdot, \cdot, \cdot)$ represent the input channel size, output channel size, and kernel size. The fused hidden states are passed through: $\text{dropout}(0.3) \rightarrow \text{MLP}(32, 5)$. The L_2 penalty and learning rate are set to 0.5 and 0.001 respectively.

For the Flow Aggregator/Sequence Generator, apart from our TCP termination fix (see §3.2.1), we set flow timeout to 30s and subflow duration 5s in CICFlowMeter. The Sequence Generator uses timeout $\tau = 30$ s and window size $\alpha = 10$.

5.1. Datasets

We experiment with two datasets published by the Canadian Institute for Cybersecurity (CIC), as described below.

CIC-IDS-2017 [30] contains most common cyber attacks, including bruteforcing, heartbleed, botnet, (D)DoS, Infiltration, and Web attacks. Traces were collected in a LAN, with benign traffic generated by profiling normal online behaviors of 25 users on different OSs, including Win Vista, Win 7, Win 8, Win 10, Mac OS, and Ubuntu 12. Attacks are produced by 4 different machines with one running Kali Linux and three Win 8.1. The traffic collection spans 5 working days, and in total 51.1 GB of pcap files are open-sourced. The feature sets of the corresponding pcap files were published, but as we reveal in Section 3.2.1, these were extracted incorrectly. Hence, we only use the raw capture files in our experiments.

CSE-CIC-IDS2018 [18] was generated in a much larger environment where an organizational LAN with five subnets simulated give different departments. 450 machines act as normal users and 50 as attackers. The dataset contains a wider range of benign traffic, including HTTPS, HTTP, SMTP, POP3, IMAP, SSH, and FTP, and contains a larger attack collection (17 types). The dataset spans 10 days.

Self-collected traffic – since *FTP-Bruteforce* and *DoS-SlowHTTPTest* attacks were erroneously collected in CSE-CIC-IDS2018, the generated traffic merely contains SYN and RST packets. To mitigate this, we collected FTP-bruteforcing traffic ourselves, generating 4,050 flows. We decide not to collect *DoS-SlowHTTPTest* traffic, since a similar type of attack, i.e., *DoS-Slowloris*, exists in CSE-CIC-IDS2018. We remove the mislabeled traffic and merge the

| | # Features | # stances | In- Anomaly ratio | Automated attack ratio |
|----------|---------------|--------------|-------------------------|------------------------------|
| IDS-2017 | 69 | 2,607,289 | 0.2 | 0.189 |
| IDS-2018 | 69 | 8,786,169 | 0.1806 | 0.1802 |

Table 2

Statistics of datasets used for experimentation. ID-based features in Table 1 not included in the training; ‘protocol’ is hot-encoded. Thus, only 65 features used in total.

self-collected traffic with the CSE-CIC-IDS2018 dataset. In the rest of our paper, we use CSE-CIC-IDS2018 to refer to the merged dataset. We employ our revised version of CICFlowMeter to generate network flows based on the pcap files. The statistics of both datasets are shown in Table 2.

5.2. Cross-Evaluation

We adopt a rigorous evaluation methodology, aiming to show the true generalization ability of our design, by **cross-evaluation**. Normally, a dataset is split into training and testing subsets, and the results on the test set compared across different algorithms. However, network environments are heterogeneous and data collected in one environment may not accurately reflect the diversity seen in practice. To test if an algorithm can truly distinguish the same type of malicious traffic in a different network topology, we also evaluate it on a second, unseen dataset. CSE-CIC-IDS2018 is split into training (70%) and test (30%) sets. To maintain time consistency, training data is selected from events that took place before the test data. CIC-IDS-2017 is purely used for cross-evaluation, after the model was trained on the former.

5.3. Benchmarks

For comparison, we implement a range of benchmarks, including basic ML/DL structures (MLP, CNN, autoencoder, RIPPER [31], Decision Tree); state-of-the-art anomaly/intrusion detectors, i.e., One Class Neural Nets (OC-NN) [32], Kitsune/KitNET [7] and Deep Autoencoding Gaussian Mixture Model (DAGMM) [33]; and three Bi-LSTM variants [27].

OC-NN [32] and DAGMM [33] are offline semi-supervised algorithms for general anomaly detection. OC-NN aims to learn a mapping for the benign samples to a kernel space where the majority of them can be enclosed by a hypersphere. During the testing phase, the distances from the samples to the center of the hypersphere represent the anomaly score of the data. Different from OC-NN, DAGMM models the benign data from a probabilistic perspective with a mixture of Gaussian distributions. The negative probability of the data being sampled from the Probabilistic Density Function (PDF) represents the anomaly score.

Kitsune [7] is an online semi-supervised NIDS. It uses an ensemble of shallow autoencoders to learn the features of benign data in different subspaces; a final autoencoder fits the correlations of the reconstruction errors from the shallow autoencoders. The neural architecture is named KitNET. During testing, the reconstruction errors are computed to

represent the degree of abnormality. For a fair comparison, KitNET is trained in an offline manner with more than one epoch.

For semi-supervised algorithms (OC-NN, DAGMM, KitNET and Autoencoder), an anomaly ratio α needs to be preset, indicating the proportion of anomalous samples, and during the testing phase, the data with the top $\alpha \times 100\%$ of the anomaly scores are classified as anomalous. The anomaly ratio is set to 0.189 and 0.1802 on CIC-IDS-2017 and CSE-CIC-IDS2018 respectively, which is the same percentage of automated attacks in the two datasets.

RIPPER [31] and Decision Tree are two basic machine learning models, where the first one aims to generate a simple ruleset for classifications while the second embeds rules in a tree by recursively finding the best splits.

The structures of Bi-LSTM and Bi-ConvLSTM resemble the units in Bi-ALSTM. For CNN-Bi-LSTM, an extra CNN block, with the structure: $Conv1D(1, 3, 3) \rightarrow MaxPooling \rightarrow Conv1D(3, 6, 3) \rightarrow MaxPooling$, is implemented. The arguments in $Conv1D(\cdot, \cdot, \cdot)$ represent input channels, output channels, and kernel sizes.

Prior to testing, we retrain all benchmarks with all the features in the CSE-CIC-IDS2018 dataset, which is richer than the datasets used for training in the original papers.

5.4. Evaluation Metrics

The average precision, recall and F1 score are commonly used to evaluate the performance of anomaly detection algorithms. These metrics can be measured based on the True Positives (TP), False Positives (FP), True Negatives (TN) and recall are computed as $precision = TP / (TP + FP)$, $recall = TP / (TP + FN)$. The precision indicates how likely the algorithm would give true alarms, and the recall measures how sensitive the algorithm is towards anomalies. There exists a trade-off between precision and recall, and to obtain an overall performance measure, their harmonic average is computed, i.e., the F1 score: $F1 = 2 \times \frac{precision \times recall}{precision + recall}$.

We do not measure accuracy i.e., the percentage of the correctly classified samples, which is unlikely to reveal the algorithms’ true NID performance: consider a dataset with 80% benign and 20% malicious instances; a model that classifies everything as benign has the same accuracy as a model that correctly recognizes all but 20% of the benign traffic.

For the ML-based algorithms that output an anomaly score for each test instance rather than just the predicted class, system administrators may choose a threshold higher than 0.5, which guarantees that the classifier has a lower False Positive Rate (FPR). We plot the Receiver Operating Characteristic (ROC) curve for sequential models, to evaluate their performance when the anomaly threshold is varied in $[0, 1]$. The ROC curve is obtained by plotting the False Positive Rate (TPR) against FPR. The closer to 1 the Area Under Curve (AUC) is, the better the classifier performs.

Given that our datasets consist of multiple types of cyber attacks, we further plot the Empirical Cumulative Distribution Function (ECDF) of the anomaly score with respect

| Algorithm | CSE-CIC-IDS2018 | | | CIC-IDS-2017 (X-eval) | | |
|--------------------------|-----------------|---------------|---------------|-----------------------|---------------|---------------|
| | precision | recall | F1 | precision | recall | F1 |
| RIPPER | 0.9983 | 0.0981 | 0.1786 | 0.0873 | 0.0106 | 0.0190 |
| Decision Tree | 0.9989 | 0.9990 | 0.9990 | 0.5385 | 0.3717 | 0.4398 |
| MLP | 0.9989 | 0.9962 | 0.9976 | 0.6736 | 0.4631 | 0.5435 |
| CNN | 0.9947 | 0.9951 | 0.9949 | 0.7705 | 0.6344 | 0.6958 |
| Autoencoder | 0.7783 | 0.7500 | 0.7639 | 0.4362 | 0.4197 | 0.4278 |
| OC-NN | 0.9722 | 0.5310 | 0.6868 | 0.7844 | 0.5136 | 0.6208 |
| Kitsune | 0.6310 | 0.6081 | 0.6193 | 0.4086 | 0.3932 | 0.4007 |
| DAGMM | 0.8666 | 0.8253 | 0.8454 | 0.4159 | 0.3116 | 0.3576 |
| Bi-LSTM | 0.9990 | 0.9979 | 0.9985 | 0.7258 | 0.4209 | 0.5317 |
| CNN-Bi-LSTM | 0.9996 | 0.9982 | 0.9989 | 0.8813 | 0.3750 | 0.5261 |
| Bi-ConvLSTM | 0.9984 | 0.9971 | 0.9977 | 0.8721 | 0.9693 | 0.9178 |
| Bi-ALSTM | 0.9994 | 0.9990 | 0.9992 | 0.9116 | 0.9446 | 0.9275 |
| Algorithm + augmentation | CSE-CIC-IDS2018 | | | CIC-IDS-2017 (X-eval) | | |
| | precision | recall | F1 | precision | recall | F1 |
| RIPPER | 0.9980 | 0.0934 | 0.1709 | 0.4998 | 0.1837 | 0.3687 |
| Decision Tree | 0.9989 | 0.9993 | 0.9991 | 0.5897 | 0.8556 | 0.6914 |
| MLP | 0.9989 | 0.9963 | 0.9976 | 0.7540 | 0.8690 | 0.8071 |
| CNN | 0.9925 | 0.9847 | 0.9886 | 0.7453 | 0.8687 | 0.8021 |
| Bi-LSTM | 0.9991 | 0.9956 | 0.9973 | 0.8555 | 0.9777 | 0.9125 |
| CNN-Bi-LSTM | 0.9996 | 0.9966 | 0.9981 | 0.8479 | 0.9683 | 0.9041 |
| Bi-ConvLSTM | 0.9996 | 0.9975 | 0.9985 | 0.8728 | 0.9780 | 0.9222 |
| Bi-ALSTM | 0.9997 | 0.9976 | 0.9987 | 0.9190 | 0.9800 | 0.9485 |

Table 3

Precision, recall, and F1 score for Bi-ALSTM and benchmarks. NB: only supervised algorithms can be trained with augmented data and only HTTP (D)DoS attacks are augmented; semi-supervised methods only require benign traffic for training.

to each type of traffic on the crossed evaluated dataset, to illustrate the confidence of each sequential model.

Beyond the metrics for classification performance, we also care about the computational overhead of our design, and therefore report Multiply-Accumulate Operation Counts (MACs), the number of parameters and the concurrent processing capacity of each model on a edge GPU. MACs and the parameter numbers reveal the complexity of different algorithms at a micro level, while the concurrent processing capacity on GPU can reflect the computational bottlenecks.

5.5. Performance without Augmented Data

We summarize our comparison in terms of threat detection performance between our Bi-ConvLSTM/-ALSTM models and the benchmarks considered, in Table 3. In the upper half, the different algorithms are trained on non-augmented data. The performance of the benchmark algorithms and those adopted by our NetSentry is similar on the CSE-CIC-IDS2018 dataset, most of them attaining average metrics above 0.99. CNN-Bi-LSTM slightly outperforms other algorithms in terms of precision, while Bi-ALSTM yields the highest recall and F1 score. An interesting finding is that the semi-supervised algorithms for general anomaly detection may not be suitable for network intrusion detection. Autoencoder, OC-NN and DAGMM cannot compete with basic supervised ML algorithms. One of the core assumptions for semi-supervised anomaly detection is that the algorithm can learn the characteristics of benign data, by

estimating the probability, reconstructing the benign samples or finding an appropriate hyper-boundary enclosing them. However, network traffic is highly heterogeneous, serving with different protocols various applications, such as email, web browsing, streaming, etc. It remains questionable whether the aforementioned assumption holds on such a large range of ‘benign data’. Besides, detecting malicious traffic, especially automated attacks (which is our objective), appears to be a time-sensitive task. Therefore, observing a single instance may be insufficient to make reliable decisions. Consequently, existing anomaly detection algorithms tend to ignore this, which leads to modest results.

The advantage of Bi-ConvLSTM and Bi-ALSTM can be clearly seen on cross-evaluation results, where our models maintain consistently competitive performance. Both attain

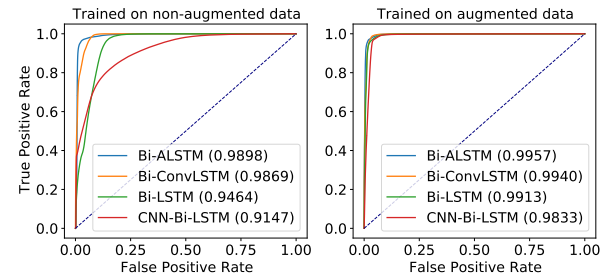


Figure 6: ROC curves of LSTM-based algorithms. AUC behind labels. Models trained w/o (left) and w/ augmented data (right).




Figure 7: Normalized confusion matrices (row values add to 1) for LSTM-based models cross-evaluated on CIC-IDS-2017. Models trained with/without augmented dataset (left/right). B represents Benign, D DoS, BF Bruteforcing and Fuzzing, and PS PortScanning. Numbers on diagonals are recalls.

F1 scores above 90%, while other supervised algorithms, including Bi-LSTM and CNN-Bi-LSTM exhibit a significant performance drop (F1 score around 50%). We notice that though both CNN-LSTM and ConvLSTM are proposed to handle spatiotemporal data, there is an obvious difference in performances, both in terms of F1 score (Table 3) and ROC (Figure 6) on intrusion detection. As shown in Figs 7 (a), (b), Bi-LSTM and CNN-Bi-LSTM do not learn a reliable decision boundary between benign traffic and DoS attacks without the augmented data. Bi-ConvLSTM clearly outperforms them, yet still exhibits a high probability of classifying DoS as port scanning attacks, as illustrated in Figure 7(c), whereas Bi-ALSTM is the most reliable (Figure 7(d)).

5.6. Performance with Augmented Data

The data augmentation procedure we propose is highly effective in helping the models generalize well. *When the supervised models are trained with the augmented dataset, a remarkable performance gain can be observed in the cross-evaluation results* (Table 3, bottom half). The F1 scores of the benchmarks increase by at least 16%, and Bi-LSTM and CNN-Bi-LSTM even jump to 90%. The improvements of Bi-ConvLSTM and Bi-ALSTM are less noticeable since outstanding results can be achieved even without augmented data, but still, the former reaches the highest recall and Bi-ALSTM is the most robust in terms of overall performance.

Observing confusion matrices in the second column in Figure 7, all models reveal roughly the same pattern, as opposed to the corresponding results on the first. This confirms that *augmentation encourages models to learn associating timing info rather than payload features* in the classification task.

5.7. Impact of Feature Arrangement

We investigate the influence of the feature arrangement and kernel size on the performance of ConvLSTM. For this, we experiment with 3 different kernel sizes, namely (3, 5, 7), and two sets of feature arrangements. Specifically, the 1D feature vectors are logically ordered and randomly shuffled. Note that most of the features listed in Table 1 in the Appendix are computed for forward traffic only, backward traffic only, and bidirectionally. Logically ordered features means that they follow an alternating order of forward, backward, and bi-direction. In each experiment, we train both unidirectional ConvLSTM and Bi-ConvLSTM with the augmented dataset for 10 epochs and repeat the process 5 times. The mean and error bars of the F1 score on the cross-evaluation dataset (CIC-IDS-2017) are illustrated in Figure 8.

Intuitively, one might expect ConvLSTM would only work with sequential data possessing clear spatial information, such as videos. However, we find that ConvLSTM is robust to 1D traffic features regardless of their arrangement. Indeed, the results in Figure 8 demonstrate that there is no significant gap between the model trained with logically ordered features or randomly shuffled ones. For most cases, the mean in the former case is slightly higher than in the latter, while the error bars show a large degree of overlap.

We also find that although F1 scores tend to rise slightly with larger kernel sizes when the ConvLSTM is trained with the ordered feature arrangement, this increase is not significant. Considering the growth in computation overhead with using a larger kernel size, we argue that training both Bi-ConvLSTM and Bi-ALSTM with a kernel size equal to 3 (which was also the case for the results presented in Table 3) is sufficient.

5.8. Performance Gains of Bi-ALSTM

Bi-ALSTM not only yields the highest overall detection rate (recall), but also reliably detects each type of cyber attacks, as illustrated in Figure 9. Both Bi-LSTM and Bi-ConvLSTM have difficulty recognizing web bruteforcing,

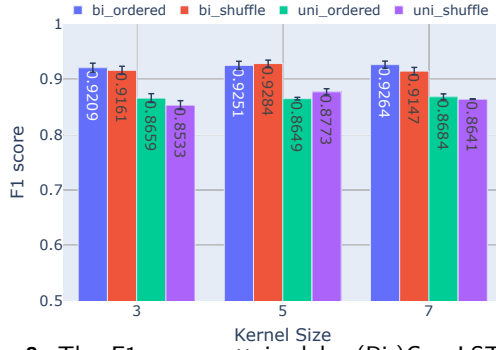


Figure 8: The F1 scores attained by (Bi-)ConvLSTM with different kernel sizes and different order of features.

XSS, and Slowloris attacks, whereas *Bi-ALSTM* attains up to $3\times$ higher detection rates. The only exception is SQL injection, which cannot be detected by all the algorithms. This is because there are only 53 instances of this attack, merely accounting for 0.0006% of the entire dataset, which is insufficient for the model to learn a reliable decision boundary for classification.




Figure 9: Detection rate (recall) of each type of traffic evaluated on CSE-CIC-IDS2018 (top) and CIC-IDS-2017 (bottom).




Figure 10: ECDF of the anomaly scores with respect to each type of traffic in CIC-IDS-2017 given by Bi-LSTM, Bi-ConvLSTM and Bi-ALSTM. All models trained with augmented dataset.

We evaluate the quality of the anomaly scores approximated by Bi-(Conv)LSTM and Bi-ALSTM. The anomaly score is the value output by the model. Since the activation function of the last layer is softmax, the output is squeezed between $[0, 1]$ and the higher the value, the more anomalous a flow is regarded. System administrators routinely customize an anomaly threshold to lower the FPR. Figure

| Model | MACs(k) | Parameters(k) | Edge GPU (Mflow/s) |
|-------------|---------|---------------|--------------------|
| MLP | 5.7 | 5.8 | 41.4 |
| CNN | 3.2 | 1.5 | 73.7 |
| Autoencoder | 10.3 | 10.6 | 22.9 |
| OC-NN | 5.2 | 5.2 | 45.4 |
| Kitsune | 0.7 | 0.8 | 337.1 |
| DAGMM | 5.3 | 5.4 | 44.5 |
| Bi-LSTM | 102.2 | 100.7 | 2.3 |
| CNN-Bi-LSTM | 116.8 | 112.7 | 2.0 |
| Bi-ConvLSTM | 51.8 | 2.7 | 4.5 |
| Bi-ALSTM | 66.8 | 41.4 | 3.5 |

Table 4

Computation overhead in terms of MACs per inference instance and number of parameters for each model. The last column presents the number of flows (in millions) an edge GPU (Jetson Nano) can process per second.

10 plots the ECDF of each type of traffic in CIC-IDS-2017 given by the three algorithms, in which the blue line corresponds to benign traffic. The black dashed line is the threshold that sets the FPR to 1.5%, and the area under the other lines to the left of the threshold line represents the proportion of attacks that would be misclassified. We find that Bi-ALSTM delivers the lowest False Negative Rate (FNR) (2.63%) compared with Bi-LSTM (10.17%) and Bi-ConvLSTM (5.87%).

5.9. Computational Overhead

While NetSentry is primarily designed as an offline NIDS, employing it for online NID is also feasible. Table 4 details the MACs the benchmark models and our Bi-ConvLSTM/ALSTM structures require for a single traffic flow inference, as well as their number of parameters. (CNN-)Bi-LSTM are the most computationally expensive, given that multiple fully-connected layers are embedded in the LSTM unit. In contrast, *Bi-ConvLSTM/ALSTM* are *relatively lightweight, both involving fewer computations and parameters*. Deploying NetSentry as an online system next to routers or organizational gateways equipped with a GPU or TPU should thus be straightforward.

Given that edge AI platforms are now available, e.g., Nvidia Jetson Nano [34], running NetSentry on constrained small-business/ home routers is within reach. Results in Table 4 reveal that Bi-ConvLSTM/ALSTM can handle 4.5/3.5 Mflows per second, which confirms our practicality assessment.

Another key merit of NetSentry is that the system inherently analyses consecutive traffic between pairs of hosts, which is easy to integrate into an Intrusion Prevention System (IPS), without the need for collecting statistics of potentially malicious hosts until reaching full confidence about decisions to enforce. Recall that our system directly gives prediction results about the traffic flows generated between two hosts during a short interval, offering comprehensive contexts to the IPS with low FP risks. Dynamic firewall rules

can also be updated effortlessly, since the atomic processing input of NetSentry originates from the same pair of hosts.

6. Discussion

Lastly, we discuss the robustness of our system against different evasion attacks.

IP Spoofing: IP addresses can be spoofed with little effort, which is also a common approach to generating DDoS attacks. Flow-based NIDS may be ineffective in preventing such traffic because ‘identities’ are changed frequently. However, it is worth noting that IP spoofing can only be used to initiate stateless DDoS, given that any responses from the victim is not guaranteed to be routed back to the attacker. Existing countermeasures such as TCP half-open and ICMP threshold are capable of mitigating those issues. For application-layer DDoS, attackers must control the real IP addresses to maintain the states, where NetSentry will not be fooled.

Traffic Encryption: Traffic Encryption was proposed for evasion attacks [35], whereby malicious payload is hidden in an encrypted channel. This is however only effective against NIDS that examine the syntax of network communications, such as BotHunter [10]. NetSentry is designed to extract and analyze timing- and protocol-based statistics. That said, manipulation of payload contents cannot bypass our design.

Adversarial Perturbations: Adding small perturbations to input data may lead to misclassification by ML models [36]. Nevertheless, the existing adversarial attacks often require access to model gradients, structures, or numerous queries for weight approximation. In reality, a ML-based NIDS would not disclose the details of its neural model and tolerate countless queries. Zhang et al. demonstrate the possibility of attacking ML-based intrusion detection algorithms by heuristic-based methods without knowing model’s information [37], but it would still take 100~11,000 queries to generate one adversarial sample. Note that NetSentry is intended for continuous and repetitive network attacks, meaning that similar queries would trigger alarms much earlier than discovering a valid adversarial sample.

Adversarial perturbations are likely to modify every individual feature to create malicious samples, which is not always practical in the networking domain, since the modified flows are not guaranteed to stem from any real traffic. Consider instead a more pragmatic attack scenario where the adversary slows the attack speed by increasing the time between the packets sent. To evaluate the potential impact of this adaptive attack on NetSentry, we first pre-process both CSE-CIC-IDS2018 and CIC-IDS-2017 datasets as follows: (1) we group the packets belonging to malicious activity in the pcap traces into flows; (2) in each flow, we alter the timestamps of the attacker’s packets by expanding the time gap between the previously received/sent packet and the current one, by a fixed multiplier $m \in \{1, 2, 4, 8\}$ (packets are not delayed if $m = 1$); and (3) we alter the timestamps of the victim’s packets to ensure the time gaps between these

| | m | Test (CSE-CIC-IDS-2018) | | | |
|---------------------|---|---------------------------|--------|--------|--------|
| | | 1 | 2 | 4 | 8 |
| Train (IDS-2018) | 1 | 0 | -0.08% | -0.02% | -0.02% |
| | 2 | -0.32% | 0 | -0.16% | -0.24% |
| | 4 | -0.35% | 0 | 0 | -0.01% |
| | 8 | -0.09% | +0.2% | 0 | 0 |
| | m | Cross Test (CIC-IDS-2017) | | | |
| | | 1 | 2 | 4 | 8 |
| Train (IDS-2018) | 1 | 0 | -0.05% | -0.11% | -0.58% |
| | 2 | +0.1% | 0 | +0.2% | -0.39% |
| | 4 | 0 | -0.01% | 0 | -0.35% |
| | 8 | -0.49% | -0.56% | -0.11% | 0 |

Table 5

Percentage Error wrt. F1 score. Attacker’s packets slowed down by factors $\{1, 2, 4, 8\}$. $m = 1$ for original timing.

and the attacker’s packets still match those in the original flows.

PortScan attacks are excluded from both datasets because the majority only consists of 1–2 packets, and applying the logic above will not change their timestamps at all. We choose a set of multipliers $m \in \{1, 2, 4, 8\}$, where the attacker’s packets are not delayed if $m = 1$. As such, We obtain three altered versions of the CSE-CIC-IDS2018 and CIC-IDS-2017 dataset. Each variant of CSE-CIC-IDS2018 is split into a training set (70% of samples) and a test set (30%), and we augment all the training sets as detailed in §3.2.3, then retrain the Bi-ALSTM. The altered CIC-IDS-2017 datasets are used for cross evaluation. We measure the Percentage Error (PE) with respect to the F1 score, to understand to what extent the model would degrade when facing malicious traffic that is purposely slowed down by different factors, to attempt evasion. Formally, PE wrt. F1 score is defined as:

$$PE_{i,j}^{F1} = \frac{F1_{i,j} - F1_{i,i}}{F1_{i,i}} \times 100\%,$$

where the first subscript denotes the multiplier $m = i$ applied in the dataset used for model training, and the second subscript to slow-down factor in the set used for testing. As shown in Table 5, we find that the maximum PE on the CSE-CIC-IDS2018 is never above 0.35% and the maximum PE on CIC-IDS-2017 is below 0.58%. This demonstrates that *manipulating the attack timing has no effective impact on the detection performance of the proposed NetSentry*.

7. Related Work

Network intrusion detection has been the focus of extensive research in the security community. In what follows, we briefly discuss the most relevant work related to ours, highlighting limitations of prior approaches and similarities with the proposed NetSentry, where appropriate.

Defenses through Offensive Footprint Profiling. Modeling the unique malicious nature of network anomalies is effective for detection. BotHunter [10] builds infection dialogues to describe the dynamic process of Botnet infection, then employs modularized detection engines to identify the footprint of each stage in an attack. BotSniffer [38]

identifies bot activity by highlighting the spatiotemporal correlations of Command and Control (C&C) traffic originating from pre-programmed behaviors. Profiling malicious code execution paths plays an important role in detecting malware [39, 40]. Likewise, stealth DDoS amplification can be fingerprinted by its unique two-stage behavior (i.e., scan and attack) [41]. These contributions demonstrate that modeling the potential links between different attack phases has merit in practice. Unlike previous works, here we reveal how *different stages are common between different large-scale attacks* and why breaking their sequence is essential to thwarting intrusions.

Time-invariant ML for NID. DL-based NIDSs learn illicit traffic patterns through a spectrum of algorithms, replacing the explicit attack modeling methodology introduced previously. In detecting anomalies, such algorithms largely performing analysis on a per-sample basis, i.e., using statistical features of a traffic flow, to determine its nature, rather than exploring any potential correlations in network traffic. *Supervised Learning* approaches, including RIPPER [31], Support Vector Machine (SVM) [42], and Random Forest [43], treat anomaly detection as a classification problem, seeking a decision boundary between benign and malicious traffic. *Semi-supervised Learning* methods discard anomalous samples during training, and only learn patterns of benign traffic. Kitsune [7] learn to reconstruct benign data via encoder (\mathcal{E}) and decoder (\mathcal{D}) networks. Samples with high reconstruction errors, i.e., $\|x - \mathcal{D}(\mathcal{E}(x))\|_2$, are deemed as malicious. One-Class Deep SVDD [32] believes benign samples can be enclosed by a hyper-sphere, whereas anomalous ones are distinct from the center. Thus, Deep SVDD learns a non-linear transformation that maps innocuous samples into a feature space where the majority of them can be surrounded by a small hyper-sphere. Statistical approaches assume that benign data in nature are densely distributed in the feature space, while anomalies (outliers) are scattered. Dense areas can be approximated by Deep Gaussian Mixture models [33] or Generative Adversarial Networks [44].

Time-sensitive ML for NID relies on temporal context along with a sample, to detect any intrusion. NIDS that employ this approach are sparse. More commonly, it is Host-based Intrusion Detection Systems (HIDS) [45, 46, 47, 48] that utilize time-sensitive models, such as LSTM and RNN, because the target data (system calls, logs and security events) present obvious semantic meaning and potential temporal dependencies. Attention-based Graph Neural Networks [49] can also be used to model high-dimensional time-series data and spot anomalies. Alternatively, USAD [50] handles time-sensitive tasks by segmenting time-series data into fixed-size windows, and uses adversely trained autoencoders to detect intrusions or anomalies. Recent studies attempt to model temporal correlations within network attacks and propose a range of RNN-based algorithms [51, 52]. However, an appropriate threat model detailing what temporal information is relevant to NID is missing. Moreover, the training inputs are often randomly sampled, which

suppresses relevant temporal information and makes NID effectiveness questionable. Our NetSentry design sets to address this particular issue and *takes a dynamic view to cyber attacks, so as to identify possible temporal relationships that exist among different types of attacks*, thereby building a well-directed defensive approach.

8. Future Work

As we discuss in Section 6, existing adversarial attacks on NIDS add perturbations to the statistical features of traffic flows. There is no guarantee that perturbed features can be mapped back to a sequence of packets to be transmitted in practice. It remains unclear whether conducting adversarial attacks by directly shaping consecutive packet sizes and inter-arrival times can deceive ML-based NIDS. We deem this topic as important, since it could further shed light on the robustness and reliability of our method.

On the other hand, Bi-ALSTM is a supervised algorithm that demands a significant amount of data for training, but acquiring up-to-date datasets is not always feasible, given the stealth nature of cyber attacks. Unfortunately, existing semi-supervised algorithms still focus on per-flow classification and neglect temporal context, resulting in the undesirable performance seen in Table 3 (Autoencoder, OC-NN, Kitsune, DAGMM). Instead of approximating the distribution of benign flows, estimating the stochastic process of consecutive benign traffic may provide higher reliability, which is also an interesting topic for future study.

9. Conclusions

In this paper, we show that large-scale network threats with potential high-impact can be tackled in their early stages, if correctly recognizing the unique temporal dependencies of malicious flows, and we propose NetSentry to effectively detect such incipient attacks. NetSentry incorporates a novel data augmentation technique to enhance the generalization ability of supervised algorithms and we design an ensemble Bi-ALSTM as the core intrusion detection logic. Extensive results demonstrate that our ensemble structure outperforms a wide range of benchmarks, attaining up to $3\times$ higher detection rates, under different network environments. Finally, we discuss computation overhead and robustness to evasion attacks, making the case for the feasibility of deploying NetSentry alongside threat prevention logic in real-world settings.

Acknowledgments

This material is based upon work supported by Arm Ltd and Scotland's Innovation Centre for sensing, imaging and Internet of Things technologies (CENSIS).

References

- [1] Spamhaus Botnet Threat Update: Q1-2021, <https://www.spamhaus.org/news/article/809/spamhaus-botnet-threat-update-q1-2021>.
- [2] C. Ventures, Global Cybercrime Damages Predicted To Reach \$6 Trillion Annually By 2021, <https://cybersecurityventures.com/cybercrime-damages-6-trillion-by-2021/> (2020).
- [3] L. Bilge, T. Dumitras, Before we knew it: An empirical study of zero-day attacks in the real world, in: Proc. ACM CCS, 2012.
- [4] A. L. Buczak, E. Guven, A survey of data mining and machine learning methods for cyber security intrusion detection, IEEE Communications Surveys & Tutorials 18 (2) (2016) 1153–1176.
- [5] Z. Lin, et al., IDSGAN: Generative adversarial networks for attack generation against intrusion detection, arXiv:1809.02077 (2019).
- [6] Y. Xia, et al., Learning discriminative reconstructions for unsupervised outlier removal, in: Proc. IEEE ICCV, 2015.
- [7] Y. Mirsky, et al., Kitsune: an ensemble of autoencoders for online network intrusion detection, in: NDSS, 2018.
- [8] R. Sommer, V. Paxson, Outside the closed world: On using machine learning for network intrusion detection, in: IEEE S&P, 2010.
- [9] Cyberint Research, British Airways Flight to DDoS Lands with Cyber Turbulence, <https://blog.cyberint.com/british-airways-flight-to-ddos-lands-with-cyber-turbulence> (March 2020).
- [10] G. Gu, et al., Bothunter: Detecting malware infection through ids-driven dialog correlation, in: USENIX Security, 2007.
- [11] L. Iffländer, et al., Hands off my database: Ransomware detection in databases through dynamic analysis of query sequences, arXiv preprint arXiv:1907.06775 (2019).
- [12] M. Antonakakis, et al., Understanding the Mirai botnet, in: USENIX Security, 2017.
- [13] E. M. Hutchins, et al., Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains, Leading Issues in Information Warfare & Security Research (2011).
- [14] E. Saad, et al., OWASP Web Security Testing Guide Version 4.1 (2019).
- [15] C. C. Vulnerabilities, Exposures, CVE-2017-0144, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2017-0144> (Sept. 2016).
- [16] D.-Y. Kao, S.-C. Hsiao, The dynamic analysis of wannacry ransomware, in: Proc. ICACT, 2018.
- [17] C. I. for Cybersecurity, CICFLOWMETER, <https://www.unb.ca/cic/research/applications.html#CICFlowMeter> (2018).
- [18] T. C. S. Establishment, C. I. for Cybersecurity, A Realistic Cyber Defense Dataset (CSE-CIC-IDS2018), <https://registry.opendata.aws/cse-cic-ids2018/> (2018).
- [19] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural computation 9 (8) (1997) 1735–1780.
- [20] I. Sutskever, O. Vinyals, Q. V. Le, Sequence to sequence learning with neural networks, Advances in neural information processing systems 27 (2014).
- [21] S. Kumar, D. Kumar, P. K. Donta, T. Amgoth, Land subsidence prediction using recurrent neural networks, Stochastic Environmental Research and Risk Assessment 36 (2) (2022) 373–388.
- [22] S. Xingjian, et al., Convolutional LSTM network: A machine learning approach for precipitation nowcasting, in: NeurIPS, 2015.
- [23] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [24] J. Yang, S. Xiao, A. Li, W. Lu, X. Gao, Y. Li, Msta-net: Forgery detection by generating manipulation trace based on multi-scale self-attention, IEEE Transactions on Circuits and Systems for Video Technology (2021).
- [25] P. Sirinam, et al., Deep fingerprinting: Undermining website fingerprinting defenses with deep learning, in: Proc. ACM CCS, 2018.
- [26] C. Zhang, et al., Multi-service mobile traffic forecasting via convolutional long short-term memories, in: Proc. IEEE International Symposium on Measurements & Networking, 2019.
- [27] K. Jiang, et al., Network intrusion detection combined hybrid sampling with deep hierarchical network, IEEE Access 8 (2020).
- [28] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: Proc. AISTATS, 2010.
- [29] M. Raghu, T. Unterthiner, S. Kornblith, C. Zhang, A. Dosovitskiy, Do vision transformers see like convolutional neural networks?, Advances in Neural Information Processing Systems 34 (2021).
- [30] I. Sharafaldin, et al., Toward generating a new intrusion detection dataset and intrusion traffic characterization, in: ICISSP, 2018.
- [31] W. Lee, S. Stolfo, Data mining approaches for intrusion detection, in: USENIX Security Symposium, 1998.
- [32] L. Ruff, et al., Deep one-class classification, in: International conference on machine learning, 2018, pp. 4393–4402.
- [33] B. Zong, Q. Song, M. R. Min, W. Cheng, C. Lumezanu, D. Cho, H. Chen, Deep autoencoding Gaussian mixture model for unsupervised anomaly detection, in: ICLR, 2018.
- [34] NVIDIA, Jetson Nano | NVIDIA Developer, <https://developer.nvidia.com/embedded/jetson-nano> (2019).
- [35] E. Stinson, J. C. Mitchell, Towards systematic evaluation of the evadability of bot/botnet detection methods, WOOT 8 (2008).
- [36] K. T. Co, et al., Procedural noise adversarial examples for black-box attacks on deep convolutional networks, in: ACM CCS, 2019.
- [37] C. Zhang, et al., Tiki-taka: Attacking and defending deep learning-based intrusion detection systems, in: ACM CCSW, 2020.
- [38] G. Gu, et al., Botsniffer: Detecting botnet command and control channels in network traffic, in: Network and Distributed System Security Symposium (NDSS), 2008.
- [39] C. Kolbitsch, et al., Effective and efficient malware detection at the end host, in: USENIX Security Symposium, Vol. 4, 2009, pp. 351–366.
- [40] A. Naderi-Afooshteh, et al., Malmax: Multi-aspect execution for automated dynamic web server malware analysis, in: ACM CCS, 2019.
- [41] J. Krupp, et al., Identifying the scan and attack infrastructures behind amplification ddos attacks, in: Proc. ACM CCS, 2016.
- [42] Y. Yi, et al., Incremental SVM Based on Reserved Set for Network Intrusion Detection, Expert Systems with Applications 38 (6) (2011) 7698–7707.
- [43] P. Sangkatsanee, et al., Practical real-time intrusion detection using machine learning approaches, Computer Communications 34 (18) (2011) 2227–2235.
- [44] D. Li, D. Chen, B. Jin, L. Shi, J. Goh, S.-K. Ng, Mad-gan: Multivariate anomaly detection for time series data with generative adversarial networks, in: International Conference on Artificial Neural Networks, Springer, 2019, pp. 703–716.
- [45] M. Du, et al., Lifelong anomaly detection through unlearning, in: Proc. ACM CCS, 2019.
- [46] Y. Shen, et al., Tiresias: Predicting security events through deep learning, in: Proc. ACM CCS, 2018.
- [47] Y. Su, Y. Zhao, C. Niu, R. Liu, W. Sun, D. Pei, Robust anomaly detection for multivariate time series through stochastic recurrent neural network, in: Proc. ACM SIGKDD international conference on knowledge discovery & data mining, 2019, pp. 2828–2837.
- [48] Y. Liu, S. Garg, J. Nie, Y. Zhang, Z. Xiong, J. Kang, M. S. Hossain, Deep anomaly detection for time-series data in industrial iot: A communication-efficient on-device federated learning approach, IEEE Internet of Things Journal 8 (8) (2020) 6348–6358.
- [49] A. Deng, B. Hooi, Graph neural network-based anomaly detection in multivariate time series, in: AAAI, Vol. 35, 2021, pp. 4027–4035.
- [50] J. Audibert, P. Michiardi, F. Guyard, S. Marti, M. A. Zuluaga, USAD: unsupervised anomaly detection on multivariate time series, in: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2020, pp. 3395–3404.
- [51] A. Diro, N. Chilamkurti, Leveraging lstm networks for attack detection in fog-to-things communications, IEEE Communications Magazine 56 (9) (2018) 124–130.
- [52] Z. Li, Z. Qin, A semantic parsing based lstm model for intrusion detection, in: ICONIP, 2018.