

SAFE-NID: Self-Attention with Normalizing-Flow Encodings for Network Intrusion Detection

Anonymous authors

Paper under double-blind review

Abstract

Machine learning models are increasingly being adopted to monitor network traffic and detect network intrusion. In this paper, we develop a deep learning architecture for traffic monitoring at the packet level. Current network intrusion detection models typically struggle when faced with zero-day attacks and concept drift. We introduce SAFE-NID, a novel approach to address these in practical deployments—in particular, with high efficacy and low latency. Contrasting with previous work, we train a relatively lightweight encoder-only transformer architecture for the packet classification task. Our deep learning framework introduces a normalizing flows safeguard that quantifies uncertainty in the decision made by the classification model. Our generative model learns class-conditional representations of the internal features of the deep neural network. We demonstrate the effectiveness of our approach by converting publicly available flow-level network intrusion datasets into packet-level ones. We release the labeled packet-level versions of these datasets with over 50 million packets each and describe the challenges in creating these datasets. We withhold from the training data certain attack categories to simulate zero-day attacks. Existing deep learning models, which achieve an accuracy of over 99% when detecting known attacks, only correctly classify 1% of the novel attacks. Our proposed transformer architecture with normalizing flows model safeguard achieves an area under the receiver operating characteristic of over 0.97 in detecting these novel inputs, outperforming existing combinations of neural architectures and model safeguards.

1 Introduction

An increasing amount of new malware originates every day with both targeted and random attacks occurring against individuals, businesses, and government entities at large scales. These attacks often exploit known vulnerabilities or attempt to overwhelm system capacity to steal privileged information, disrupt legitimate traffic, and cause economic hardship to the victims. With an ever-increasing number of connected devices, manual monitoring of network traffic to identify malicious behavior is simply too expensive and infeasible. This has motivated increased adoption of machine learning (ML) (Leevy & Khoshgoftaar, 2020), in particular deep learning (DL) (Ferrag et al., 2020), to aid and increase the bandwidth of cybersecurity professionals.

Traditionally, ML models for network traffic operate at the *flow-level* (Wang, 2015; Kim et al., 2019; Sarhan et al., 2020). Network flow represents a continued connection between two devices parameterized by the five-tuple: (src_ip, src_port, dest_ip, dest_port, protocol). Flow summary tools such as CICFlowMeter (Draper-Gil et al., 2016; Lashkari et al., 2017) and Zeek (zeek) produce features from the bidirectional flows on which an ML model can be trained. These features include flow duration, bytes per second, and flag settings. These methods are inherently *ex post facto*, however, as the flow must conclude before we know the final statistics and decide whether it was benign or identify its attack category. Furthermore, the features themselves are highly variable to the flow capture mechanism and its setup (Sarhan et al., 2020). For these reasons, we focus on the *packet-level* classification pipeline that labels traffic into benign and different attack categories at a per-packet granularity. Expanding on existing work of packet-level classification (Shenfield et al., 2018; De Lucia et al., 2021; Bierbrauer et al., 2022; Bizzarri et al., 2024; Rani & Bastian, 2024), we

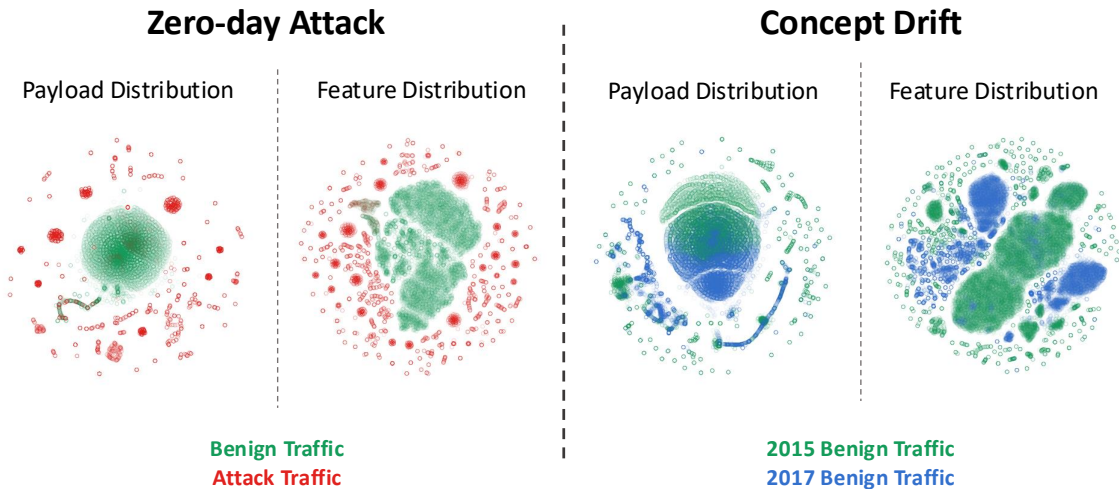


Figure 1: *Contrasting the input and feature space.* Deep neural networks trained on packet-level data can, with high accuracy, discriminate between benign and malicious traffic. However, these models fail in the presence of zero-day attacks and evolving network traffic. These drifts in payload byte distribution are often small, causing difficulty for OOD detection algorithms. We propose a “safeguard” generative model that learns the distribution of hidden layers from a DNN to identify novel inputs that the model cannot adequately evaluate. Shown above: t-SNE plots of the payload distributions and extracted feature distributions for benign and attack traffic (left) and evolving benign traffic (right).

include header context in inputs to our models to leverage additional information alongside the raw payload bytes. Furthermore, we design a light-weight encoder-only transformer model for the network intrusion detection task with a simple custom tokenization scheme.

While DNNs achieve high accuracy on inputs from the training distribution, they fail to generalize to novel inputs outside this distribution. Thus, their success on existing datasets does not entail their effectiveness in a real-world deployment where the network traffic is continuously evolving, and novel threats emerge routinely. We address this challenge by proposing a model-agnostic safeguard for DNNs that quantifies the uncertainty in the decision of any discriminative classification model by assigning a confidence score to each decision made by the DNN. The novel inputs are assigned low confidence and for such inputs, the decision by the machine learning model cannot be trusted. A high rate of detection of novel inputs by the monitor ensures wrong decisions by the DNN are avoided, and a low false positive rate ensures that the needed bandwidth of intervention (such as human experts) remains low. Thus, the proposed model safeguard improves the robustness of the DNN network intrusion detection models. This approach differs from unsupervised anomaly detection algorithms which look for concept drift before inference (Rabanser et al., 2019; Han et al., 2023). We note that this is not an either-or situation. One could first detect anomalous inputs in an unsupervised manner before using a discriminative classifier. However, the classifier would still be susceptible to non-flagged out-of-distribution inputs, requiring a model safeguard. We learn the intermediate latent representations of the inputs as the in- and out-of-distribution traffic are more easily separable in the latent space (Figure 1). Furthermore, we highlight the success of SAFE-NID, which uses a combination of encoder-only transformer model for packet classification and normalizing flows for the model safeguard within this framework.

Without such robust models, system operators relying on DNN models may have overconfidence in their network security. For example, incorrectly labeled benign network traffic may contain malicious payloads. In particular, DNN models need to remain robust in the presence of novel inputs such as zero-day exploits and distribution shift as network traffic and payload distribution evolves. Although we cannot currently expect models to generalize to correctly classify all packets from zero-day exploits without additional finetuning, it is critical to recognize these novel inputs as out-of-distribution (Lee et al., 2018; Geng et al., 2020; Bulusu et al., 2020; Shen et al., 2021; Yang et al., 2021). By adding model safeguards to the DNN models, we can

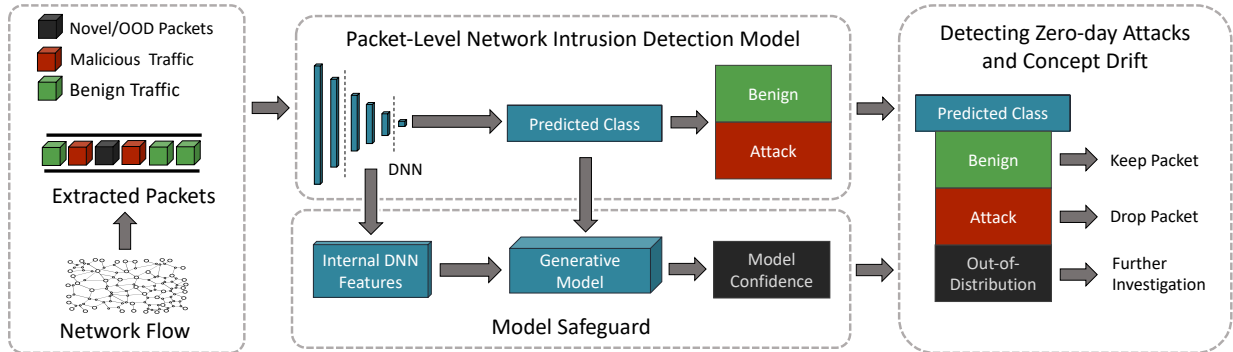


Figure 2: *Packet-level network intrusion detection system with model safeguards.* Our architecture ingests raw packet data and extracted information from the internet (IPv4) and transport layer (TCP/UDP) headers. We exclude information that is specific to the configuration used in data collection and not generalizable to the real-world. The predicted class and the internal features of the packet-level network intrusion detection model are used by the model safeguard. The monitor learns a class-condition distribution of the internal features and quantifies the uncertainty in prediction as a confidence score (that reflects the distance of internal features from the distribution of the features of the predicted class for any new input). The novel and out-of-distribution inputs are assigned low confidence score and can, hence, be detected for further investigation. We show strong results on detecting novel inputs when using SAFE-NID—a transformer architecture for our discriminative classifier and normalizing flows for our model safeguard.

quantify uncertainty in the predictions from the models as a confidence score. Previous research has shown DNN models make overconfident wrong predictions on novel out-of-distribution inputs when considering softmax output (Guo et al., 2017) as the confidence score. Furthermore, any such uncertainty quantification should occur in near real-time (Cai & Koutsoukos, 2020) to augment current cybersecurity workflows where attacks can significantly impair operations within milliseconds (Shan et al., 2017).

The overall uncertainty-quantified and robust deep learning architecture proposed in this paper is described in Figure 2. We complement the DNN model used for detecting attacks with a model safeguard that learns the class-conditional distribution of the internal features of the DNN model over the training data. For any new input at inference time, the predicted class and the internal features from the DNN model are used to query the learned distribution model for the likelihood of the input being in-distribution. The confidence of the prediction is high for inputs which have higher likelihood. The model safeguard-based architecture proposed in this paper applies to any DNN model and architecture. We observe very high in-distribution accuracy for the network intrusion datasets for relatively simple DNN architectures such as feed-forward neural networks and convolution neural networks, as well as more complex architectures such as transformers. Our model safeguard and uncertainty quantification approach is agnostic to the DNN architecture, and could even be used with non deep learning approaches for latency restricted applications. However, we note best results with SAFE-NID: a light-weight encoder-only transformer architecture and normalizing flows safeguard within our framework.

We make the following contributions in this paper:

- We construct a network intrusion detection system that analyzes information at the packet-level. Traditional flow-level detection needs the flow to conclude before an attack can be detected. In contrast, the packet-level detection enables online nearly real-time detection of attacks and also avoids the need for engineering flow features (De Lucia et al., 2021; Bierbrauer et al., 2022).
- We extract packet-level network intrusion detection datasets from flow-level datasets (Sharafaldin et al., 2018b; Moustafa & Slay, 2015), and make these available to the wider research community. We enumerate several challenges of using the flow-level datasets in this manner and provide guidance for extracting packet-level data from other similar datasets.

- We train three DNN models, a CNN, FNN, and transformer, on these datasets that achieve high accuracy. We empirically demonstrate that these models are not robust to inputs from novel attacks and concept drift. The baseline models, despite having 99% accuracy on the in-distribution inputs, exhibit very poor accuracy ($< 1\%$) on novel attack classes.
- We develop model safeguards to protect our DNNs against OOD inputs and concept drift. We show effectiveness in identifying suspect inputs using two different generative models (Lee et al., 2018; Papamakarios et al., 2021) to learn and represent the class-conditional distribution of internal features of the discriminative deep learning model being used for classification. The model safeguards can quantify uncertainty in the decision made by the underlying DNN models, and detect novel out-of-distribution inputs, despite the high imbalance of in-distribution to out-of-distribution samples (600–2000 : 1).
- SAFE-NID: Our light-weight encoder-only transformer discriminative classifier with our normalizing flows model safeguard outperforms existing baseline combinations of classifiers and out-of-distribution detection methods. The confidence score from our model safeguard achieves a high AU ROC of over 97% on detecting novel classes. This makes our framework robust to new attack classes.
- We make our processed packet-level dataset freely available and our code open source to encourage further research on packet-level network intrusion detection.¹

2 Related Works

A significant body of research concerns the identification of malware at different points of the attack pipeline. Some ML models consider detecting malicious software at the earliest possible point by analyzing raw binary files. Early methods focused on identifying structural patterns within attributes from the PE headers in Windows executables (Saxe & Berlin, 2015; Anderson & Roth, 2018; Vinayakumar & Soman, 2018). More recent strategies attempt to exploit the advancements in computer vision by converting the raw binaries into images using N-grams (Raff et al., 2018; Mohammed et al., 2021). As an alternative, Ling *et al.* construct control graphs by unpacking and disassembling binaries (Ling et al., 2022). Although these methods achieve high classification accuracy, they are limited in scope as many frequent attacks exploit buffer overflows by tailoring inputs to otherwise benign programs (Butt et al., 2022). Furthermore, some disruption-oriented attacks such as the many variants of denial-of-service exploit the very processes that enable devices to interact like the SYN-ACK attack (Schuba et al., 1997). Network intrusion detection systems (IDS) classify network traffic as benign or malicious (Moustafa & Slay, 2015; Sharafaldin et al., 2018b). These setups can identify ongoing denial-of-service attacks as well as advanced persistent threats (APTs) that have already infiltrated a compromised system.

Most existing efforts on learning-based network intrusion focus on classifying traffic flows with several available datasets (Moustafa & Slay, 2015; Sharafaldin et al., 2018b). However, there are several difficulties in creating accurate training datasets for network intrusion detection such as adequately anonymizing data and correctly identifying malicious flows, among other problems. Therefore, the Canadian Institute for Cybersecurity (CIC) and Intelligent Security Group (ISG) at the University of New South Wales (UNSW) produced several datasets that capture both benign and malicious network traffic in small, simulated environments (Moustafa & Slay, 2015; Sharafaldin et al., 2018b; Koroniotis et al., 2019; Moustafa, 2021). These datasets typically provide raw pcap files summarizing all of the packets that pass through a victim network. Frequently, the authors provide bidirectional flow-level summaries produced by tools such as CICFlowMeter (Draper-Gil et al., 2016; Lashkari et al., 2017) or Zeek (zeek). These summaries include features such as flow duration, number of forward and backward packets, and number of flags set, among others.

A significant amount of research considers the problem of manually engineering features to capture traffic flows and using these features for the detection of malicious flows. Several different neural network architectures have been explored (Wankhede & Kshirsagar, 2018; Vinayakumar et al., 2019; Maxwell et al., 2019; Pelletier & Abualkibash, 2020) in addition to the traditional machine learning models such as Adaboost (Yulianto et al., 2019) or random forests (Wankhede & Kshirsagar, 2018). Kim et al. (2019) train

¹Link omitted for review.

a CNN classifier on these features by transforming the 78-dimensional feature space into images. Since flow-level detection is *ex post facto*, real-time detection requires packet-level classification. There has been some work on packet-level detection of malicious traffic (Wang, 2015; Shenfield et al., 2018; De Lucia et al., 2021; Bierbrauer et al., 2022; Bizzarri et al., 2024; Rani & Bastian, 2024) wherein, the model takes as input the raw payload data, typically from the transport layer. Care must be taken to omit any potentially trivial information that could cause the ML model to learn the system configuration used in the data gathering process such as the IPs of the attack devices. This would lead to artificially high accuracy on the collected data which would not generalize to real-world environments. These previous approaches have generally focused on the discriminative models, with some analysis on transfer learning in the face of concept drift (Bierbrauer et al., 2022). However, they have not considered the problem of detecting inputs outside of the training distribution.

The problem of overconfident incorrect prediction by deep learning models on novel inputs that are out-of-distribution has been previously reported in domains such as computer vision (Guo et al., 2017) and several methods to compute confidence of deep learning models in these domains have been proposed (Hendrycks & Gimpel, 2016; Jiang et al., 2018; Jha et al., 2019; Ren et al., 2019). Broadly, these methods can be categorized into two classes. The first are the supervised techniques (Lee et al., 2017; Hendrycks et al., 2019; Meinke & Hein, 2019; Kaur et al., 2021) that require some exposure to the out-of-distribution inputs. This is impractical for cybersecurity where we cannot assume even a small number of packets corresponding to novel attack classes would be available at training time. Unsupervised approaches use only the in-distribution data. A direct approach is to use the softmax score of the classifier as a confidence metric (Hendrycks & Gimpel, 2016). ODIN (Liang et al., 2017) enhances the softmax score by adding perturbations to the input and using temperature scaling to the classifier’s confidence. Recently, (Macedo et al., 2021) proposed replacing softmax scores with isomax scores and entropy maximization for detection. Other unsupervised detection techniques based on energy scores (Liu et al., 2020), trust scores (Jiang et al., 2018), and likelihood ratio (Ren et al., 2019) between the in-distribution and OOD data points have been proposed for detection. In contrast, we use a normalizing flow generative model for learning the distribution of internal features. These uncertainty-quantification and out-of-distribution detection approaches rely on learning the manifold or distribution of training data and are known to be susceptible to adversarial attacks (Jang et al., 2020).

In this paper, we focus on robustness to novel inputs and out-of-distribution data. The goal is to build a detection pipeline where the generative safeguard assigns high uncertainty to recognize novel kinds of traffic or attacks which are different from the training data and on which the discriminative classifier cannot be trusted. Our attack model does not consider the scenario where an active attacker uses knowledge about our model safeguard and the detection pipeline to craft an attack specifically for our framework. The robustness to novelty is an important characteristic for building a robust ML-enabled system and is distinct from resilience to adversarial attacks. Our work is similar to novelty detection, out-of-distribution detection, and open-set recognition studied in vision and natural language domains. The robustness to novel and out-of-distribution is even more critical in cybersecurity applications, and our framework addresses this challenge.

3 Overview

We provide an overview of our system architecture in Figure 3. We take as input pre-existing flow-level IDS datasets that also provide raw pcap files (Section 4.1). Typically, these flow-level datasets use a flow summary tool to create a set of labeled features such as flow duration, and the number of forward and backward packets, among others. The flow labels come from existing knowledge of the test-bed architectures where set IPs and timestamps correspond to malicious activity of a certain type. Since the pcap files provided with the flow-level datasets are not labeled, we process the data to create packet-level datasets (Section 4.2). Since these existing datasets did not intend for this type of conversion, we identify a series of issues and solutions in this process. We convert the packets into input features for our neural networks by first taking the first 1,500 bytes of the payload, and second, extracting relevant features (Section 4.3). We need to take care to not extract features that are trivially correlated with benign and malicious behavior, such as source or destination IP. We consider an encoder-only transformer neural network for the binary classification task for packets (Section 4.4). Although our transformer performs with high accuracies (> 99%)




Figure 3: *System architecture overview*. We first take existing labeled flow-level datasets (box 1) and convert them into packet-level datasets by matching packets in a pcap file to their corresponding flows (box 2). We extract header features from the packets to augment our ML feature space, taking care to not include features that could trivially classify packets as benign or malicious, such as source and destination IP (box 3). A transformer takes as input the payloads and header features to produce a binary classification of packets as benign or malicious (box 4). Despite high accuracy on the binary classification task, our model fails when given novel inputs, such as zero-day exploits (box 5) and inputs from a different time (box 6). We propose a normalizing flows model safeguard that takes the internal features from the DNN, models their distribution, and produces confidence in the predicted class (box 7). This enables us to have a robust decision-making process where we can classify inputs as benign, attack, or novel (box 8).

on in-distribution data, we find that the networks perform poorly when we mimic novel inputs (e.g., zero-day exploits) (Section 4.5), or when the data consists of packets from a different year (Section 4.6). We extract internal features from the neural network and model their distributions using normalizing flows to produce a confidence (or novelty) score to produce during inference (Section 4.7). This end-to-end system enables us to produce packet-level predictions with an uncertainty-quantification. Although we focus on SAFE-NID (the encoder-only transformer packet classifier and normalizing flows safeguard), the overall framework is agnostic to both models, allowing us to switch models based on real-world computational constraints.

4 Technical Approach

Throughout our work, we use the following terminology: a *packet label* refers to the binary classification of a packet (i.e., benign or malicious), whereas a *packet category* refers to the multiclass classification of a packet that further granularizes attack types (i.e., benign, denial-of-service, heartbleed, fuzzers, etc.).

4.1 Packet-level Capture

Most existing network IDS datasets provide flow summaries with accompanying benign and attack category labels (Moustafa & Slay, 2015; Sharafaldin et al., 2018b). We instead focus on classifying packets as benign and malicious for a couple of reasons.

First, flow capture devices extract different features from pcap files which creates difficulties when designing ML solutions that must work across a wide range of differently configured networks. Sarhan *et al.* use NetFlow (Claise, 2004) to extract features from four publicly available IDS datasets (Sarhan et al., 2020). NetFlow is easy to configure and generates features using almost solely packet headers; this ease of use, however, leads to a less expressive feature space compared to hand-tailored flow capture systems. In contrast, different packet capture technology running on different operating systems produce standardized PCAP files.

Second, learning benign and attack characteristics at the flow level is inherently a retroactive process. One has to wait for a flow to conclude before extracting features such as the number of forward and backward packets sent, mean payload length, and average time between packets, amongst others. Although this can be very useful in postmortem contexts where a network administrator wants to analyze fault points in a system’s security after an attack concludes, it is a less viable paradigm for identifying ongoing and incoming attacks. By classifying packets as they arrive as benign or malicious, one can actively reject traffic before the payload arrives at its intended destination.

4.2 Preprocessing Data

We first must transform the raw pcap files into packet-level labeled datasets for training and inference. Since most of the existing literature on intrusion detection in network traffic focuses on flow-level features, existing datasets typically provide category labels (i.e., benign, denial-of-service, exploits, heartbleed, etc.) only at the flow-level.

Our method ingests the outputs from a flow capture system and the raw pcap files. We characterize each flow with the following five-tuple:

$$(\text{src_ip}, \text{src_port}, \text{dest_ip}, \text{dest_port}, \text{protocol})$$

Note that this five-tuple is often not discriminative enough on the IDS datasets since the victim and attack devices have set IPs by the architectural design of the testbed. That is, multiple flows will have the same five-tuple, and not all of those flows will have the same packet category or even label. Therefore, we need to not only identify a flow with the matching five-tuple but also select the flow with the proper timeframe.

Flow capture systems can be either unidirectional or bidirectional (Li et al., 2013). Bidirectional systems, such as CICFlowMeter (Lashkari et al., 2017; Draper-Gil et al., 2016) and NetFlow v9 (Claise, 2004), create a single flow once a host initiates a connection. These bidirectional flows cover any response packets to the initiating host. Therefore, these response packets have corresponding flows with the reciprocal of the standard five-tuple:

$$(\text{dest_ip}, \text{dest_port}, \text{src_ip}, \text{src_port}, \text{protocol})$$

Thus, we need to consider flows with both the standard five-tuple and its reciprocal during packet matching.

The pcap files record the packet timestamps in Coordinate Universal Time to microsecond precision. Note, we assume that the capture device has a correctly calibrated internal time. Unfortunately, flow summary tools do not have one standardized representation of the start and end times of the flows. Therefore, we once again need to be cognizant of the internal workings of the flow capture system. For example, the flows published with the CIC-IDS-2017 dataset use local start times (UTC-3) (Sharafaldin et al., 2018b) while those in the UNSW-NB15 dataset use universal time (Moustafa & Slay, 2015).

The disparity between the precision of the packet timestamps in the pcap files versus the precision of the start and end times in the flow files is perhaps as big an issue when accurately assigning packets to flows for category and label assignment. The flows that we consider for this work have time precision to the minute at worst and the second at best. However, the average duration of a flow is 1.48 (± 3.37) and 0.66 (± 13.93) seconds for the CIC-IDS-2017 and UNSW-NB15 datasets, respectively. The packet timestamps from the pcap files have microsecond precision. Therefore, we will often not identify any valid flows for a packet given the five-tuple with timeframe:

$$(\text{start_time}, \text{start_time} + \text{duration})$$

For example, we cannot match any packets that belong to a flow that starts at the half second (which the summary file rounds down) and only lasts a quarter second. Therefore, when matching packets to flows, we consider any five-tuple matching flow where the packet timestamp is in the range:

$$(\text{start_time}, \text{start_time} + \text{duration} + \text{precision})$$

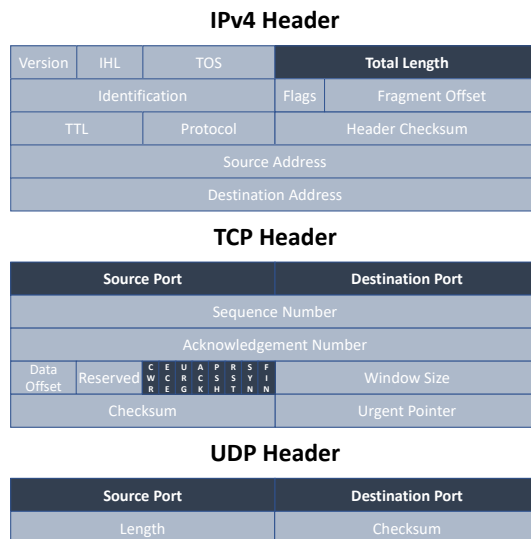


Figure 4: *Extracted header features*. We extract features from the packet headers (darkened fields) for our binary classification task. We do not include header information, such as source and destination IP, that could cause the system to learn the testbed architectures. We convert the source and destination ports into a binary vector that bins certain ports together (Table 1).

where precision is the precision of the start_time attribute of the flow summary files. However, this wider range of acceptable times often causes a given packet to match with several flows. We apply a conservative rule where a packet receives a packet category if and only if every valid flow has the same category as the one being applied to the packet. This prevents us from misclassifying packets to the wrong category or label.

We drop any packets that do not follow the IPv4 protocol at the internet layer and TCP or UDP protocols at the transport layer. Lastly, we drop all packets with empty payloads.

4.3 Extracting Header Features

We extract information from the packet headers to augment our input feature space. However, we cannot simply input the entire header as our models will learn the setup configuration of the testbed architectures. For example, an ML model could trivially learn the source and destination IP addresses that correspond to devices in the attack network. Therefore, we carefully consider the header fields that provide useful contextual information that is transferable to other network configurations. We construct features from the darkened header fields shown for the IPv4, TCP, and UDP packet protocols in (Figure 4).

We only make use of the “Total Length” field in IPv4 headers. We divide this field by 65,536 to standardize the value between 0 and 1. Notably, we do not consider information in the “TTL”, “Protocol”, or “Address” fields. Although time-to-live might be useful in real-world applications with diverse network topologies, the publicly available datasets have very few plausible attack paths because of the limited number of victim and attack devices. We found that one of the most cited IDS datasets (Sharafaldin et al., 2018b) almost exclusively uses the TCP transport protocol for attacks (only two UDP attack packets). Therefore, we exclude the protocol field since the great imbalance is not indicative of the real world where some attacks, such as UDP Flood, use the UDP protocol (Dittrich, 1999). Lastly, most network IDS datasets have set IP addresses for the victim and attack devices. Training data with those fields would produce a model unable to generalize to other configuration settings.

For the transport layer headers, we produce a set of features based on the source and destination ports. We cannot simply produce a one-hot encoding for each port since many ports would indicate benign or attack based on the scripts that generate attacks on IDS datasets. After considering some commonly used network

Table 1: *Port features for header context.* We construct seven indicator features from the source and destination ports. Features for sets that contain a given port number receive a value of one.

Port Number	Description
21	File Transfer Protocol (FTP)
22	Secure Shell (SSH)
80/8080	Hypertext Transfer Protocol (HTTP)
443/444	Hypertext Transfer Protocol Secure (HTTPS)
0 - 1,023	Well Known Ports
1,024 - 49,151	Registered Ports
49,152 - 65535	Dynamic/Private Ports

IDS datasets, we generate seven binary features for both the source and destination ports of every packet. These features are non-exclusive, i.e., at least one but perhaps two of the features can receive a value of one. We can summarize the features as which of the sets in Table 1 include the port number. Note that we group ports 443 and 444—the heartbleed attacks on the CIC-IDS-2017 dataset exploit the SSL vulnerability on port 444. Lastly, we include the following eight flags from the TCP header: CWR, ECE, URG, ACK, PSH, RST, SYN, and FIN. We simply set these features to zero for UDP packets.

4.4 Encoder-Only Transformers for Packet Data

We consider an encoder-only transformer architecture for the sequence to classification task of predicting whether a payload is benign or malicious. We use a simple tokenization scheme to convert our payloads into vectors of tokens. We convert each byte in the payload into the corresponding ASCII number. Thus, we have a vocabulary size of 256 corresponding to the 256 character codes. We find existing subword tokenization strategies (Kudo & Richardson, 2018; Song et al., 2020) provide sub-optimal results on payload data, in part because encryption and compression create high entropy payloads that do not correlate to the structured text in natural languages. Our encoded payloads are then input into a series of transformer encoder blocks to create an embedding for each token (input byte) (Vaswani et al., 2017). We then use mean pooling to convert the array of embeddings into a single sentence embedding (Reimers & Gurevych, 2019). We concatenate the header context if applicable to the sentence embedding and input the resultant vector into three hidden fully connected layers with 256, 128, and 64 units. Each of the hidden layers has a LeakyRELU activation with $\alpha = 0.01$ (Maas et al., 2013) and batch normalization (Ioffe & Szegedy, 2015) and dropout regularization (Srivastava et al., 2014). The two output neurons have a softmax activation function.

4.5 Handling Zero-day Exploits

Machine learning models generally perform well when given in-distribution testing data, i.e., data similar to the training data. However, real-world cybersecurity defense implementations will eventually receive data that falls out of the training distribution. For example, zero-day exploits cannot by definition exist in the training data. It is imperative that any machine learning systems do not fail when such exploits appear. Here, we focus on false negatives where our model classifies packets from previously unseen attack paradigms as benign.

We see similar characteristics between the zero-day exploit problem and the open-set recognition one in the broader machine learning community (Geng et al., 2020). In open-set recognition, a model trains on incomplete data that does not include all classes that exist during inference. We model this behavior by training ML models on the benign traffic and all but one type of attack. We repeat this procedure for each attack category, to create N trained models where each model corresponds to a single missing attack type. We consider two different metrics for success for our models: first, we can correctly classify the unseen attack types as malicious, and second, we can recognize that the attack types are unfamiliar and flag them as out-of-distribution. Correctly classifying unseen attacks as malicious indicates that our model is learning general attack patterns in the payloads themselves. For example, some of the brute force attack methods have similar characteristics and so removing one from training does not change the quality of inference. However, more

Table 2: *Dataset port distribution.* We extract seven binary features from the source and destination ports each. A feature has a value of one if the corresponding source/destination port belongs to the indicated set.

Port(s)	UNSW-NB15		CIC-IDS-2017	
Source Port	No. Benign	No. Attack	No. Benign	No. Attack
21	689,008	19,288	40,417	17,975
22	1,592,854	284	128,195	43,651
80/8080	15,021,418	450,780	13,912,273	757,410
443/444	0	268	7,916,684	19,020
Well Known (0 - 1023)	19,186,346	1,015,630	24,025,879	837,249
Registered (1024 - 49151)	28,350,557	1,098,046	1,436,391	166,559
Dynamic/Private (49152 - 65535)	2,217,170	344,656	3,722,886	216,220
Destination Port	No. Benign	No. Attack	No. Benign	No. Attack
21	500,000	17,192	30,470	17,947
22	1,187,386	314	124,003	46,594
80/8080	354,910	112,911	589,520	286,401
443/444	0	394	2,305,603	31,026
Well Known (0 - 1023)	4,366,251	1,384,884	5,197,451	380,195
Registered (1024 - 49151)	34,715,404	789,883	3,618,349	440,957
Dynamic/Private (49152 - 65535)	10,672,418	283,565	20,369,356	398,876

often, when we remove an attack from the training data, we classify those packets on inference as benign. Thus, it is imperative to identify those packets as out-of-distribution (Section 4.7).

4.6 Handling Distribution Shift

Internet most-common and best practices continually evolve, especially as privacy and security concerns become higher societal priorities. As the landscape morphs, the structure of the payloads changes, and previously trained models can become outdated. In recent years, an ever-increasing number of websites have transitioned from HTTP connections to HTTPS (signified by the rows for ports 80/8080 and 443/444 in Table 2). Sometimes this evolution occurs organically as more websites adopt existing libraries to improve privacy. In such instances, we expect a moderate distribution shift over small stretches of time. We model the natural distribution shift of network traffic by training on a dataset from Q1 2015 and inferring on one from Q3 2017, and vice versa. This 2.5-year difference corresponds to a significant period of change from unencrypted (HTTP) to encrypted web traffic (HTTPS), from approximately 30% in Q1 2015 to 60% in Q2 2017, per the percentage of web pages loaded by Firefox using HTTPS (letsencrypt). However, occasionally targeted government mandates or external business pressures encourage the rapid adoption of a new framework or protocol, such as when Google Chrome began labeling all HTTP connections as “Not Secured” in 2018 (Robbins, 2021). These “seismic” events can cause a significant break between the distributions of current network traffic from previous months (by Q2 2018, over 80% of web pages in the United States loaded by Firefox used HTTPS (letsencrypt)).

4.7 Normalizing Flows as a Model Safeguard

In the normalizing flows paradigm, a model learns a series of bijective transformations that can perform a one-to-one mapping from a simple distribution, such as a multivariate Gaussian, into a complex target distribution (Papamakarios et al., 2021; Kobyzev et al., 2020). Since each transformation block is invertible, we can take a feature vector and easily determine its corresponding value in the probability density function of the complex distribution by applying a series of matrix multiplications followed by the inverse of the simple activation functions. In this way, we extract features from the intermediate layers of our neural networks for the in-distribution training data and train normalizing flows that transform a multivariate Gaussian distribution into this complex space. During inference, we extract features from our network and transform the data into the multivariate Gaussian space. By looking at the loss (i.e., the probability that the vector belongs to the complex distribution), we can order our inputs by the probability that they

Table 3: *UNSW-NB15 packet distribution*. The UNSW-NB15 dataset contains over 52 million non-empty TCP and UDP payloads. The dataset captures nine attack categories as well as benign background traffic over two days.

Packet Category	No. Packets	Avg. Payload Length
Benign	49,754,073	2,140.15 ($\pm 1,195.68$)
Attack	2,458,332	2,236.90 ($\pm 1,140.02$)
Analysis	1,394	561.35 (± 273.80)
Backdoor	1,246	733.00 (± 929.74)
DoS	479,275	2,326.59 ($\pm 1,026.94$)
Exploits	1,619,721	2,250.03 ($\pm 1,148.98$)
Fuzzers	144,114	1,688.62 ($\pm 1,330.87$)
Generic	192,901	2,429.40 (± 971.32)
Reconnaissance	9,601	189.90 (± 140.33)
Shellcode	828	271.03 (± 229.82)
Worms	9,252	2,575.79 (± 851.72)

belong to the training set distribution. Inputs with a lower negative log-likelihood loss are more likely to be in-distribution. We train two normalizing flows for each model, one each for benign and attack network traffic. When evaluating the distance from our training distributions during inference, we only consider the normalizing flow model that matches the output label from our NN classifier.

There are several normalizing flow blocks common in the literature (Dinh et al., 2016; 2014; Kingma & Dhariwal, 2018; Sorrenson et al., 2020), with each offering advantages and trade-offs during training and forward and reverse inference. For our purposes, we use RealNVP blocks (Dinh et al., 2016) that we can parameterize with the following equation:

$$y = R\Psi(s_{\text{global}}) \odot \text{Coupling}(R^{-1}x) + t_{\text{global}} \quad (1)$$

where R is a (deterministic) permutation matrix that allows each feature to influence the others, \odot is the Hadamard Product (element-wise multiplication), x is the input vector for each block, s_{global} and t_{global} are learnable parameters, and Coupling is the following function that first evenly divides the input vector into halves x_1 and x_2 :

$$u = \text{concat}(u_1, u_2) \quad (2)$$

$$u_1 = x_1 \odot \exp(\alpha \tanh(s(x_2))) + t(x_2) \quad (3)$$

$$u_2 = x_2 \quad (4)$$

α is a clamping value that restricts the range of possible values in the exponent, and s and t are learnable parameters. Note that in this coupling block, features in x_2 can influence outputs in u_1 but not vice versa. The permutation matrix allows each feature to influence the others when stacking multiple blocks during training.

To avoid overfitting to our training distribution, we add a small amount of Gaussian noise to our features during training and inference with the following equation:

$$X = X + \mathcal{N}(0, 0.05) \quad (5)$$

This improves the stability of the training procedure and simultaneously washes out any features with no discernible signal. Such features are not unexpected as others have previously observed feature collapse in deeper layers (Zhu et al., 2021).

Table 4: *CIC-IDS-2017 packet distribution*. The CIC-IDS-2017 dataset contains over 30 million non-empty TCP and UDP payloads. The dataset captures fourteen attack categories as well as benign background traffic over the course of a week.

Packet Category	No. Packets	Avg. Payload Length
Benign	29,185,156	3,017.95 ($\pm 2,444.65$)
Attack	1,220,028	4,392.96 ($\pm 4,794.26$)
Bot	2,584	4,088.69 ($\pm 6,859.15$)
DDoS	220,160	6,166.16 ($\pm 5,822.03$)
DoS GoldenEye	26,198	5,484.42 ($\pm 5,509.26$)
DoS Hulk	747,340	4,810.35 ($\pm 4,467.81$)
DoS Slowhttptest	5,906	1,049.66 (± 965.19)
DoS Slowloris	23,593	404.80 (± 204.47)
FTP-Patator	35,922	44.06 (± 14.41)
Heartbleed	20,181	7,665.59 ($\pm 4,815.72$)
Infiltration	29,865	908.23 (± 551.25)
Port Scan	417	4,761.18 ($\pm 3,781.50$)
SSH-Patator	90,245	307.56 (± 454.89)
Web Attack-Brute Force	14,400	1,148.79 (± 590.99)
Web Attack-SQL Injection	19	2,511.68 ($\pm 1,966.98$)
Web Attack-XSS	3,198	2,280.02 ($\pm 1,398.51$)

5 Experimental Setup

5.1 Datasets

We evaluate our system using two different network intrusion detection datasets that publish both raw pcap data and corresponding hand-labeled flow data. Both datasets configure a testbed infrastructure with an attack and a victim network comprised of multiple connected devices. The attack networks initiate a series of different attack profiles throughout data collection.

UNSW-NB15. The UNSW-NB15 dataset (Moustafa & Slay, 2015) captures raw network traffic over two full (fifteen and sixteen hour) days in January and February 2015. In the testbed architecture, an IXIA traffic generator uses three virtual servers, two that spread benign network traffic and one that forms malicious activity. The servers pass all traffic through two routers connected to a firewall that allows all packets to pass through. The dataset authors installed tcpdump (Jacobson, 1989) on one of the routers to capture all packet data. The IXIA tool simulated nine different attack categories including analysis, backdoor, denial-of-service, exploits, fuzzers, generic, reconnaissance, shellcode, and worms. Table 3 provides an additional summary of each packet category after our processing method.

CIC-IDS-2017. The CIC-IDS-2017 dataset (Sharafaldin et al., 2018b) contains raw pcap data captured over a week-long period in July 2017. The devices on the victim network run different versions of the three most common operating systems (Windows, Mac, and Linux). Three of the attacking PCs have the Windows 8.1 operating system and the fourth has Kali Linux. The dataset authors configured one of the ports of the main switch as a mirror port that completely captures all traffic traversing into and out of the victim network. A B-Profile system profiles the abstract (benign) behavior of 25 different users and an automated agent derived from these profiles generates realistic benign events for packet capture (Sharafaldin et al., 2018a). The CIC-IDS-2017 dataset contains fourteen attack types from seven common attack families: brute force, heartbleed, botnet, denial-of-service, distributed denial-of-service, web, and infiltration. Table 4 provides a summary of the number of packets and average payload lengths for each packet category after our processing method.

We note that others (Engelen et al., 2021; Liu et al., 2022) have found issues with this particular dataset caused in part by bugs in the CICFlowMeter (Draper-Gil et al., 2016; Lashkari et al., 2017) tool. We avoid

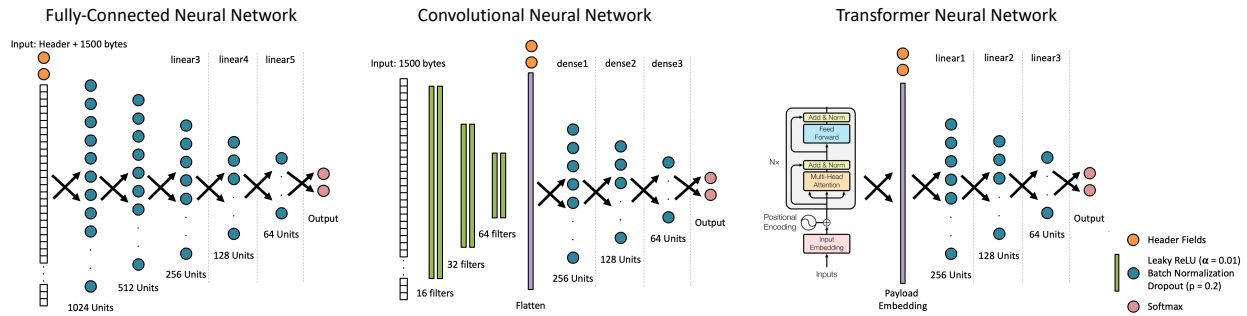


Figure 5: *DNN architectures for packet classification.* We train an FNN, CNN, and transformer architecture for the packet label classification class. All networks take as input the first 1,500 bytes of the packet payload. We concatenate the header features to the payload for the FNN on input, the CNN after the last convolution, and the transformer after the sentence embedding. Schematic for transformer block is from Vaswani et al. (2017). All networks have two output neurons with a softmax activation.

many of the troublesome artifacts by approaching the problem at a packet-level. For example, we drop any packets that have a “TCP appendix” error simply since we do not consider empty payloads. The other main artifacts come from attempted attacks that fail to deliver their malicious payloads. We still include these packets in our analysis, as watching an attack unfold remains a valuable goal, even before the malicious content arrives.

5.2 Data Division

Both datasets have a high class imbalance with 20–25 times as many benign packets as attack ones. We split our data into training/validation and testing sets. For the training and validation data, we take half of the malicious samples and an equal number of benign ones. Of this data, we use 75% for training and 25% for testing. We take ten random splits in this fashion and show averages and standard deviations for all published results. Our test data contains the remaining malicious and benign samples. Thus, during inference, we have 39.5 and 46.8 times more benign packets than malicious ones for the UNSW-NB15 and CIC-IDS-2017 datasets, respectively. For the concept drift results, we use balance inference data with an equal number of benign and malicious packets for both datasets during inference.

5.3 Baselines

5.3.1 Discriminative Classifiers

We consider fully-connected neural networks (FNNs) and convolutional neural networks (CNNs) as baselines for the binary classification task (De Lucia et al., 2021; Bierbrauer et al., 2022). For both architectures, we truncate payloads at 1500 bytes following practices in existing literature. Many ethernet networks have a maximum transmission unit size of 1500 bytes. We zero-pad payloads with fewer than 1500 bytes. We normalize all bytes to be between 0 and 1 for our FNNs and CNNs. We train all models with and without header context. Figure 5 summarizes the model architectures.

Fully-connected Neural Network. Our FNN model has five hidden layers starting with 1024 neurons and halving deeper into the network. We concatenate the header features to the 1500 payload bytes. Each neuron has a LeakyRELU activation with $\alpha = 0.01$ (Maas et al., 2013) and batch normalization (Ioffe & Szegedy, 2015) and dropout regularization (Srivastava et al., 2014). Our final output neurons use a softmax activation function.

Convolutional Neural Network. Our CNN model inputs the 1500 bytes into three blocks of two convolutions with a kernel size of three and 16, 32, and 64 filters, respectively, followed by a max-pooling by a factor of two. We flatten the output of the final max-pooling and concatenate the header features. We input

the flattened vector into three hidden fully connected layers with 256, 128, and 64 units. Each of the hidden layers has a LeakyRELU activation with $\alpha = 0.01$ (Maas et al., 2013) and batch normalization (Ioffe & Szegedy, 2015) and dropout regularization (Srivastava et al., 2014). The two output neurons have a softmax activation function.

5.3.2 Model Safeguards

Maximum softmax probability. The maximum softmax probability (MSP) method for out-of-distribution detection uses the highest output from the final layer of the neural network (Hendrycks & Gimpel, 2017). This early baseline for out-of-distribution detection assumes a well calibrated model which will assign low probabilities for samples it cannot accurately classify. Since we have only two output classes, benign or attack, our normalized softmax probabilities range between 0.5 and 1.0.

Gaussian kernel density. Gaussian kernel density out-of-distribution methods obtain class conditional Gaussian distributions for the extracted intermediate features from our NNs (Lee et al., 2018). These methods assume that a class-conditional Gaussian distribution can adequately fit these features, and therefore inputs whose features are farther from the closest class-conditional Gaussian distribution are more likely to fall out-of-distribution (Lee et al., 2018). We approximate the covariance matrix of the dataset using the maximum likelihood estimate (empirical covariance). Once we obtain the covariance matrix, we use the Mahalanobis distance to measure the distance between our input points P and the class-conditional distributions D_i (Mahalanobis, 1936). Mahalanobis distances measure the number of standard deviations from P to the mean of each D_i . We only calculate the distance for each point P to the distribution corresponding to the classifier prediction (benign or attack).

5.4 Training Parameters

Our FNN and CNN baseline models each use batch normalization (Ioffe & Szegedy, 2015) and dropout ($p = 0.2$) (Srivastava et al., 2014) regularization techniques. Each hidden layer activation function is LeakyRELU ($\alpha = 0.01$). Both of the classification networks use the AMSGrad variant of the Adam optimizer (Kingma & Ba, 2014; Reddi et al., 2019) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a learning rate of $1e-4$. We use the binary cross-entropy loss function. We train each network for 20 epochs and use the weights with the lowest validation loss.

For our transformer architecture, we use an embedding size of 384, used previously in sentence transformer sequence to classification tasks (Reimers & Gurevych, 2019). We only stack two transformer blocks, and each block has six self-attention heads for 64-dimensional key, value, and query vectors (Vaswani et al., 2017). Similar to the FNN and CNN models, we use batch normalization (Ioffe & Szegedy, 2015) and dropout ($p = 0.2$) (Srivastava et al., 2014) regularization techniques for our fully connected layers after the transformer block. We use the AMSGrad variant of the Adam optimizer (Kingma & Ba, 2014; Reddi et al., 2019) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a learning rate of $3e-4$. We use the binary cross-entropy loss function and train each network for six epochs.

For our normalizing flow models, we stack 20 RealNVP blocks with an affine clamping $\alpha = 2$. We learn our parameters for s and t using a simple fully connected network with two hidden layers with 128 features each and LeakyRELU activation with $\alpha = 0.01$. The input and output dimensions of these learnable blocks are dependent on the size of the extracted NN layers (either 256, 128, or 64 dimensions). We use the Adam optimizer (Kingma & Ba, 2014) with $\beta_1 = 0.8$, $\beta_2 = 0.9$, a learning rate of $1e-4$, and weight decay of $2e-5$. We train each normalizing flow model for 512 epochs. We use 25% of our in-distribution data as validation data, stratifying by packet category.

5.5 Implementation

We implement our system in Python using Payload-Byte (Farrukh et al., 2022) for extracting and labeling packet capture (PCAP) files of modern NIDS datasets, Pytorch (Paszke et al., 2019) for our neural net-

Table 5: *Discriminative classifier results.* Each neural architecture achieves a high accuracy on the binary packet-level classification task. The transformer architecture outperforms the others on the CIC-IDS-2017 dataset, whereas the FNN achieves the highest accuracy and F1-Score on the UNSW-NB15 dataset. These datasets are both highly imbalanced with 39.5 and 46.8× more benign samples than malicious ones.

Architecture	Accuracy (\uparrow)	AU ROC (\uparrow)	F1-Score (\uparrow)
UNSW-NB15			
CNN	0.9950 (± 0.0003)	0.9997 (± 0.0000)	0.9952 (± 0.0003)
FNN	0.9951 (± 0.0001)	0.9998 (± 0.0000)	0.9953 (± 0.0001)
Transformer	0.9938 (± 0.0011)	0.9999 (± 0.0000)	0.9941 (± 0.0010)
CIC-IDS-2017			
CNN	0.9940 (± 0.0015)	0.9955 (± 0.0016)	0.9944 (± 0.0013)
FNN	0.9926 (± 0.0020)	0.9955 (± 0.0014)	0.9931 (± 0.0017)
Transformer	0.9965 (± 0.0011)	0.9973 (± 0.0007)	0.9966 (± 0.0010)

Table 6: *Modeling zero-day exploits.* The “In-Distribution” column refers to one model trained with all attack types and provides the recall for each attack on withheld testing data. The “Out-of-Distribution” column refers to 14 models with the indicated attack excluded from training. Some attack types have significant similarities with others and so excluding those does not greatly decrease their recall.

Attack Category	No. Test Packets	In-Distribution (\uparrow)	Out-of-Distribution (\uparrow)
Bot	1,292	0.9424 (± 0.0270)	0.0950 (± 0.0935)
DDoS	110,080	1.0000 (± 0.0000)	0.2145 (± 0.0258)
DoS GoldenEye	13,099	0.9999 (± 0.0001)	0.1995 (± 0.0426)
DoS Hulk	373,670	1.0000 (± 0.0000)	0.7529 (± 0.1626)
DoS Slowhttptest	2,953	0.9980 (± 0.0025)	0.5248 (± 0.1659)
DoS Slowloris	11,796	0.9996 (± 0.0002)	0.1737 (± 0.2859)
FTP-Patator	17,961	0.9994 (± 0.0002)	0.0030 (± 0.0019)
Heartbleed	10,090	0.4418 (± 0.0660)	0.0002 (± 0.0002)
Infiltration	14,932	0.9946 (± 0.0051)	0.0780 (± 0.0312)
Port Scan	208	0.9856 (± 0.0071)	0.8058 (± 0.0234)
SSH-Patator	45,123	0.9998 (± 0.0001)	0.0027 (± 0.0026)
Web Attack-Brute Force	7,200	0.9997 (± 0.0001)	0.9439 (± 0.0701)
Web Attack-SQL Injection	9	0.7222 (± 0.1745)	0.4678 (± 0.1382)
Web Attack-XSS	1,599	0.9959 (± 0.0030)	0.6305 (± 0.1717)

works and the Framework for Easily Invertible Architectures (FrEIA) (Ardizzone et al., 2018-2022) for our normalizing flows. We publish our code to encourage further investigation by the research community.²

6 Results

6.1 Packet-level Classification Accuracy

Table 5 shows the accuracy for the all neural network architectures. In all configurations, our model accurately classifies packets with over 99.26% accuracy. These results come on highly imbalanced datasets, as discussed in Section 5.2. On the UNSW-NB15 dataset, the FNN outperform the CNN and transformer architectures on the accuracy and F1-score metrics. The transformer architecture records the best results on all metrics on the CIC-IDS-2017 dataset.

²Link omitted for review.

Table 7: *Modeling distribution shift.* The overall accuracy on the UNSW-NB15 and CIC-IDS-2017 test sets decreases significantly when trained on the other. The left columns show the accuracy on the testing sets of each dataset when the corresponding training data is included during training. The right two columns show the inference accuracy when the training process excludes any examples from the given dataset.

Architecture	Accuracy When In-Distribution		Accuracy When Out-of-Distribution	
	UNSW-NB15	CIC-IDS-2017	UNSW-NB15	CIC-IDS-2017
CNN	0.9969 (± 0.0000)	0.9923 (± 0.0003)	0.4861 (± 0.0031)	0.5221 (± 0.0285)
FNN	0.9968 (± 0.0000)	0.9913 (± 0.0003)	0.4733 (± 0.0144)	0.5114 (± 0.0085)
Transformer	0.9967 (± 0.0006)	0.9940 (± 0.0010)	0.4891 (± 0.0095)	0.4024 (± 0.0363)

6.2 Handling Zero-day Exploits

We reformulate the problem of classifying inputs from zero-day exploits into the more general open-set recognition task in machine learning. Using this strategy, we train models on all benign packets and all but one type of attack. We then can evaluate how well our model predicts packets from the withheld classes. In Table 6, the “In-Distribution” column shows the recall of predicting a packet from the category in the first column as an attack. For example, we classify 99.94% of FTP-Patator packets as malicious when we include FTP-Patator samples in the training data. The last “Out-of-Distribution” column shows the recall for a given attack when we exclude that attack from training while still including the other attacks.

We highlight three specific attack categories where recalls fall precipitously when excluded from training: FTP-Patator, Infiltration, and SSH-Patator. When included in the training, we classify between 99.46–99.98% of these malicious payloads. However, our models achieve accuracies between 0.27–7.80% when excluding the various classes from training. Perhaps more surprisingly, we can correctly classify several attack categories as malicious despite removing them from training, suggesting a general hierarchy of attacks where some are merely derivative of others. The WebAttack-BruteForce, DoS Hulk, and Port Scan attacks see only limited degradation in recall.

We show the results for the CNN in Table 6. However, we note similar trends with the other configurations. We focus solely on FTP-Patator, Infiltration, and SSH-Patator when evaluating our real-time ML monitoring since those attack categories had both a significant number of packets and a large drop in recall when in- and out-of-distribution.

6.3 Handling Distribution Shift

We note substantial declines in accuracy when we train on one dataset and infer on the other (Table 7). In this table, the leftmost two results columns show the accuracies on the test set from the UNSW-NB15 and CIC-IDS-2017 datasets when we use the corresponding training data on our models. The rightmost two results columns show the accuracies when we train on one dataset and infer on the other (i.e., we train on UNSW-NB15 data and infer on the test set of the CIC-IDS-2017 dataset). Note, the accuracies do not match those from Table 5 since we use a balanced testing dataset with 50% benign and 50% attack packets. We do not find these results surprising since distribution shift is a well-known phenomenon in the vision domain Taori et al. (2020); Kulinski & Inouye (2022) with similar results in network traffic data Bierbrauer et al. (2022). However, it does highlight the need for the safeguarding of any cybersecurity ML models.

Table 8 provides the breakdown for packet category accuracy. We note two opposite issues when inferring on the CIC-IDS-2017 and UNSW-NB15 datasets when trained on data from the other. Most of the CIC-IDS-2017 packets are labeled as malicious when inferred on models trained on UNSW-NB15 data. Thus, the vast majority of benign packets are erroneously labeled as attack. Conversely, most of the UNSW-NB15 packets are labeled as benign when inferred on models trained on CIC-IDS-2017. We again emphasize the drastic shift in most common protocols between the creation of the two datasets (Table 2). By 2017, most websites shifted to encrypted protocols such as HTTPS. Thus, very few of the benign examples in UNSW-NB15 are encrypted compared to the CIC-IDS-2017 dataset. This appears to have caused issues during inference for the CIC-IDS-2017 benign packets when out-of-distribution.

Table 8: *Out-of-distribution category-wise accuracy.* When inferring on the other dataset, we note significant decreases in accuracy. In particular, the models fail to classify correctly the benign packets in CIC-IDS-2017 and attack ones in UNSW-NB15.

CIC-IDS-2017		UNSW-NB15	
Category	Accuracy (\uparrow)	Category	Accuracy (\uparrow)
Benign	0.1781 (± 0.0112)	Benign	0.9579 (± 0.0278)
Bot	0.7925 (± 0.1436)	Analysis	0.1056 (± 0.1545)
DDoS	0.7155 (± 0.0309)	Backdoor	0.0541 (± 0.0488)
DoSGoldenEye	0.6523 (± 0.0469)	DoS	0.0230 (± 0.0244)
DoSHulk	0.6604 (± 0.1084)	Exploits	0.0210 (± 0.0160)
DoSSlowhttptest	0.8478 (± 0.0501)	Fuzzers	0.0160 (± 0.0099)
DoSSlowloris	0.9224 (± 0.0720)	Generic	0.0087 (± 0.0062)
FTP-Patator	0.4878 (± 0.0783)	Reconnaissance	0.0363 (± 0.0525)
Heartbleed	0.7306 (± 0.0204)	Shellcode	0.0031 (± 0.0068)
Infiltration	0.8703 (± 0.2108)	Worms	0.0070 (± 0.0070)
PortScan	0.6378 (± 0.1356)		
SSH-Patator	0.0005 (± 0.0011)		
WebAttack-BruteForce	0.5268 (± 0.0735)		
WebAttack-SQLInjection	0.4800 (± 0.0588)		
WebAttack-XSS	0.5006 (± 0.0100)		

6.4 Detecting Out-of-Distribution Inputs

We evaluate our model safeguards with three metrics: area under the receiver operating characteristic (AU ROC), true positive rate at a true negative rate of 95% (TPR (TNR 95%)), and true positive rate at a true negative rate of 85% (TPR (TNR 85%)) (Tables 9 and 10). We first calculate the Mahalanobis distance and normalizing flow loss for each inference packet. Note, that in both cases, higher distances or losses indicate an inference input that we deem farther from the target in-distribution data. For AU ROC, we order the inputs by increasing score (farther from the target distribution) with binary labels of one corresponding to the out-of-distribution classes. For the TPR (TNR) scores, we adopt a hard threshold for what constitutes in- and out-of-distribution. Under these metrics, we allow ourselves to dispose of 5% and 15% of the in-distribution inference data (i.e., achieve a true negative rate of 95% and 85%). We then take the corresponding distance at those two thresholds and assign any inputs with scores greater than that as out-of-distribution. Our true positive rate thus refers to the proportion of out-of-distribution data that had a measured distance greater than those values. For all three metrics, higher values are better. We note that for some networks, losing nearly 5% of in-distribution data (including benign data), may not be tolerable. One can adapt these thresholds to the pain tolerances of the network users. For the maximum softmax probability baseline, we order the inputs by the output value for the predicted class.

Table 9 shows the results for detecting out-of-distribution inputs using maximum softmax probability, Gaussian kernel density, and normalizing flows. The final rows within each subproblem show the results using SAFE-NID. These tables contain the highest performing extracted layers for each network architecture. A complete set of results can be found in the Appendix. We note high AU ROC scores for Gaussian kernel density and normalizing flows with optimal values between 0.9668 and 0.9950 for each withheld attack class. Note, these results are with a high class imbalance with many more benign samples than attack ones. Therefore, our imbalance of in-distribution to out-of-distribution samples for FTP-Patator, Infiltration, and SSH-Patator is 1623.9 : 1, 1953.5 : 1, and 645.8 : 1, respectively.

Most of the FTP-Patator and SSH-Patator payloads are simply zeros causing a large amount of redundancy in the feature space before the dense layers for the CNN (with average payload lengths of 44.06 (± 14.41), 307.56 (± 454.89), respectively). Perhaps unsurprisingly then, we note that it is better to extract features from the FNN for FTP-Patator and SSH-Patator than the CNN, although best yet to use the transformer architecture. We note the opposite trend for the CNN and FNN for Infiltration, which has an average payload length of 908.23 (± 551.25). The transformer architecture with both Gaussian kernel estimation and

Table 9: *Out-of-distribution detection results for zero-day exploits.* We achieve high AU ROC scores using the transformer architecture with a normalizing flows safeguard. Additional results in the Appendix.

Architecture	Layer	AU ROC (\uparrow)	TPR (TNR 95%) (\uparrow)	TPR (TNR 85%) (\uparrow)
FTP-Patator				
Maximum Softmax Probability				
CNN	Output	0.7931 (± 0.1534)	15.45% ($\pm 15.48\%$)	64.19% ($\pm 37.97\%$)
FNN	Output	0.6149 (± 0.0893)	0.12% ($\pm 0.37\%$)	9.89% ($\pm 11.39\%$)
Transformer	Output	0.4034 (± 0.2232)	0.21% ($\pm 0.54\%$)	14.61% ($\pm 13.37\%$)
Gaussian Kernel Density				
CNN	dense3	0.8297 (± 0.1157)	26.47% ($\pm 24.29\%$)	58.92% ($\pm 29.76\%$)
FNN	linear3	0.8473 (± 0.0259)	3.55% ($\pm 8.46\%$)	59.20% ($\pm 19.59\%$)
Transformer	linear1	0.9950 (± 0.0018)	100.00% ($\pm 0.00\%$)	100.00% ($\pm 0.00\%$)
Normalizing Flows				
CNN	dense2	0.8058 (± 0.1224)	4.53% ($\pm 5.42\%$)	48.96% ($\pm 31.88\%$)
FNN	linear3	0.8640 (± 0.0152)	6.88% ($\pm 12.78\%$)	61.51% ($\pm 14.50\%$)
Transformer	linear1	0.9936 (± 0.0014)	100.00% ($\pm 0.00\%$)	100.00% ($\pm 0.00\%$)
Infiltration				
Maximum Softmax Probability				
CNN	Output	0.6611 (± 0.1897)	26.16% ($\pm 20.54\%$)	53.95% ($\pm 23.83\%$)
FNN	Output	0.7676 (± 0.0567)	12.39% ($\pm 11.58\%$)	47.52% ($\pm 16.07\%$)
Transformer	Output	0.3167 (± 0.2055)	3.37% ($\pm 6.66\%$)	9.56% ($\pm 19.88\%$)
Gaussian Kernel Density				
CNN	dense3	0.9622 (± 0.0217)	79.11% ($\pm 15.19\%$)	96.29% ($\pm 2.99\%$)
FNN	linear4	0.8748 (± 0.0405)	41.70% ($\pm 15.30\%$)	61.45% ($\pm 16.68\%$)
Transformer	linear2	0.9552 (± 0.0318)	71.26% ($\pm 28.99\%$)	94.93% ($\pm 13.73\%$)
Normalizing Flows				
CNN	dense3	0.9638 (± 0.0153)	77.08% ($\pm 17.79\%$)	98.22% ($\pm 0.78\%$)
FNN	linear4	0.8820 (± 0.0300)	30.41% ($\pm 15.53\%$)	67.72% ($\pm 14.22\%$)
Transformer	linear2	0.9668 (± 0.0182)	77.89% ($\pm 23.41\%$)	99.18% ($\pm 2.25\%$)
SSH-Patator				
Maximum Softmax Probability				
CNN	Output	0.7821 (± 0.0511)	17.19% ($\pm 10.35\%$)	66.61% ($\pm 15.90\%$)
FNN	Output	0.6568 (± 0.0727)	1.53% ($\pm 0.87\%$)	31.75% ($\pm 11.64\%$)
Transformer	Output	0.1924 (± 0.1103)	0.20% ($\pm 0.35\%$)	2.24% ($\pm 2.45\%$)
Gaussian Kernel Density				
CNN	dense3	0.7254 (± 0.1269)	23.72% ($\pm 10.20\%$)	46.08% ($\pm 12.92\%$)
FNN	linear3	0.8967 (± 0.0100)	26.14% ($\pm 4.38\%$)	77.99% ($\pm 5.01\%$)
Transformer	linear1	0.9821 (± 0.0037)	94.90% ($\pm 2.61\%$)	99.05% ($\pm 1.09\%$)
Normalizing Flows				
CNN	dense2	0.7086 (± 0.1117)	11.23% ($\pm 4.00\%$)	39.33% ($\pm 13.11\%$)
FNN	linear3	0.9119 (± 0.0062)	30.81% ($\pm 3.04\%$)	84.98% ($\pm 3.99\%$)
Transformer	linear1	0.9901 (± 0.0007)	100.00% ($\pm 0.00\%$)	100.00% ($\pm 0.00\%$)

Table 10: *Out-of-distribution detection results for concept drift.* We achieve slightly lower AU ROC scores on the out-of-distribution CIC-IDS-2017 and UNSW-NB15 datasets, in part because we have higher overall accuracy (more overlap of distributions). Additional results can be found in the Appendix.

Architecture	Layer	AU ROC (\uparrow)	TPR (TNR 95%) (\uparrow)	TPR (TNR 85%) (\uparrow)
CIC-IDS-2017				
Maximum Softmax Probability				
CNN	Output	0.7818 (± 0.0254)	36.43% ($\pm 4.48\%$)	63.02% ($\pm 4.47\%$)
FNN	Output	0.7187 (± 0.0289)	23.25% ($\pm 3.21\%$)	50.69% ($\pm 3.18\%$)
Transformer	Output	0.6731 (± 0.0737)	23.41% ($\pm 8.93\%$)	38.37% ($\pm 14.72\%$)
Gaussian Kernel Density				
CNN	dense2	0.8981 (± 0.0170)	55.39% ($\pm 5.94\%$)	77.16% ($\pm 5.28\%$)
FNN	linear4	0.7343 (± 0.0597)	34.96% ($\pm 4.17\%$)	45.25% ($\pm 5.73\%$)
Transformer	linear2	0.8753 (± 0.0552)	58.05% ($\pm 11.80\%$)	73.07% ($\pm 11.11\%$)
Normalizing Flows				
CNN	dense2	0.8960 (± 0.0151)	52.88% ($\pm 3.63\%$)	76.19% ($\pm 5.61\%$)
FNN	linear3	0.8130 (± 0.0226)	27.27% ($\pm 6.58\%$)	50.10% ($\pm 2.03\%$)
Transformer	linear2	0.9259 (± 0.0046)	73.01% ($\pm 2.10\%$)	81.55% ($\pm 1.11\%$)
UNSW-NB15				
Maximum Softmax Probability				
CNN	Output	0.5590 (± 0.0063)	34.87% ($\pm 2.45\%$)	44.47% ($\pm 3.19\%$)
FNN	Output	0.7349 (± 0.0417)	25.98% ($\pm 2.04\%$)	43.43% ($\pm 5.41\%$)
Transformer	Output	0.6255 (± 0.1033)	24.21% ($\pm 11.24\%$)	40.14% ($\pm 14.78\%$)
Gaussian Kernel Density				
CNN	dense2	0.9263 (± 0.0139)	61.61% ($\pm 4.44\%$)	83.14% ($\pm 4.44\%$)
FNN	linear4	0.9079 (± 0.0109)	56.06% ($\pm 5.32\%$)	80.31% ($\pm 2.68\%$)
Transformer	linear2	0.9636 (± 0.0126)	79.23% ($\pm 8.25\%$)	95.73% ($\pm 2.68\%$)
Normalizing Flows				
CNN	dense2	0.9263 (± 0.0063)	55.98% ($\pm 1.29\%$)	83.03% ($\pm 2.72\%$)
FNN	linear4	0.8963 (± 0.0112)	45.10% ($\pm 8.46\%$)	77.66% ($\pm 3.28\%$)
Transformer	linear3	0.9583 (± 0.0090)	77.71% ($\pm 4.49\%$)	94.44% ($\pm 2.24\%$)

normalizing flows performs well for all OOD inputs. The maximum softmax probability method performs poorly across the board with no high scores for any attack class.

Table 10 shows our results for detecting the CIC-IDS-2017 and UNSW-NB15 datasets as OOD. The last rows for each dataset grouping contain results for SAFE-NID. We note lower AU ROC scores than for the withheld CIC attack classes. We attribute this, in part, to the higher accuracies when these datasets are left out of training. Although there was a distribution shift between the datasets, we still achieve higher accuracies. Thus some of these packets may not truly be out-of-distribution. The transformer architecture outperforms the others during this task as well, with the normalizing flows and Gaussian kernel density safeguards performing best on CIC-IDS-2017 and UNSW-NB15, respectively.

6.5 Latency Analysis

The CNN, FNN, and Transformer require 172.58 μ sec, 144.60 μ sec, and 2418.02 μ sec to process one packet, respectively. We note that the transformer architecture takes significantly more time during inference than the other two networks, with the fully-connected neural network requiring the least overhead. We note a trade off between latency and accuracy. A real-world implementation of our system might employ different classifiers on network traffic depending on the perceived importance of the receiving devices.

We show the amount of processing time per packet in Table 11 for our model safeguards. The feature dimensionality depends on the extracted layer from our DNNs. Both safeguard implementations take on

Table 11: *Model safeguard latency.* Average number of microseconds to process a packet for the model safeguards

Gaussian Kernel Density		Normalizing Flows	
Feature Dimensionality	Latency (\downarrow)	Feature Dimensionality	Latency (\downarrow)
256	59.60 μ sec	256	40.45 μ sec
128	29.59 μ sec	128	35.37 μ sec
64	23.96 μ sec	64	32.95 μ sec

the order of tens of microseconds to process each packet. The timing analysis was run on an AMD Ryzen Threadripper PRO 5965WX 24-Cores processor at 1.8 GHz for Gaussian kernel density and on an NVIDIA RTX A6000 for normalizing flows. The average time from packet to flow termination in CIC-IDS-2017 is 31.29 seconds (median time is 11.56 seconds). Thus, packet-level detection has very low latency when compared with flow-level detection. The maximum softmax probability system requires limited additional overhead cost as it merely takes the output from the last layer.

7 Conclusions

There are several challenges to using deep learning methods for classifying packets as benign or malicious, including the creation of such datasets from existing publicly available flow-level datasets. We identify these challenges and publish the data for future research. We demonstrate that DNNs trained for network intrusion detection on these datasets achieve high accuracy (over 99%) on in-distribution inputs, but fail drastically (accuracy below 1%) in the presence of zero-day attacks, with severe degradation in the presence of concept drift. We address this robustness challenge by constructing a model safeguard that uses the prediction of the DNN classifier and its internal features to quantify the uncertainty in predictions made by the classifier and to detect novel and out-of-distribution inputs. We propose SAFE-NID: a lightweight encoder-only transformer architecture for the packet classification task with a normalizing flows model safeguard. The normalizing flows model learns the distribution of internal features of the neural networks on in-distribution input samples. We then compare features from new inputs to this distribution to identify out-of-distribution examples. Our proposed solution achieves an AU ROC score of over 0.9668 in detecting novel inputs from withheld attack classes.

References

- Hyrum S Anderson and Phil Roth. Ember: an open dataset for training static pe malware machine learning models. *arXiv preprint arXiv:1804.04637*, 2018.
- Lynton Ardizzone, Till Bungert, Felix Draxler, Ullrich Köthe, Jakob Kruse, Robert Schmier, and Peter Sorrenson. Framework for Easily Invertible Architectures (FrEIA), 2018-2022. URL <https://github.com/vislearn/FrEIA>.
- David A Bierbrauer, Michael De Lucia, Krishna Reddy, Paul Maxwell, and Nathaniel D Bastian. Transfer learning for raw network traffic detection. *Expert Systems with Applications*, 211(118641):1, 2022.
- Alice Bizzarri, Brian Jalaian, Fabrizio Riguzzi, and Nathaniel D Bastian. A neuro-symbolic artificial intelligence network intrusion detection system. In *2024 33rd International Conference on Computer Communications and Networks (ICCCN)*, pp. 1–9. IEEE, 2024.
- Saikiran Bulusu, Bhavya Kailkhura, Bo Li, Pramod K Varshney, and Dawn Song. Anomalous example detection in deep learning: A survey. *IEEE Access*, 8:132330–132347, 2020.
- Muhammad Arif Butt, Zarafshan Ajmal, Zafar Iqbal Khan, Muhammad Idrees, and Yasir Javed. An in-depth survey of bypassing buffer overflow mitigation techniques. *Applied Sciences*, 12(13):6702, 2022.

- Feiyang Cai and Xenofon Koutsoukos. Real-time out-of-distribution detection in learning-enabled cyber-physical systems. In *2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPS)*, pp. 174–183. IEEE, 2020.
- Benoit Claise. Cisco systems netflow services export version 9. Technical report, Cisco, 2004.
- Michael J De Lucia, Paul E Maxwell, Nathaniel D Bastian, Ananthram Swami, Brian Jalaian, and Nandi Leslie. Machine learning raw network traffic detection. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications III*, volume 11746, pp. 185–194. SPIE, 2021.
- Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- David Dittrich. The dos project’s ‘trinoo’ distributed denial of service attack tool. 1999.
- Gerard Draper-Gil, Arash Habibi Lashkari, Mohammad Saiful Islam Mamun, and Ali A Ghorbani. Characterization of encrypted and vpn traffic using time-related. In *Proceedings of the 2nd international conference on information systems security and privacy (ICISSP)*, pp. 407–414, 2016.
- Gints Engelen, Vera Rimmer, and Wouter Joosen. Troubleshooting an intrusion detection dataset: the cicids2017 case study. In *2021 IEEE Security and Privacy Workshops (SPW)*, pp. 7–12. IEEE, 2021.
- Yasir Ali Farrukh, Irfan Khan, Syed Wali, David Bierbrauer, John A. Pavlik, and Nathaniel D. Bastian. Payload-Byte: A Tool for Extracting and Labeling Packet Capture Files of Modern Network Intrusion Detection Datasets. *Proceedings of the 9th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT2022)*, 12 2022.
- Mohamed Amine Ferrag, Leandros Maglaras, Sotiris Moschoyiannis, and Helge Janicke. Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. *Journal of Information Security and Applications*, 50:102419, 2020.
- Chuanxing Geng, Sheng-jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3614–3631, 2020.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pp. 1321–1330. PMLR, 2017.
- Dongqi Han, Zhiliang Wang, Wenqi Chen, Kai Wang, Rui Yu, Su Wang, Han Zhang, Zhihua Wang, Minghui Jin, Jiahai Yang, et al. Anomaly detection in the open world: Normality shift detection, explanation, and adaptation. In *30th Annual Network and Distributed System Security Symposium (NDSS)*, 2023.
- Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017.
- Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *International Conference on Learning Representations*, 2019.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456. PMLR, 2015.
- Van Jacobson. Tcpcdump. *ftp://ftp. ee. lbl. gov*, 1989.
- Uyeong Jang, Susmit Jha, and Somesh Jha. On the need for topology-aware generative models for manifold-based defenses. *8th International Conference on Learning Representations (ICLR) 2020*, 2020. URL <https://par.nsf.gov/biblio/10181037>.

- Susmit Jha, Sunny Raj, Steven Fernandes, Sumit K Jha, Somesh Jha, Brian Jalaian, Gunjan Verma, and Ananthram Swami. Attribution-based confidence metric for deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 11826–11837, 2019.
- Heinrich Jiang, Been Kim, Melody Guan, and Maya Gupta. To trust or not to trust a classifier. In *Advances in neural information processing systems*, pp. 5541–5552, 2018.
- Ramneet Kaur, Susmit Jha, Anirban Roy, Sangdon Park, Oleg Sokolsky, and Insup Lee. Detecting oods as datapoints with high uncertainty. *arXiv preprint arXiv:2108.06380*, 2021.
- Jiyeon Kim, Yulim Shin, and Eunjung Choi. An intrusion detection model based on a convolutional neural network. *Journal of Multimedia Information System*, 6(4):165–172, 2019.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.
- Ivan Kobyzev, Simon JD Prince, and Marcus A Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3964–3979, 2020.
- Nickolaos Koroniotis, Nour Moustafa, Elena Sitnikova, and Benjamin Turnbull. Towards the development of realistic botnet dataset in the internet of things for network forensic analytics: Bot-iot dataset. *Future Generation Computer Systems*, 100:779–796, 2019.
- Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*, 2018.
- Sean Kulinski and David I Inouye. Towards explaining image-based distribution shifts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4788–4792, 2022.
- Arash Habibi Lashkari, Gerard Draper-Gil, Mohammad Saiful Islam Mamun, Ali A Ghorbani, et al. Characterization of tor traffic using time based features. In *ICISSp*, pp. 253–262, 2017.
- Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. *arXiv preprint arXiv:1711.09325*, 2017.
- Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018.
- Joffrey L Leevy and Taghi M Khoshgoftaar. A survey and analysis of intrusion detection models based on cse-cic-ids2018 big data. *Journal of Big Data*, 7(1):1–19, 2020.
- letsencrypt. Let’s encrypt stats. <https://letsencrypt.org/stats/>. Accessed: December 2, 2022.
- Bingdong Li, Jeff Springer, George Bebis, and Mehmet Hadi Gunes. A survey of network flow applications. *Journal of Network and Computer Applications*, 36(2):567–581, 2013.
- Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv preprint arXiv:1706.02690*, 2017.
- Xiang Ling, Lingfei Wu, Wei Deng, Zhenqing Qu, Jianguy Zhang, Sheng Zhang, Tengfei Ma, Bin Wang, Chunming Wu, and Shouling Ji. Malgraph: Hierarchical graph neural networks for robust windows malware detection. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*, pp. 1998–2007. IEEE, 2022.

- Lisa Liu, Gints Engelen, Timothy Lynar, Daryl Essam, and Wouter Joosen. Error prevalence in nids datasets: A case study on cic-ids-2017 and cse-cic-ids-2018. In *2022 IEEE Conference on Communications and Network Security (CNS)*, pp. 254–262. IEEE, 2022.
- Weitang Liu, Xiaoyun Wang, John D Owens, and Yixuan Li. Energy-based out-of-distribution detection. *arXiv preprint arXiv:2010.03759*, 2020.
- Andrew L Maas, Awni Y Hannun, Andrew Y Ng, et al. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, pp. 3. Atlanta, Georgia, USA, 2013.
- David Macedo, Tsang Ing Ren, Cleber Zanchettin, Adriano L. I. Oliveira, and Teresa Ludermir. Entropic out-of-distribution detection, 2021.
- Prasanta Chandra Mahalanobis. On the generalized distance in statistics. National Institute of Science of India, 1936.
- Paul Maxwell, Elie Alhajjar, and Nathaniel D Bastian. Intelligent feature engineering for cybersecurity. In *2019 IEEE International Conference on Big Data (Big Data)*, pp. 5005–5011. IEEE, 2019.
- Alexander Meinke and Matthias Hein. Towards neural networks that provably know when they don’t know. *arXiv preprint arXiv:1909.12180*, 2019.
- Tajuddin Manhar Mohammed, Lakshmanan Nataraj, Satish Chikkagoudar, Shivkumar Chandrasekaran, and BS Manjunath. Malware detection using frequency domain-based image visualization and deep learning. *arXiv preprint arXiv:2101.10578*, 2021.
- Nour Moustafa. A new distributed architecture for evaluating ai-based security systems at the edge: Network ton_iiot datasets. *Sustainable Cities and Society*, 72:102994, 2021.
- Nour Moustafa and Jill Slay. Unsw-nb15: a comprehensive data set for network intrusion detection systems (unsw-nb15 network data set). In *2015 military communications and information systems conference (MilCIS)*, pp. 1–6. IEEE, 2015.
- George Papamakarios, Eric T Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *J. Mach. Learn. Res.*, 22(57):1–64, 2021.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Zachariah Pelletier and Munther Abualkibash. Evaluating the cic ids-2017 dataset using machine learning methods and creating multiple predictive models in the statistical computing language r. *Science*, 5(2): 187–191, 2020.
- Stephan Rabanser, Stephan Günnemann, and Zachary Lipton. Failing loudly: An empirical study of methods for detecting dataset shift. *Advances in Neural Information Processing Systems*, 32, 2019.
- Edward Raff, Jon Barker, Jared Sylvester, Robert Brandon, Bryan Catanzaro, and Charles K Nicholas. Malware detection by eating a whole exe. In *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Pooja Rani and Nathaniel D Bastian. Zero-shot learning for raw network traffic detection. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications VI*, volume 13051, pp. 441–452. SPIE, 2024.
- Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*, 2019.

- Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. URL <https://arxiv.org/abs/1908.10084>.
- Jie Ren, Peter J Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark Depristo, Joshua Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. In *Advances in Neural Information Processing Systems*, pp. 14707–14718, 2019.
- Michelle Robbins. Effective july 2018, google’s chrome browser will mark non-https sites as ‘not secure’. <https://searchengineland.com/effective-july-2018-googles-chrome-browser-will-mark-non-https-sites-as-not-secure-291623>, Aug 2021. Accessed: December 2, 2022.
- Mohanad Sarhan, Siamak Layeghy, Nour Moustafa, and Marius Portmann. Netflow datasets for machine learning-based network intrusion detection systems. In *Big Data Technologies and Applications*, pp. 117–135. Springer, 2020.
- Joshua Saxe and Konstantin Berlin. Deep neural network based malware detection using two dimensional binary program features. In *2015 10th international conference on malicious and unwanted software (MALWARE)*, pp. 11–20. IEEE, 2015.
- Christoph L Schuba, Ivan V Krsul, Markus G Kuhn, Eugene H Spafford, Aurobindo Sundaram, and Diego Zamboni. Analysis of a denial of service attack on tcp. In *Proceedings. 1997 IEEE Symposium on Security and Privacy (Cat. No. 97CB36097)*, pp. 208–223. IEEE, 1997.
- Huasong Shan, Qingyang Wang, and Qiben Yan. Very short intermittent ddos attacks in an unsaturated system. In *International Conference on Security and Privacy in Communication Systems*, pp. 45–66. Springer, 2017.
- Iman Sharafaldin, Amirhossein Gharib, Arash Habibi Lashkari, and Ali A Ghorbani. Towards a reliable intrusion detection benchmark dataset. *Software Networking*, 2018(1):177–200, 2018a.
- Iman Sharafaldin, Arash Habibi Lashkari, and Ali A Ghorbani. Toward generating a new intrusion detection dataset and intrusion traffic characterization. *ICISSp*, 1:108–116, 2018b.
- Zheyang Shen, Jiashuo Liu, Yue He, Xingxuan Zhang, Renzhe Xu, Han Yu, and Peng Cui. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:2108.13624*, 2021.
- Alex Shenfield, David Day, and Aladdin Ayeshe. Intelligent intrusion detection systems using artificial neural networks. *Ict Express*, 4(2):95–99, 2018.
- Xinying Song, Alex Salcianu, Yang Song, Dave Dopson, and Denny Zhou. Fast wordpiece tokenization. *arXiv preprint arXiv:2012.15524*, 2020.
- Peter Sorrenson, Carsten Rother, and Ullrich Köthe. Disentanglement by nonlinear ica with general incompressible-flow networks (gin). *arXiv preprint arXiv:2001.04872*, 2020.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1): 1929–1958, 2014.
- Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. *Advances in Neural Information Processing Systems*, 33:18583–18599, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- R Vinayakumar and KP Soman. Deepmalnet: evaluating shallow and deep networks for static pe malware detection. *ICT express*, 4(4):255–258, 2018.

- Ravi Vinayakumar, Mamoun Alazab, KP Soman, Prabakaran Poornachandran, Ameer Al-Nemrat, and Sitalakshmi Venkatraman. Deep learning approach for intelligent intrusion detection system. *Ieee Access*, 7:41525–41550, 2019.
- Zhanyi Wang. The applications of deep learning on traffic identification. *BlackHat USA*, 24(11):1–10, 2015.
- Shreekh Wankhede and Deepak Kshirsagar. Dos attack detection using machine learning and neural network. In *2018 Fourth International Conference on Computing Communication Control and Automation (ICCUBEA)*, pp. 1–5. IEEE, 2018.
- Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021.
- Arif Yulianto, Parman Sukarno, and Novian Anggis Suwastika. Improving adaboost-based intrusion detection system (ids) performance on cic ids 2017 dataset. In *Journal of Physics: Conference Series*, volume 1192, pp. 012018. IOP Publishing, 2019.
- zeek. The zeek network security monitor. URL <https://zeek.org/>.
- Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. *Advances in Neural Information Processing Systems*, 34:29820–29834, 2021.

A Complete Out-of-Distribution Results

Table 12: *FTP-Patator out-of-distribution results*. The transformer architecture with either the Gaussian kernel density or normalizing flows safeguard achieves very high AU ROC scores > 0.9936 . The FNN model performs better than the CNN on this data, but still is significantly worse than when using a transformer architecture.

FTP-Patator				
Maximum Softmax Probability				
CNN	Output	0.7931 (± 0.1534)	15.45% ($\pm 15.48\%$)	64.19% ($\pm 37.97\%$)
FNN	Output	0.6149 (± 0.0893)	0.12% ($\pm 0.37\%$)	9.89% ($\pm 11.39\%$)
Transformer	Output	0.4034 (± 0.2232)	0.21% ($\pm 0.54\%$)	14.61% ($\pm 13.37\%$)
Gaussian Kernel Density				
CNN	dense1	0.7563 (± 0.2059)	0.53% ($\pm 0.51\%$)	57.66% ($\pm 26.93\%$)
CNN	dense2	0.8207 (± 0.0684)	4.09% ($\pm 2.69\%$)	51.64% ($\pm 30.72\%$)
CNN	dense3	0.8297 (± 0.1157)	26.47% ($\pm 24.29\%$)	58.92% ($\pm 29.76\%$)
FNN	linear3	0.8473 (± 0.0259)	3.55% ($\pm 8.46\%$)	59.20% ($\pm 19.59\%$)
FNN	linear4	0.7511 (± 0.0406)	2.56% ($\pm 4.37\%$)	19.70% ($\pm 15.77\%$)
FNN	linear5	0.7585 (± 0.0591)	0.37% ($\pm 0.47\%$)	24.53% ($\pm 17.33\%$)
Transformer	linear1	0.9950 (± 0.0018)	100.00% ($\pm 0.00\%$)	100.00% ($\pm 0.00\%$)
Transformer	linear2	0.9717 (± 0.0083)	89.63% ($\pm 16.73\%$)	100.00% ($\pm 0.00\%$)
Transformer	linear3	0.9255 (± 0.0406)	44.22% ($\pm 33.00\%$)	89.51% ($\pm 16.49\%$)
Normalizing Flows				
CNN	dense1	0.7442 (± 0.2052)	0.50% ($\pm 0.72\%$)	51.70% ($\pm 35.34\%$)
CNN	dense2	0.8058 (± 0.1224)	4.53% ($\pm 5.42\%$)	48.96% ($\pm 31.88\%$)
CNN	dense3	0.7998 (± 0.1178)	15.46% ($\pm 14.41\%$)	56.90% ($\pm 29.89\%$)
FNN	linear3	0.8640 (± 0.0152)	6.88% ($\pm 12.78\%$)	61.51% ($\pm 14.50\%$)
FNN	linear4	0.7917 (± 0.0389)	3.37% ($\pm 8.52\%$)	29.55% ($\pm 14.60\%$)
FNN	linear5	0.7732 (± 0.0580)	7.25% ($\pm 7.74\%$)	38.37% ($\pm 13.60\%$)
Transformer	linear1	0.9936 (± 0.0014)	100.00% ($\pm 0.00\%$)	100.00% ($\pm 0.00\%$)
Transformer	linear2	0.9726 (± 0.0064)	92.71% ($\pm 13.52\%$)	100.00% ($\pm 0.00\%$)
Transformer	linear3	0.9480 (± 0.0199)	58.83% ($\pm 26.97\%$)	97.50% ($\pm 4.21\%$)

Table 13: *Infiltration out-of-distribution results.* The transformer with normalizing flows safeguard performs worse here than with the other two zero-day exploits. However, is still outperforms baseline methods on AU ROC. The CNN architecture performs well here compared to FTP-Patator and SSH-Patator results.

Architecture	Layer	AU ROC (\uparrow)	TPR (TNR 95%) (\uparrow)	TPR (TNR 85%) (\uparrow)
Infiltration				
Maximum Softmax Probability				
CNN	Output	0.6611 (± 0.1897)	26.16% ($\pm 20.54\%$)	53.95% ($\pm 23.83\%$)
FNN	Output	0.7676 (± 0.0567)	12.39% ($\pm 11.58\%$)	47.52% ($\pm 16.07\%$)
Transformer	Output	0.3167 (± 0.2055)	3.37% ($\pm 6.66\%$)	9.56% ($\pm 19.88\%$)
Gaussian Kernel Density				
CNN	dense1	0.9581 (± 0.0062)	83.74% ($\pm 10.23\%$)	98.70% ($\pm 0.53\%$)
CNN	dense2	0.9511 (± 0.0218)	75.73% ($\pm 19.16\%$)	96.55% ($\pm 2.93\%$)
CNN	dense3	0.9622 (± 0.0217)	79.11% ($\pm 15.19\%$)	96.29% ($\pm 2.99\%$)
FNN	linear3	0.8740 (± 0.0337)	27.72% ($\pm 13.36\%$)	67.78% ($\pm 13.74\%$)
FNN	linear4	0.8748 (± 0.0405)	41.70% ($\pm 15.30\%$)	61.45% ($\pm 16.68\%$)
FNN	linear5	0.8345 (± 0.0439)	29.58% ($\pm 19.39\%$)	47.85% ($\pm 18.80\%$)
Transformer	linear1	0.9455 (± 0.0248)	62.58% ($\pm 20.95\%$)	93.79% ($\pm 10.34\%$)
Transformer	linear2	0.9552 (± 0.0318)	71.26% ($\pm 28.99\%$)	94.93% ($\pm 13.73\%$)
Transformer	linear3	0.8996 (± 0.0580)	35.48% ($\pm 28.28\%$)	75.45% ($\pm 25.56\%$)
Normalizing Flows				
CNN	dense1	0.9464 (± 0.0061)	67.72% ($\pm 15.48\%$)	98.90% ($\pm 0.05\%$)
CNN	dense2	0.9519 (± 0.0057)	72.07% ($\pm 18.42\%$)	98.91% ($\pm 0.06\%$)
CNN	dense3	0.9638 (± 0.0153)	77.08% ($\pm 17.79\%$)	98.22% ($\pm 0.78\%$)
FNN	linear3	0.8655 (± 0.0115)	5.89% ($\pm 4.88\%$)	66.73% ($\pm 11.63\%$)
FNN	linear4	0.8820 (± 0.0300)	30.41% ($\pm 15.53\%$)	67.72% ($\pm 14.22\%$)
FNN	linear5	0.8530 (± 0.0385)	29.11% ($\pm 18.50\%$)	59.60% ($\pm 14.84\%$)
Transformer	linear1	0.9555 (± 0.0146)	63.20% ($\pm 24.30\%$)	99.64% ($\pm 0.62\%$)
Transformer	linear2	0.9668 (± 0.0182)	77.89% ($\pm 23.41\%$)	99.18% ($\pm 2.25\%$)
Transformer	linear3	0.9361 (± 0.0455)	56.13% ($\pm 31.83\%$)	89.54% ($\pm 15.68\%$)

Table 14: *SSH-Patator out-of-distribution results*. The transformer classifier with normalizing flows safeguard outperforms all existing methods with a very high AU ROC of 0.9901.

Architecture	Layer	AU ROC (\uparrow)	TPR (TNR 95%) (\uparrow)	TPR (TNR 85%) (\uparrow)
SSH-Patator				
Maximum Softmax Probability				
CNN	Output	0.7821 (± 0.0511)	17.19% ($\pm 10.35\%$)	66.61% ($\pm 15.90\%$)
FNN	Output	0.6568 (± 0.0727)	1.53% ($\pm 0.87\%$)	31.75% ($\pm 11.64\%$)
Transformer	Output	0.1924 (± 0.1103)	0.20% ($\pm 0.35\%$)	2.24% ($\pm 2.45\%$)
Gaussian Kernel Density				
CNN	dense1	0.6158 (± 0.1096)	7.14% ($\pm 1.62\%$)	27.81% ($\pm 13.33\%$)
CNN	dense2	0.6993 (± 0.1041)	11.45% ($\pm 7.93\%$)	39.42% ($\pm 11.65\%$)
CNN	dense3	0.7254 (± 0.1269)	23.72% ($\pm 10.20\%$)	46.08% ($\pm 12.92\%$)
FNN	linear3	0.8967 (± 0.0100)	26.14% ($\pm 4.38\%$)	77.99% ($\pm 5.01\%$)
FNN	linear4	0.8381 (± 0.0169)	19.34% ($\pm 3.42\%$)	52.81% ($\pm 6.62\%$)
FNN	linear5	0.8298 (± 0.0265)	18.36% ($\pm 3.01\%$)	56.14% ($\pm 7.42\%$)
Transformer	linear1	0.9821 (± 0.0037)	94.90% ($\pm 2.61\%$)	99.05% ($\pm 1.09\%$)
Transformer	linear2	0.9629 (± 0.0048)	76.63% ($\pm 9.36\%$)	100.00% ($\pm 0.00\%$)
Transformer	linear3	0.9097 (± 0.0313)	27.03% ($\pm 14.36\%$)	86.99% ($\pm 12.27\%$)
Normalizing Flows				
CNN	dense1	0.5739 (± 0.1273)	6.90% ($\pm 0.51\%$)	18.94% ($\pm 6.43\%$)
CNN	dense2	0.7086 (± 0.1117)	11.23% ($\pm 4.00\%$)	39.33% ($\pm 13.11\%$)
CNN	dense3	0.7010 (± 0.1336)	21.97% ($\pm 11.52\%$)	44.46% ($\pm 15.20\%$)
FNN	linear3	0.9119 (± 0.0062)	30.81% ($\pm 3.04\%$)	84.98% ($\pm 3.99\%$)
FNN	linear4	0.8810 (± 0.0144)	30.83% ($\pm 3.27\%$)	70.58% ($\pm 6.78\%$)
FNN	linear5	0.8596 (± 0.0312)	30.49% ($\pm 3.33\%$)	67.67% ($\pm 8.18\%$)
Transformer	linear1	0.9901 (± 0.0007)	100.00% ($\pm 0.00\%$)	100.00% ($\pm 0.00\%$)
Transformer	linear2	0.9668 (± 0.0062)	84.62% ($\pm 12.51\%$)	100.00% ($\pm 0.00\%$)
Transformer	linear3	0.9323 (± 0.0180)	45.45% ($\pm 13.92\%$)	92.96% ($\pm 6.04\%$)

Table 15: *CIC-IDS-2017 out-of-distribution results*. The highest performing combination for detecting the out-of-distribution samples is the the transformer discriminative classifier with the normalizing flows safeguard.

Architecture	Layer	AU ROC (\uparrow)	TPR (TNR 95%) (\uparrow)	TPR (TNR 85%) (\uparrow)
CIC-IDS-2017				
Maximum Softmax Probability				
CNN	Output	0.7818 (± 0.0254)	36.43% ($\pm 4.48\%$)	63.02% ($\pm 4.47\%$)
FNN	Output	0.7187 (± 0.0289)	23.25% ($\pm 3.21\%$)	50.69% ($\pm 3.18\%$)
Transformer	Output	0.6731 (± 0.0737)	23.41% ($\pm 8.93\%$)	38.37% ($\pm 14.72\%$)
Gaussian Kernel Density				
CNN	dense1	0.8814 (± 0.0144)	49.93% ($\pm 4.44\%$)	75.21% ($\pm 3.38\%$)
CNN	dense2	0.8981 (± 0.0170)	55.39% ($\pm 5.94\%$)	77.16% ($\pm 5.28\%$)
CNN	dense3	0.8580 (± 0.0394)	55.26% ($\pm 4.30\%$)	65.82% ($\pm 5.45\%$)
FNN	linear3	0.6986 (± 0.0464)	25.80% ($\pm 5.13\%$)	40.43% ($\pm 6.12\%$)
FNN	linear4	0.7343 (± 0.0597)	34.96% ($\pm 4.17\%$)	45.25% ($\pm 5.73\%$)
FNN	linear5	0.6801 (± 0.0520)	25.72% ($\pm 2.39\%$)	35.15% ($\pm 3.33\%$)
Transformer	linear1	0.8323 (± 0.0649)	50.61% ($\pm 8.24\%$)	64.29% ($\pm 10.40\%$)
Transformer	linear2	0.8753 (± 0.0552)	58.05% ($\pm 11.80\%$)	73.07% ($\pm 11.11\%$)
Transformer	linear3	0.7893 (± 0.0541)	42.14% ($\pm 6.95\%$)	53.56% ($\pm 9.77\%$)
Normalizing Flows				
CNN	dense1	0.8744 (± 0.0121)	30.50% ($\pm 3.95\%$)	71.26% ($\pm 5.71\%$)
CNN	dense2	0.8960 (± 0.0151)	52.88% ($\pm 3.63\%$)	76.19% ($\pm 5.61\%$)
CNN	dense3	0.8384 (± 0.0594)	53.77% ($\pm 5.61\%$)	67.10% ($\pm 8.80\%$)
FNN	linear3	0.8130 (± 0.0226)	27.27% ($\pm 6.58\%$)	50.10% ($\pm 2.03\%$)
FNN	linear4	0.7625 (± 0.0210)	35.61% ($\pm 1.12\%$)	48.22% ($\pm 2.08\%$)
FNN	linear5	0.7139 (± 0.0582)	31.64% ($\pm 2.98\%$)	42.69% ($\pm 5.90\%$)
Transformer	linear1	0.9011 (± 0.0115)	66.87% ($\pm 3.37\%$)	76.70% ($\pm 2.31\%$)
Transformer	linear2	0.9259 (± 0.0046)	73.01% ($\pm 2.10\%$)	81.55% ($\pm 1.11\%$)
Transformer	linear3	0.8945 (± 0.0085)	63.15% ($\pm 4.94\%$)	77.45% ($\pm 2.06\%$)

Table 16: *UNSW-NB15 out-of-distribution results*. Although transformer with normalizing flows performs well, the best performer comes from the transformer with the Gaussian kernel density safeguard.

Architecture	Layer	AU ROC (\uparrow)	TPR (TNR 95%) (\uparrow)	TPR (TNR 85%) (\uparrow)
UNSW-NB15				
Maximum Softmax Probability				
CNN	Output	0.5590 (± 0.0063)	34.87% ($\pm 2.45\%$)	44.47% ($\pm 3.19\%$)
FNN	Output	0.7349 (± 0.0417)	25.98% ($\pm 2.04\%$)	43.43% ($\pm 5.41\%$)
Transformer	Output	0.6255 (± 0.1033)	24.21% ($\pm 11.24\%$)	40.14% ($\pm 14.78\%$)
Gaussian Kernel Density				
CNN	dense1	0.9137 (± 0.0208)	53.07% ($\pm 2.51\%$)	79.71% ($\pm 7.07\%$)
CNN	dense2	0.9263 (± 0.0139)	61.61% ($\pm 4.44\%$)	83.14% ($\pm 4.44\%$)
CNN	dense3	0.9217 (± 0.0118)	62.52% ($\pm 3.50\%$)	81.29% ($\pm 4.04\%$)
FNN	linear3	0.8600 (± 0.0178)	27.30% ($\pm 4.04\%$)	65.76% ($\pm 6.38\%$)
FNN	linear4	0.9079 (± 0.0109)	56.06% ($\pm 5.32\%$)	80.31% ($\pm 2.68\%$)
FNN	linear5	0.8759 (± 0.0235)	43.55% ($\pm 9.04\%$)	71.70% ($\pm 6.78\%$)
Transformer	linear1	0.8762 (± 0.0227)	43.21% ($\pm 6.52\%$)	64.55% ($\pm 8.58\%$)
Transformer	linear2	0.9636 (± 0.0126)	79.23% ($\pm 8.25\%$)	95.73% ($\pm 2.68\%$)
Transformer	linear3	0.9550 (± 0.0094)	71.49% ($\pm 8.26\%$)	95.27% ($\pm 1.73\%$)
Normalizing Flows				
CNN	dense1	0.9071 (± 0.0047)	48.95% ($\pm 1.34\%$)	74.30% ($\pm 2.58\%$)
CNN	dense2	0.9263 (± 0.0063)	55.98% ($\pm 1.29\%$)	83.03% ($\pm 2.72\%$)
CNN	dense3	0.8982 (± 0.0130)	58.78% ($\pm 1.38\%$)	72.69% ($\pm 3.38\%$)
FNN	linear3	0.8567 (± 0.0137)	22.42% ($\pm 3.62\%$)	62.69% ($\pm 6.80\%$)
FNN	linear4	0.8963 (± 0.0112)	45.10% ($\pm 8.46\%$)	77.66% ($\pm 3.28\%$)
FNN	linear5	0.8549 (± 0.0219)	36.02% ($\pm 7.18\%$)	64.46% ($\pm 5.96\%$)
Transformer	linear1	0.9263 (± 0.0071)	77.83% ($\pm 0.39\%$)	81.84% ($\pm 1.61\%$)
Transformer	linear2	0.9536 (± 0.0072)	77.96% ($\pm 3.03\%$)	91.07% ($\pm 2.59\%$)
Transformer	linear3	0.9583 (± 0.0090)	77.71% ($\pm 4.49\%$)	94.44% ($\pm 2.24\%$)