

# Deep Learning and Zero-Day Traffic Classification: Lessons learned from a commercial-grade dataset

Lixuan Yang, Alessandro Finamore, Feng Jun, Dario Rossi  
Huawei Technologies, France

**Abstract**—The increasing success of Machine Learning (ML) and Deep Learning (DL) has recently re-sparked interest towards traffic classification. While supervised techniques provide satisfactory performance when classifying *known* traffic, the detection of *zero-day* (i.e., unknown) traffic is a more challenging task.

At the same time, zero-day detection, generally tackled with unsupervised techniques such as clustering, received less coverage by the traffic classification literature which focuses more on deriving DL models via supervised techniques. Moreover, the combination of supervised and unsupervised techniques poses challenges not fully covered by the traffic classification literature.

In this paper, we share our experience on a commercial-grade DL traffic classification engine that combines supervised and unsupervised techniques to identify known and zero-day traffic. In particular, we rely on a dataset with *hundreds* of very fine grained application labels, and perform a thorough assessment of two state of the art traffic classifiers in commercial-grade settings. This pushes the boundaries of traffic classifiers evaluation beyond the *few tens* of classes typically used in the literature.

Our main contribution is the design and evaluation of GradBP, a novel technique for zero-day applications detection. Based on gradient backpropagation and tailored for DL models, GradBP yields superior performance with respect to state of the art alternatives, in both accuracy and computational cost. Overall, while ML and DL models are both equally able to provide excellent performance for the classification of known traffic, the non-linear feature extraction process of DL models backbone provides sizable advantages for the detection of unknown classes over classical ML models.

**Index Terms**—Network measurements, Machine learning, Deep learning, Gradient backpropagation, Features extraction.

## I. INTRODUCTION

Internet traffic classification is a well investigated subject. Research in this area started in the early 2000s, aiming to create novel statistical tools to characterize both broad traffic classes and the specific applications within each class, in order to supplant both light and Deep Packet Inspection (DPI), i.e., the so called port-based and payload-based classification. Seminal works such as [1] ignited a *first wave* of approaches [2]–[7] and focused on classic Machine Learning (ML) to classify a relatively small set of applications based on a careful and hand-driven feature engineering process. This first wave culminated with very simple yet effective techniques, referred to as “early traffic classification” [3], [4], that readily used time series information (e.g., the size and direction of the first few packets in a flow) to take classification decisions.

The tremendous successes of Convolutional Neural Networks (CNN) in the image recognition field [8] ignited a *second wave* of traffic classification approaches leveraging Deep Learning (DL) techniques [9]–[18]. DL is becoming

particularly appealing in reason of domain-specific CNN hardware accelerators (known as “tensor processing units”) that started appearing in the last few years, and make CNN a viable and appealing option for real-time traffic classification [19]. In reason of the tremendous push towards encryption in the Post-Snowden era, this second wave of research is particularly relevant to industry and telco vendors actively looking at deploying statistical classification approaches — as recently pointed out in [20], until this point traffic classification remained mostly an academic exercise due to a gap between industry and academia interests, but such scenario is changing.

To better understand the reasons behind this gap, Fig. 1 sketches the fundamental blocks of a classification engine. It is based on two functions: an application identification function  $f(x)$ , and a zero-day application detection function  $g(x)$ . First, the goal of  $\ell = f(x)$  is to determine, from an input  $x$ , an application label  $\ell \in [1, K]$  among a set of  $K$  known classes. Supervised ML/DL techniques fit well this task, and learn  $f(x)$  in a process called training. However, whereas commercial DPI tools are able to handle *hundreds to thousands* of application classes [21]–[24], statistical techniques developed in the academic world consider only a *few tens of classes* — this is significantly simpler than commercial needs, hence the gap between industry and academia.

Second, the Internet applications landscape keeps evolving. Thus, commercial engines require a function  $g(x)$  to *detect zero-day applications*:  $g(x)$  assesses whether the supervised label  $\ell = f(x)$  should be rejected, implying that the input  $x$  likely belongs to a class never seen during the training of  $f(x)$ . In this case  $g(x)$  enforces  $\ell = 0$  to “signal” a zero-day detection. Clearly, for this task, unsupervised techniques (such as clustering) are a better fit than supervised ones — the hard separation of methodologies to perform  $f(x)$  and  $g(x)$  is introducing frictions, and is recognized as a major blocking point for the deployment of statistical traffic classifiers [20].

In this work, we discuss key aspects on the design of both  $f(x)$  and  $g(x)$ . First, we share insights from our experience on creating a classifier  $f(x)$  for a commercial-grade dataset comprising tens of millions of flows, and thousands of application labels, significantly larger than what is typically used in the literature. Rather than proposing new techniques, we contrast state of the art ML and DL techniques on exactly the same input (packets size time series for early classification).

Second, and most important, we design GradBP, a novel zero-day detection technique tailored for DL models, that we thoroughly compare against four state of the art alternatives. In a nutshell, GradBP exploits gradient backpropagation to assess the amount of novelty in the feed-forward inference: the larger

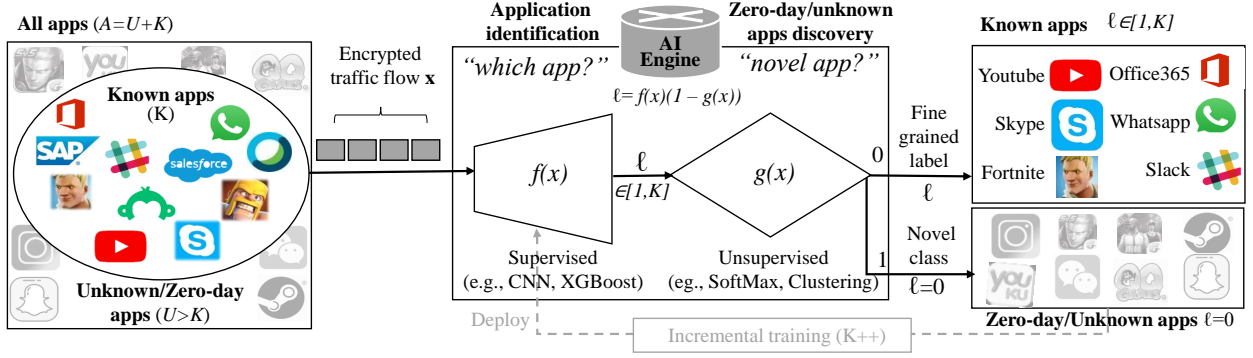


Fig. 1: Synoptic of known application identification  $f(x)$  and zero-day application detection  $g(x)$ .

the gradient of the first backpropagation step, the smaller the chance that an input  $x$  has been seen during training, hence the larger the likelihood that  $x$  belongs to a zero-day application. Our analysis shows that GradBP is (i) more *accurate* (it detects nearly 80% of zero-day traffic, compared to less than 11% for alternatives), (ii) *lightweight* (as simple as SoftMax, and over  $10\times$  faster than alternatives), and (iii) *incremental* (it does not require ad-hoc training, or model architecture changes). Overall, while we find ML and DL techniques being equally capable of accurate identification of known classes, DL models equipped with GradBP have a significant advantage over ML methods for zero-day traffic detection.

In the remainder of this paper we first cover the state of the art on both known applications identification  $f(x)$  and zero-day traffic detection  $g(x)$  (Sec. II). Then, we introduce our commercial-grade dataset (Sec. III), and the ML/DL methods selected for  $f(x)$  and  $g(x)$  (Sec. IV). We continue with a separate and thorough evaluation of the identification of known (Sec. V) and zero-day applications (Sec. VI). Finally, we summarize our findings, discuss their limits, and introduce open questions for future research avenues (Sec. VII).

## II. RELATED WORK

### A. Identification of known traffic classes $f(x)$

The early 2000’s witnessed the first wave of traffic classification, with methodologies aiming to identify fine-grained applications (e.g., YouTube, Skype, Whatsapp) or coarse-grained services (e.g., video streaming, video call, messaging). In this section we first summarize the main lessons learned from this period, and then we focus on the second wave where DL-based methods are introduced. Table I summarizes the relevant literature for both periods.

1) *First wave – ML*: Numerous surveys cover the first wave of traffic classification [7], [26]. This period focused on “classic” ML techniques, relying on either Flow Features (FF) as in [27], or Packet Payload features (PP) as in [2], [5], and it culminated with the adoption of simple, yet effective, Time Series (TS) features based on properties of the first packets of a flow such as packet size ( $S$ ), direction ( $\pm$ ), and seldom interarrival time ( $\Delta T$ ) [3], [4]. These lightweight TS approaches are particularly important since they (i) operate

“early” at the beginning of a flow, as opposite to “post mortem” techniques which compute FF after a flow ends, and (ii) sustain line rate operation with minimal additional computational costs. Indeed, whereas payload-based techniques require large amount of memory (even when processing is done on GPUs [28]), early TS techniques [3], [4] have been amenable to line rate classification in excess of 5 Mfps [29] using general purpose CPUs.

2) *Second wave – DL*: The second wave of research reconsidered all inputs features previously introduced — from PP [9], [11], [12], [18], to FF [10], [14], [15], to TS [13], [17], and hybrid FF+TS [14] — but still suffered from the same pitfalls experienced during the first wave.

For instance, studies exploiting PP essentially learn text strings, such as the Server Name Indication (SNI) [25], [30], i.e., they bind a flow to the hostname advertised (in clear) by the SNI header field. Ultimately, 1d-CNN approaches leveraging PP result into a *complex mean to do HTTPS protocol dissection, while this can be done trivially via pattern matching*. In other words, while DL is an elegant and automated way to statistically learn SNI “dictionaries”, the open question is whether CNN inference can be brought to an operational point with a lower computational footprint than traditional pattern matching techniques.

Similarly, some works are biased by the use of *port numbers* in their TS [13], [14]. For instance, authors in [13] show a striking 93% accuracy for CNN models operating on a single packet input. This ties to the fact that, in academic datasets, HTTP, SSL, and DNS (known to bound to port 80, 443 and 53 respectively) account for over 80% of flows — the DL architectures in those works *ultimately exploit the very same port-based information that they are supposed to supplant*.

More generally, the second wave shares the same key weakness already noted in the first wave: the difficulty of cross comparing models in a fair manner. Indeed, as it clearly emerges from Table I, every work uses different datasets, with different sample size (from 20k to 750k samples), and with different target classes (from 2 to 100), achieving performance in excess of  $>99\%$  (under specific conditions), making an apple-to-apple comparison difficult.

Fortunately, commendable work such as [16] started appearing, aiming at an independent evaluation of previously

TABLE I: Related work  $f(x)$ : Supervised identification of known classes

1st: ML	Input	Architecture	Samples	Classes	Performance	Notes
2004 [1]	2 FF (avg. D, avg S)	k-NN vs LDA	5.1M	4 (7)	95% (91%)	100% adding an extra feature related to $\Delta T$
2005 [2]	PP	9 DPI heuristics	573M		10 99%	macro-classes and manual ground truth creation
2006 [3]	TS <sub>5</sub> ( $\pm S$ )	K-means	n.a.		10 $\approx 90\%$	
2007 [4]	TS <sub>3</sub> ( $\pm S$ , $\Delta T$ )	2d gaussian filter	30k		4 $\approx 90\%$	multi-dimensional statistical fingerprint
2nd: DL	Input	Architecture	Samples	Classes	Performance	Notes
2015 [9]	PP (1000B)	1D CNN	300k	58	90% in top-25	Introduces DL to TC
2016 [10]	FF (40)	RF (non-DL)	131k	110	99%	Introduces APP classification
2017 [11]	PP (784B=28 $\times$ 28)	2D CNN	750k	20	99%	payload as image
2017 [12]	PP (784B)	1D CNN	750k	12	$\approx 90\%$	payload as blob
2017 [13]	TS <sub>20</sub> ( $\pm S$ , $\Delta T$ , $p$ )	LSTM + 2D-CNN	266k	15	96% (82% w/o $p$ )	several models: (CNN+RNN-2a)
2017 [14]	TS <sub>10</sub> ( $\pm S$ , $\Delta T$ , $p$ ) + FF (28)	CNN	22k	5+(5 real)	99% (88%)	
2017 [15]	FF (22)	GAN (vs DT and RF)	682k	2	all 99%	SSH vs non-SSH
2018 [16]	PP, FF, TS	MLP, SAE LSTM, CNN	138k	49	80-86%	DL [9], [11]–[13], [18] vs RF [10]
2019 [17]	TS $\rightarrow$ 2D histo (1500 <sup>2</sup> )	LeNet-5 like	21k	10	99%	
2019 [25]	TS <sub>6</sub> ( $\pm S$ ) + PP (256 $\times$ 6)	1D CNN + LSTM	unclear	80	95%	large scale dataset
2020 [18]	Payload	1D CNN and SAE	unclear	17	98%	dataset is public, but flows are not specified
<hr/>						
<b>This work</b>	<b>TS</b>	<b>1D CNN</b>	<b><math>\approx 10M</math></b>	<b>200+635</b>	<b>91% top-200</b>	<b>private commercial-grade dataset</b>

**Input:** flow features (FF), packet payload (PP), flow duration (D); time series (TS) of packet size (S), direction ( $\pm$ ), interarrival ( $\Delta T$ ) and ports ( $p$ ).  
**Architecture:** Random Forest (RF); Multi-layer perceptron (MLP); Stacked Autoencoders (SAE); Convolutional neural networks (CNN); Long-short term memory (LSTM); Generative adversarial networks (GAN)

published work. The comparison carried out in [16] (of [10]–[13], and by mean of datasets different from the ones used by original authors) reveals a different scenario from the one pictured by the original publications: (i) the expected performance drops significantly below  $<90\%$  for any architecture; (ii) there is no clear winner, although 1d-CNN models have consistently better results among the candidate approaches; (iii) 1d-CNN has a limited gain against shallow Multi Layer Perceptron (MLP) over the same input (+6%) or against Random Forest (RF) over FF input (+3%). We underline that such insights are possible only when broadening the evaluation scope beyond the typical race for the 100% classification accuracy.

Despite its merits, [16] still partially falls into an apple-vs-orange comparison. For instance, the classic RF model inherited from [10] is based on engineered FF, whereas the CNN models are either based on PP [11], [12] or packet TS [13]. As such, it is extremely difficult to attribute improvements to either the learning technique (ML vs DL) or the model input (FF vs TS). To counter this problem, sharing the same spirit of [16] *we perform an independent evaluation of two state of the art ML/DL techniques, applied to exactly the same input, on a commercial-grade dataset*. By so doing, we reveal problems that only appear at scale, and that are often not properly captured by the academic literature.

### B. Detection of zero-day applications $g(x)$

Zero-day applications detection is a network-domain problem arising when either new applications appear, or known applications change their behavior. In both cases a zero-day detection enabled classifier identifies a “new” sample as “unknown”, and avoids to label it with one of the known classes (in the remainder, we also refer to an unknown class as a zero-day class). Zero-day detection is also known as *open-set recognition* in the knowledge discovery domain. In this work, we contribute a new zero-day detection method

(Sec.IV-B) with better performance than state of the art alternatives (Sec VI).

We summarize the relevant literature about zero-day detection in Table. II dividing it into three categories based on whether  $g(x)$  is performed on the input  $x$ , at the output  $\ell = f(x)$ , or at inner-layers of the model. Note that techniques acting at the input or at the output are decoupled from the design of  $f(x)$ , which is a desirable property as altering existing models would make deployment more difficult. Conversely, this is not necessarily true for techniques using models inner-layers. Moreover, the zero-day detection computational complexity is another important aspect:  $g(x)$  is applied on each input, hence it needs to be computed faster than  $f(x)$  to avoid slowing down the overall classification process. In the following we introduce the relevant literature, while deferring a more in-depth discussion to Sec. IV-B.

1) *Input*: A first set of techniques focus on input space properties. This includes seminal works for zero-day detection in the traffic classification context such as [31], as well as the current state of the art [32] where the authors integrate into a classifier a zero-day detection module to continuously update traffic classes knowledge. Shortly, [31], [32] use K-means to cluster input data, and they identify unknown applications by thresholding the distance between an input sample and the clusters centroid. These approaches are the state of the art in the network domain, thus they are the reference benchmark.

A key challenge when using input clustering is the creation of clusters that fit well the different classes. To better control this, some works have recently proposed to apply transformations on the input so to control models output by mean of a Mahalanobis distance based score [34] or a temperature scaled SoftMax score [33]. The drawback of those proposals is their higher computational cost compared to [31], [32] which makes them less appealing from a practical perspective.

2) *Inner*: At their core, DL methods project input data into a *latent space* where it is easier to separate data based on class

TABLE II: Related work  $g(x)$ : Unsupervised zero-day application detection

	Year [ref]	Technique	Applicability	$f(x)$ Modification	Complexity
Input	2007 [31]	Input clustering	ML	None	Not assessed
	2015 [32]	✓ Input clustering	ML	None	High (hundreds of clusters for small $K$ )
	2017 [33]	Input modification wrt. Temperature scaling	DL	None	High (full backpropagation + inference)
	2018 [34]	Input modification wrt. Mahalanobis	DL	None	High (full backpropagation + inference)
Inner	2017 [35]	GAN-OpenMax	ML/DL	Yes	Medium(Weibull inference)
	2018 [36]	Classifier K+1;	ML/DL	Yes	Low(threshold)
	2018 [37]–[40]	Clustering loss Functions	ML/DL	Yes	Low(threshold)
	2018 [41]	Confidence learning	ML/DL	Yes	Low(threshold)
	2019 [42]	CNN + AE	ML/DL	None	High(AE inference)
	2019 [43]	AE	ML/DL	None	High(AE inference)
	2020 [44]	Sigmoid activation	ML/DL	Yes	Medium(Weibull inference)
Output	2015 [45]	Extreme Value Machine (FV)	ML/DL	None	Medium (Weibull inference)
	2015 [46]	✓ OpenMax (AV)	ML/DL	None	Medium (Weibull inference)
	2017 [47]	✓ SoftMax	ML/DL	None	Low (threshold)
	2020 [48]	✓ Clustering (FV)	DL	None	High (clustering)
	<i>This work</i>	<i>Gradient backpropagation</i>	<i>DL</i>	<i>None</i>	<i>Low (backpropagation on last layer)</i>

**Technique:** Feature Vector (FV); Activation Vector (AV); Autoencoders (AE), Convolutional neural networks (CNN); Generative adversarial networks (GAN)  
✓: Technique we compare against in this work

labels. A set of works then propose specific ways to alter this latent space to ease zero-day detection.

For instance, [43] uses AutoEncoders (AE) to transform input data and applies clustering to the transformed input, while [42] uses latent representation along with OpenMax [46] activation vectors. Generative Adversarial Networks (GAN) are used in [35], [36] to explore the latent space in order to generate “unknown classes” data to train a K+1 classifier for the zero-day class  $\ell=0$ . For instance, [35] generates unknown classes by mixing the latent representation of known classes, while [36] uses optimization methods to create counterfactual samples that are close to training samples but do not belong to training data. All these methods require specific architectures and extra training, thus they are less appealing.

Other works alter activation [44] or loss functions [37]–[40]. For instance, in [44] authors replace the SoftMax activation with a Sigmoid, and fit a Weibull distribution for each activation output to revise the output activation. Special clustering loss functions [37]–[40] can be used to further constraint points of the same class to be close to each other, so to project unknown classes into sparse regions far from known classes. All these methods require special DL architectures, and are not incremental. Additionally such architectural modifications can alter the accuracy of the supervised classification task, thus we deem this category inapt for the zero-day detection.

3) *Output*: Classifiers commonly complement the output classification label with a *confidence score*, i.e., the probability that the input  $x$  belongs to each of the  $K$  target classes. This is typically obtained applying SoftMax to the output layer activation vector. As such, a popular zero-day detection on output data is based on thresholding SoftMax outputs [47], i.e., rejecting a classification if the SoftMax score of the selected class is lower than a threshold. OpenMax [46] revises SoftMax by adding a special “synthetic” unknown class (induced by a Weibull modeling). Alternative approaches include Extreme Value Machine (EVM) [45], or a combination of Principal Component Analysis (PCA) with clustering [48] to reduce the “dimensionality curse”. All these approaches are appealing

given their limited complexity and do not require to change models. Our proposed method GradBP fits into this class too.

### III. DATASET

In this work we use a dataset collected from four Huawei’s customer deployments in China. The dataset does not provide any sensible information about end-users, nor raw pcap traces have been collected. Each vantage point records per-flow logs where each entry relates to a flow 5-tuple (anonymized ipSrc, ipDst, portSrc, portDst, l4Proto) with aggregate metrics (bytes, packets, mean RTT, etc.) and per-packet information ( $\pm$ PS and  $\Delta$ T time series of the first 100 packets). Each log entry is also annotated with application labels provided by a Huawei commercial-grade DPI engine [21]–[24]. Overall, the dataset corresponds to traffic activity across four weeks by tens of thousands network devices. We used the per-packet TS and the provided applications label to train our ML/DL models, without any specific pre-processing.

1) *Collection environments*: The dataset is collected via two Huawei’s product line equipments enabling network monitoring and management services for the *Enterprise campus* and *Customer OLT/ONT* market segments in China. We underline that traffic encryption in China is not as pervasive as in the Western world yet, so DPI technologies still offer fine-grained view on traffic. In our dataset, HTTPS/TLS corresponds to 33% of bytes, and 55% of flows for TCP traffic, while QUIC is negligible with only 2% of bytes for less than 1% of flows. This explains the availability of a very large number of labels: the dataset comprises 3,231 application labels that is  $30\times$  the largest number of classes considered in the literature [10]. As often remarked in the literature, the lack of public datasets is a major limitation: while this dataset is a private Huawei asset, we are investigating the possibility to release a highly anonymized (e.g., shuffled and normalized time series, etc.) and semantically deprived (e.g., no textual labels) version to the community, which we discuss further in Sec.VII.

TABLE III: Commercial-grade dataset description

Scope <sup>1</sup>	Classes	%	Flows	%	Bytes	%
$f(x)$	10	0.3%	3.4M	32.6%	5.6 TB	53.0%
	20	0.6%	4.6M	43.6%	7.2 TB	67.1%
	$K' = 50$	1.5%	7.2M	68.3%	8.8 TB	82.7%
	100	3.1%	8.7M	82.9%	9.8 TB	91.9%
	$K'' = 200$	6.2%	9.9M	94.0%	10.3 TB	97.1%
$g(x)$	250	7.7%	10.2M	95.5%	10.4 TB	98.2%
	835	25.9%	10.4M	99.3%	10.6 TB	99.8%
noise	1,000	30.9%	10.5M	99.5%	10.6 TB	99.9%
	3,231	100%	10.5M	100%	10.6 TB	100%

<sup>1</sup>Denotes the portion of the dataset which is relevant for the supervised traffic identification  $f(x)$  vs unsupervised zero-day detection  $g(x)$ .

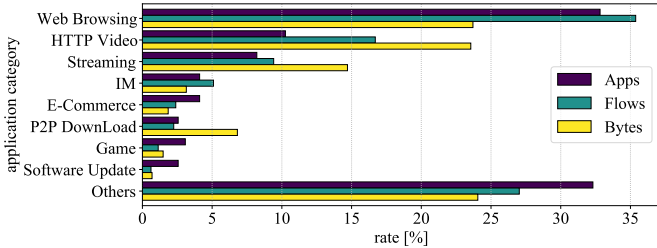


Fig. 2: Breakdown of fraction of bytes and flows by category for the top-200 applications.

2) *Traffic imbalance*: We summarize the dataset properties in Table III. The typical number of classes in academic literature  $K=10-50$  only covers a tiny portion of the applications catalog (0.3–1.5%), but captures a sizable portion of flows (32.6–68.3%) and byte-wise volume (53.0–82.7%). Yet, even when using  $K'=50$ , about 1/3 of flows (1/5 of bytes) are not covered: a commercial product needs to target about  $K''=200$  to cover roughly 95% of traffic (specifically, 94.0% of flows and 97.1% of bytes). Fig. 2 further zooms into the composition of the dataset, reporting a coarse-grained breakdown of the number of flows, bytes, and labels of the top-200 applications. We can see that applications of the *Web Browsing*, *HTTP Video*, and *Streaming* categories<sup>1</sup> represent 51% of the top-200 apps, and are responsible for 62% and 61% of flows and bytes respectively. At the same time, about 32% of apps, 27% of flows, and 24% of bytes are from *Others*, testifying the great diversity of the data, even when considering only the composition of the top-200 classes. In other words, the common scenarios studied in the literature are faraway from business needs: while academic models work well for  $K=10-50$  classes (accuracy >99%), it is not obvious how to project those results to scenarios with hundreds of classes — this is at the core of our investigation in Sec.V.

3) *Dataset scope*: As suggested by the column *Scope* in Table III, we are going to focus on the top-200 applications when discussing known application classification  $f(x)$ . We then extend the scope to include the top-835 applications (i.e., the applications with at least 100 flows) for zero-day

<sup>1</sup>The logs offer fine-grained labels as well as their coarse-grained category.

TABLE IV: Breakdown of most popular ( $K$ ) vs Zero-day TCP/UDP ( $U$ ) applications.

Proto	Known apps $K$	Zero-day apps $U$	Tot $U + K$	Perc $\frac{U}{U+K}$	Openness $1 - \sqrt{\frac{2K}{2K+U}}$
TCP	162	500	662	75.5%	37.3%
UDP	38	135	173	78.0%	40.0%
Total	200	635	835	76.0%	37.8%

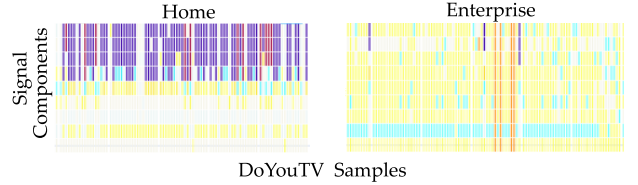


Fig. 3: Example of multi-modal behavior for the same application across two environments.

application detection  $g(x)$ . To better highlight this, Table IV further details the split between known ( $K$ ) and unknown ( $U$ ) applications for TCP and UDP traffic. The table additionally reports the percentage of unknown classes  $U/(U+K)$  and the openness  $1 - \sqrt{2K/(2K+U)}$ , which is a metric traditionally used to assess the difficulty of the open-set recognition task.

As mentioned, the top-200 applications (162 TCP and 38 UDP) cover 94% of flow and 97% of bytes, which justifies their use for  $f(x)$ . Indeed, increasing the number of classes beyond 200 minimally affects the coverage (using  $K=250$  improves the dataset coverage by only 1%). We introduce DL and ML models for known application identification in Sec.IV-A, and evaluate them in Sec.V.

Conversely, a second portion of 635 applications (500 TCP and 135 UDP) increases the overall coverage to 99.8% of bytes. We underline that the number of flows per class is too small to properly train and validate a supervised model  $f(x)$  including these classes: while the top-200 classes have on average 50,000 labeled flows, these additional 635 classes have only 750 flows on average. Conversely, this set of applications is well suited to assess zero-day detection  $g(x)$ . We stress that this small fraction of roughly 5% (2%) of flows (bytes) represents the wide majority (76%) of the overall labels. We introduce the zero-day discovery techniques in Sec.IV-B, and evaluate them in Sec.VI.

Finally, a long tail of applications (75% of the classes, with about 13% of classes having only one sample) accounts for a tiny fraction of the flows and bytes (about 0.1%). We consider those as “noise” and discard them from the analysis given their limited statistical and practical relevance.

4) *Multi-modality induced by the environment*: An interesting phenomenon we observe in our dataset relates to traffic “modes”, i.e., applications TS can change due to the network environment. We showcase this for the popular application *DoYouTv* in Fig.3. The heatmaps depict the top-100 TS across all flows observed from a residential (left) and enterprise campus (right) vantage point. Each column represents one TS, with rows visually encoding the TS values (10 packet



sizes) by mean of a color scale. Although not explicitly pointed out in previous literature, multi-modality is an intrinsic effect of access type (encapsulation, maximum segment size, etc.) and the presence of middle boxes (firewalls, NATs, etc.) that may alter packet properties (packet size in our case). This has relevance for both operating and training traffic classifiers, since those differences need to be incorporated into the training set (e.g., by including samples from all vantage points, or by distributed learning).

#### IV. METHODOLOGY

##### A. Known application identification $f(x)$

In this section we discuss the models we use in this study. The rationale of our selection is two-fold. First, we seek to avoid yet-another-solution just for the sake of the “arm race” against literature. Rather, we consider consolidated modeling techniques, focusing on an apple-vs-apple comparison on a large dataset. Second, we seek to understand models lower bound for a conservative performance assessment. As we shall see, this does not limit us from providing key observations.

1) *Output*: We focus on fine-grained traffic classification, with models targeting the identification of  $K$  classes. Recalling Table I, the DL classifiers proposed in literature consider  $K < 50$  classes. Works considering a larger set of applications either report the accuracy of the top classes (top-25 [9]) or the dataset used is practically limited to fewer classes (e.g., the top-15 classes represent over 99% of the traffic in [13]). Only a few works study  $K=50$  classes [10], [16]. Conversely, Table III shows that we need to consider  $K=200$  classes to cover 95% of flows (97% of bytes) in our dataset. We thus consider  $K'=50$  and  $K''=200$  to both compare against academic state-of-the art [16], and cover business needs. We leave as future work an in-depth evaluation of how to identify all 3,231 classes available in the dataset.

2) *Input*: Models input are TS of the first 10 UDP (100 TCP) packets size and direction  $\pm S \in \mathbb{Z}$ , which values are normalized into  $[0, 1] \in \mathbb{R}$  using the network Maximum Transmission Unit (MTU). Such TS are easy to collect, and have been consistently found to yield excellent performance across both ML [3], [4] and DL [13], [14], [16] models.

Notice that we could complement  $\pm S$  input with other packet information (e.g., packets interarrival, TCP header flags, TLS SNI strings). For instance, the first wave of traffic classification found out that the inter packets time  $\Delta T$  could improve classification performance [3], [4]. Yet, semantically this is prone to error as  $\Delta T$  can in practice represent the time between a packet sent and its response (the Round Trip Time – RTT) which correlates more with the distance between the endpoints than capturing applications behavior.

Overall, we opt for simple one-dimensional  $\pm S$  time series as to focus on models lower bound performance (desirable from a scientific standpoint), and to avoid using sensitive/privacy-related info such as TLS SNI strings (desirable from a business standpoint). Differently from previous literature [16], we use  $\pm S$  for both ML and DL models to have an apple-to-apple comparison between modeling techniques.

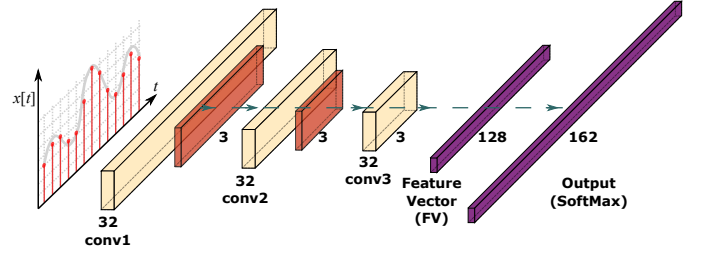


Fig. 4: Baseline CNN architecture for known application identification  $f(x)$  (hyperparameters for TCP traffic portion).

3) *DL model (1d-CNN)*: As DL model, we use a 1d-CNN given that (i) it fits well for time series input, (ii) it provides superior performance compared to 2d-CNN [11], [12], and (iii) its related literature was independently evaluated [16]. We are aware that our choice of a single 1d-CNN architecture is narrow, as more recent designs emerged, such as different 2d-CNN inputs [17], sparse-LSTM [49] or pointwise correlation [50] just to name a few. While quantitative results reported in this paper pertain to a single architecture, we discuss a broader selection of DL models in Sec.VII.

In more detail, to address the different classes and time series length, we model TCP and UDP traffic separately. As per Table IV, the top-200 classes include 162 TCP and 38 UDP applications. Fig.4 details the TCP model architecture, but the UDP model has the same structure, with layers size as annotated in parenthesis below. The input layer feeds a stack of 32 (16) convolutional filters of depth 3 and  $1 \times 3$  size, with ReLU activations, and max-pooling layers, followed by one fully connected layer of size 128 (64) for a total depth of 4 layers before the final SoftMax classification layer of size 162 (38). We underline that this architecture is rather typical and is adopted with minor differences (hyperparameters selection, and layers size) by many works [11]–[14], [17]. Further notice that we annotated one fully connected layer as FV: this layer plays an important role for zero-day detection (Sec. IV-B).

4) *ML model (XGBoost)*: As ML model, we use extreme gradient boosting (XGBoost), a tree-based ensemble technique widely acknowledged as a ML state of the art classifier in numerous domains [51], [52].

5) *Models complexity*: We define the “space” model complexity as the number of parameters  $W_{DL}$  and  $W_{ML}$  used by the DL and ML models respectively. While this can be directly measured for the CNN models, we define  $W_{ML}$  as the total number of nodes in the trained XGBoost ensemble. Moreover, as a mean to control  $W_{ML}$ , we fix the number of trees  $T=100$  and cap the individual tree depth  $d \in \{2, 4, 7\}$ .

Conversely, we intentionally avoid time complexity measurements (e.g., training duration, or inference latency) as they directly relate to implementation details and GPU/TPU hardware acceleration [19]. This still enables a fair comparison against state of the art, since only space complexity is typically discussed in (or can be easily derived from) the literature.

6) *Models comparison*: Both ML and the DL models are trained using the same dataset folds. This allows a punctual

comparison of classification accuracy and complexity (Sec.V). Despite the punctual quantification, we underline that the defined models complexity is intended for a qualitative comparison since the individual operations in the ML domain (e.g., memory comparison and branching) and DL domain (e.g., tensor products) intrinsically differ.

### B. Zero-day application detection $g(x)$

To identify zero-day applications, the  $g(x)$  function needs to reject the classification label for those input  $x$  belonging to classes never exposed to  $f(x)$  during training (see Fig.1). In other words,  $g(x)$  acts as a binary classifier: when  $g(x)=1$  the sample is considered novel and the classification result is rejected by enforcing  $\ell = f(x)(1 - g(x))=0$ . We consider  $g(x)$  functions from the literature based on input data, inner/hidden model layers telemetry, or the models output. Additionally, we contribute GradBP, a novel technique which we find superior to alternative methodologies (Sec. VI).

1) *Input based clustering (ML and DL)*: Clustering is a common technique to identify unknown classes using input data [31], [32]. Clusters are formed using training samples of the known applications. Comparing the distance between an input  $x$  and the clusters centroid with a predefined threshold allows to accept/reject the  $f(x)$  label. Specifically, the reject criterion is:

$$g^{IN}(x) = \mathbb{1} \left( \min_{c \in [1, C]} d(x, c) > \epsilon_{IN} \right) \quad (1)$$

where  $d(x, c)$  is a distance metric (Euclidean, Manhattan, etc.), and  $\epsilon_{IN}$  is an arbitrary threshold. Intuitively, a small  $\epsilon_{IN}$  increases the chance to reject a correctly classified sample of a known application (false negative), whereas setting a too large threshold leads to unknown applications going undetected (false positive). The reference state of the art technique is [32], which uses K-means to identify a large number of clusters  $C \gg K$  so to fit the input space. Clearly, the number of clusters  $C$  trades off computational complexity (as a point needs to be compared against  $C$  centroids) and accuracy (as the larger is  $C$ , the better the coverage of the input space).

2) *Inner layers based clustering (DL only)*: When using a DL model, clustering can be applied also to FVs, i.e., the output generated by one of the hidden layers when processing an input sample (see Fig.4). Essentially, the training process of DL models performs non-linear transformations that improve its discriminating power. The output of the FV layer is a *high-dimensional latent space* where the projected input samples are more easily separable with respect to their class. It is possible to induce clustering in the latent space by either using a clustering loss function [37]–[40], or by regularizing the space (e.g., by normalizing the layer into a hyper-sphere). The  $C$  clusters are constructed at training time applying K-means to  $FV(x)$ . The rejection criterion is:

$$g^{FV}(x) = \mathbb{1} \left( \min_{c \in [1, C]} d(FV(x), c) > \epsilon_{FV} \right) \quad (2)$$

where, as previously,  $d(\cdot)$  and  $\epsilon_{FV}$  are a distance metric and an arbitrary threshold respectively. This technique is standard

in the DL domain, and can be considered as the state of the art for DL-based clustering [48].

3) *Output based rejection (ML and DL)*: Output-based techniques leverage additional information from the ML/DL model output, such as the SoftMax [47] or OpenMax [46] probabilities of each class. Denoting with  $v(x)$  the activation vectors of an input sample  $x$ , and with  $v_i(x)$  the  $i$ -th component of the vector, then the SoftMax value for class  $k$  given input  $x$  is:

$$P(y = k|x) = \frac{e^{v_k(x)}}{\sum_{i=1}^K e^{v_i(x)}} \quad (3)$$

Denoting further with  $c = \arg \max_k P(y = k|x)$  the most likely class selected by the supervised model, then the SoftMax rejection criterion is:

$$g^{SM}(x) = \mathbb{1} \left( P(y = c|x) < \epsilon_{SM} \right) \quad (4)$$

that is, the SoftMax of the output class  $c = f(x)$  is required to be larger than an arbitrary threshold  $\epsilon_{SM}$ . This technique should be considered a naïve baseline since supervised models are knowingly *overconfident* — the class confidence generated by SoftMax can be high despite the classification is wrong, which stems from the nature of SoftMax to saturate [53].

Part of the problem is rooted in the fact that the SoftMax function normalizes only considering the space of known applications. To solve this issue, OpenMax [46] introduces an extra “synthetic zero label” and re-normalizes the activation vector before computing distances. In more detail, OpenMax introduces a weight vector  $w_b$  that captures the distribution of the activation vectors  $v(x)$  of the  $K$  known classes by fitting a Weibull distribution. The vector  $w_b$  includes also an extra synthetic activation value  $\hat{v}_0(x)$ , namely the *novel class*  $\ell = 0$ . A new activation vector  $\hat{v}(x)$  is then derived as follows:

$$\hat{v}(x) = v(x) \circ w_b(x) \quad (5)$$

$$\hat{v}_0(x) = \sum_{k=1}^K v_k(x)(1 - w_b^k(x)) \quad (6)$$

OpenMax probabilities are then derived with the re-normalized activation vector, including for the novel class  $c=0$

$$P'(y = k|x) = \frac{e^{\hat{v}_k(x)}}{\sum_{i=0}^K e^{\hat{v}_i(x)}} \quad (7)$$

Notice also how (7) normalizes starting from  $i=0$  to reflect the extra synthetic activation, while (3) normalizes starting from  $i=1$ . Denoting with  $c' = \arg \max_k P'(y = k|x)$  the class with the largest OpenMax values, the rejection criterion then becomes:

$$g^{OM}(x) = \mathbb{1} \left( c' = 0 \vee P'(y = c'|x) < \epsilon_{OM} \right) \quad (8)$$

which rejects when either the novel class  $c = 0$  has the largest value, or when the normalized OpenMax value for the most likely class  $c \in [1, K]$  is smaller than a threshold  $\epsilon_{OM}$ .

4) *GradBP– Gradient based rejection (DL only)*: Finally, we introduce GradBP, our novel zero-day detection method. The idea is to perform a “shadow” training step: (i) evaluate the output label  $\ell = f(x)$ ; (ii) treat the label  $\ell$  as the groundtruth and compute the magnitude of the first backpropagation step  $\delta^{L-1}$  but without actually altering the model weights; (iii) use  $\delta^{L-1}$  to assess if the input is of a known or unknown application. More precisely, let us consider

$$\delta^L = \nabla_a \mathcal{L} \odot \sigma'(z^L) \quad (9)$$

where  $\nabla_a \mathcal{L}$  is the partial derivative of the model loss function  $\mathcal{L}$  with respect to the activation (i.e., the magnitude of model weights update) and  $\sigma'(z^L)$  is the activation vector  $v(x)$  at the model layer  $L$ . We limit the backpropagation to the output of the convolutional layers (i.e., the output of the *model backbone*), which already captures the majority of the information the model uses to represent the input. This corresponds to

$$\delta^{L-1} = (W^{(L-1)T} \delta^L) \odot \sigma'(z^{L-1}) \quad (10)$$

where  $W^{(L-1)T}$  is the transpose of the weight matrix of the  $(L-1)$ -th layer. Intuitively, the larger the  $\delta^{L-1}$ , the more likely the input relates to an unknown class. Based on this intuition, we conceive a simple family of gradient-based rejection criteria:

$$g_n^{GR}(x) = \mathbb{1} \left( \|\delta^{L-1}\|_n > \epsilon_{GR} \right) \quad (11)$$

where the norm  $\|\cdot\|_n$  and  $\epsilon_{GR}$  are free hyperparameters. In particular, for L1 (i.e., max gradient) and L2 norms (i.e., the square root of the squared gradient sum) we have:

$$g_1^{GR}(x) = \mathbb{1} \left( \max_i \delta_i > \epsilon_{GR} \right) \quad (12)$$

$$g_2^{GR}(x) = \mathbb{1} \left( \sqrt{\sum_i \delta_i^2} > \epsilon_{GR} \right) \quad (13)$$

Clearly, backpropagation is an essential tool, and gradients have been used in many aspects of DL, from training (e.g., to speed up convergence [54] possibly in distributed settings [55]) to extract side-channel information (e.g., to gather information about clients participating into a federated learning cohort [56]). However, we are not aware of such use in neither zero-day applications detection, nor in open-set recognition.

## V. IDENTIFICATION OF KNOWN APPLICATIONS

Models in literature are trained to identify  $K' \leq 50$  classes, and evaluated using a test set containing  $X \in \{10, 50\}$  classes. In this section, we extend this approach towards a more business-driven scenario with  $K'' \leq 200$  and  $X \leq 200$  classes. We train and evaluate TCP (162 applications) and UDP classifiers (38 applications) separately to better address their traffic diversity, but we also report their combined performance (top-200 applications).

The use of a large number of classes  $K$ , that we consider here as a simple proxy of class *variety* (including effects such as class imbalance, multi-modality, inter-class similarity, etc.), allows us to better capture macroscopic trends that we believe

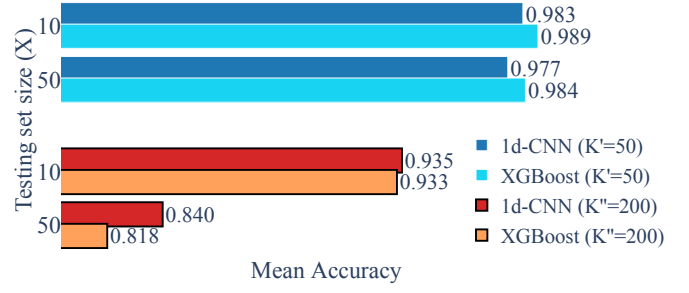


Fig. 5: *Class cardinality bias*: When training the same ML/DL model for a large number of classes  $K'' \gg K'$  accuracy drops also for the top  $X \leq K$  classes, which is unnoticed in literature.

were not well represented in previous literature. In particular, we start by illustrating how restricting the analysis to the top- $K$  classes can bias the classification accuracy (Sec.V-A). We continue by digging into classification confounding factors and labels semantic (Sec.V-B). Finally, we conclude by discussing models complexity (Sec.V-C).

### A. High-level view

1) *Class cardinality bias*: We start by comparing ML and DL models trained so that when  $X=K'=50$  their flow-level accuracy matches what reported in the literature for the top- $X=50$  classes ( $\approx 99\%$ ). Fig.5 shows the models mean classification accuracy when using  $X \in \{10, 50\}$ . We can see that models performance is stable when training with 50 classes. Conversely, when increasing  $4\times$  the number of classes ( $K''=200$ ), the accuracy notably degrades for the top-10 classes (even further for the top-50). *In other words, focusing on a small number of classes (as common in literature) makes the problem trivial to solve (even without DL).*

2) *Mean accuracy bias*: To further elaborate, Fig.6 shows the evolution of the mean accuracy for a model with  $K''=200$  classes (combining TCP and UDP models). We consider the top- $X$  classes when ranked by their popularity (classes with larger number of flows come first), and performance (best classified class comes first). The picture shows that despite classification accuracy worsen as  $X$  gets closer to  $K''$  (dashed line), the overall number of correctly classified flows remains satisfactory in reason of the application popularity skew (solid line). This effect appears only for a number of classes larger than the one typically used in the academic literature. *Otherwise stated, focusing on the average accuracy of a limited number of classes  $K' \ll K''$  hides phenomena typical in commercial settings.*

3) *DL bias*: We argue that the second wave of traffic classification published until this moment may have exaggerated the quality of DL methodologies, as already observed in computer science [57], and other fields [58]. We exemplify this in Fig.6 showing that XGBoost models performance can (significantly) surpass CNN models performance by simply tuning the maximum tree depth — although we could have, deliberately or unwittingly, adopted the opposite viewpoint.



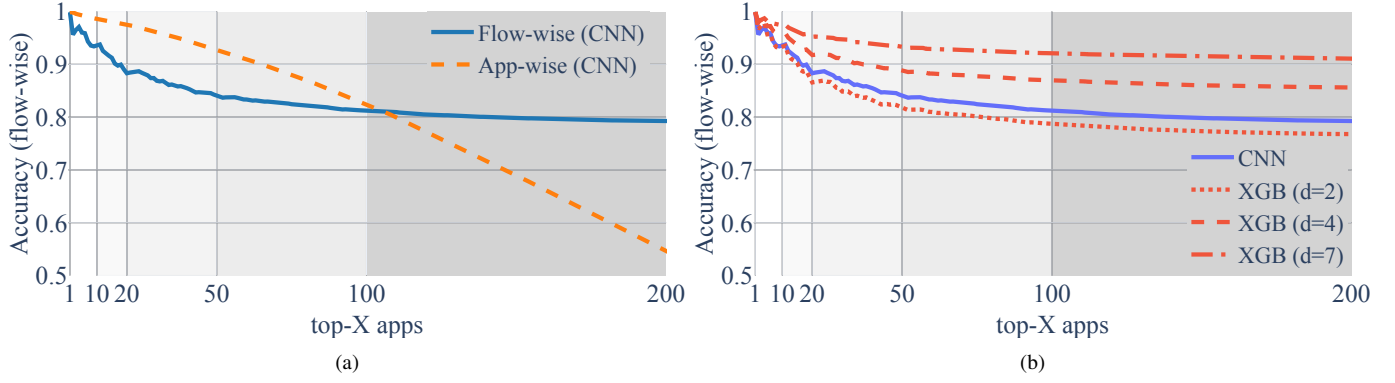


Fig. 6: (a) Mean accuracy bias. Average per-flow (affected by class imbalance) or per-application accuracy (each application counted equally). Notice that the per-application accuracy overestimates performance in the small class regime  $K < 100$ , and vice-versa happens for  $K > 100$ . (b) Model bias. By tuning model hyperparameters, it is easy to obtain operational points that show superiority of one class of approaches (purposely biased to ML in this example).

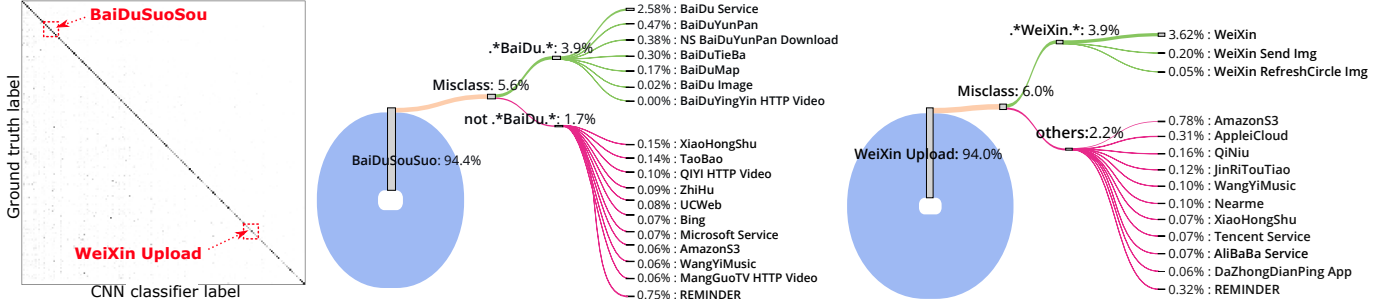


Fig. 7: Investigating classification performance for TCP traffic: (left) classification confusion matrix, with applications sorted by volume of bytes — performance is very good, but application with less traffic can suffer from lower accuracy; (center) BaiDuSouSou classification breakdown — misclassifications mostly relate to alternative labeling of traffic of the same application; (right) WeiXin Upload — a large portion of misclassifications are due to other behaviors of the same application.

Thus, as far as classification of known traffic is concerned, we find that ML and DL models are equally well fitted to support large-scale commercial scenarios.

### B. Deeper insights

We next dive into labels semantic to dissect misclassification causes. Again, we consider  $K''=200$  classes by combining TCP and UDP models. Fig.7-(left) shows the classification confusion matrix as a heatmap, with applications sorted alphabetically (labels not shown for readability): the sharp diagonal indicates that flows are classified with the expected label.

To further drill down, we pick a representative class for each side of the diagonal, *BaiDuSouSou* (the BaiDu search service) and *WeiXin Upload* (the file upload service of the WeiXin messaging application), and we dissect their classification results by mean of a sankey diagram. For *BaiDuSouSou*, Fig.7-(center) shows that 5.6% of flows are misclassified, but 3.9% of flows are labeled as *BaiDu*-related services (e.g., searches for image, maps, or social) and overall only 1.7% of flows are completely misclassified. The same considerations hold for *WeiXin Upload* where only 2.2% of misclassifications are imputed to non *WeiXin* related classes, and other popular applications (not shown). In other words, a non negligible part

of misclassifications are for neighboring services of the same application provider.

On the one hand this phenomenon ties to the very fine-grained nature of the groundtruth labels in our dataset; on the other hand it relates to the existence of different “modes” of the same application. Intuitively, the  $\pm S$  time series capture application-level signaling at the beginning of a flow. Different applications might share common patterns when being part of the same “umbrella” of services and relying on the same codebase and libraries (as in the case of *Baidu*), leading to “soft misclassifications”. Overall, this suggests that fine-grained applications identification is possible, with high accuracy, even for a larger sets of known applications than what was considered in the literature.

This holds particularly true for applications where sufficient number of samples are provided to the training process: notice indeed the degradation of flow-wise accuracy for less popular applications in the top-200, as early shown in Fig.6-(b), confirming that class imbalance in the real traffic helps maintaining satisfactory accuracy even in commercial-grade cases. At the same time, recall that labels are scarce beyond the top-200 applications (Sec. III). Thus, training  $f(x)$  with a larger number of classes calls for experimentation with few-

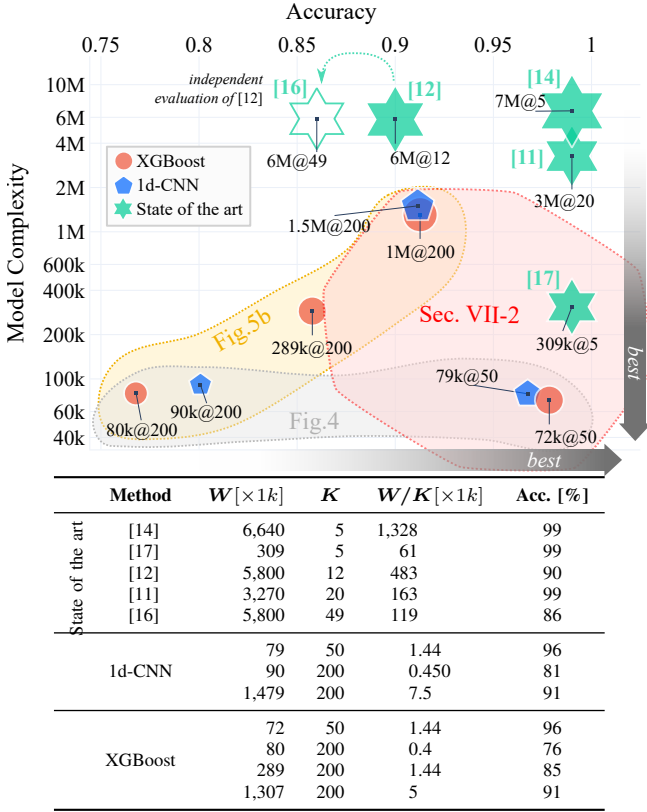


Fig. 8: *Models performance*. Scatter plot of model complexity vs accuracy, for 1d-CNN (pentagons), and XGBoost (circles). References from the State of the Art (green stars) are included for *qualitative* comparison purposes. Shapes are directly annotated ( $W@K$ ) with the model size  $W$  and classes  $K$ , and the shape size is proportional to the weights-per-class  $W/K$  value (the table details the scatter plot values).

shot learning [59], and ensemble-based methodologies [60], which we discuss further in Sec. VII.

### C. Complexity

In our view, traffic classification literature tends to underestimate models complexity. As explained in Sec.IV-A, for DL models the space complexity is the number of model parameters  $W_{DL}$ , while for ML models we use the number of nodes in the tree ensemble  $W_{ML}$ . We reiterate that space complexity does neither directly translate into computational complexity (as this depends on the ML/DL architecture, and the specific operations executed at inference time), nor energy expenditures (as this depends on the available hardware). Yet, space complexity allows to abstract from specific implementations (i.e., software frameworks, hardware acceleration, system design choices), enabling a *qualitative* comparison between models. In particular, as models have disparate capabilities, we compare not only the absolute models size  $W$ , but especially the model size normalized over the number of output classes  $W/K$ .

The scatter plot in Fig.8 illustrates the accuracy-vs-complexity trade-off for both our and literature models. We

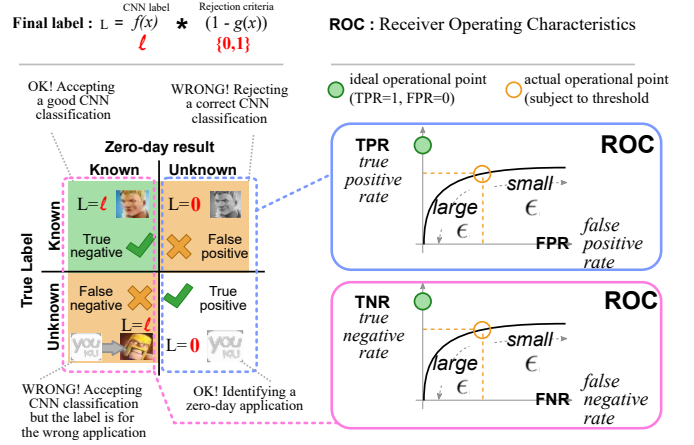


Fig. 9: *Quantifying zero-day application detection*: confusion matrix based on zero-day detection (left); ROC curve for zero-day application detection (right).

underline that while the picture provides quantitative assessment of the accuracy and complexity performance, comparison among literature models should be interpreted *qualitatively* (as the datasets on which these results have been gathered differ). Results show that models in the literature are quite heavyweight, using up to  $W/K = 6.6M/15 = 1.3M$  weights per-class and generally  $W/K \gg 100k$  weights per-class, with the most parsimonious approach employing about 61k weights per-class [17]. In contrast, our CNN and XGBoost models can achieve about 90% accuracy on the top-200 classes with at least two orders of magnitude less weights (just 5-7k weights per-class), while still covering a  $4\times$  bigger number of classes.

Two other observations emerge from the picture. On the one hand, when the number of classes is large ( $K \approx 200$ ), it becomes necessary to increase models size to maintain accuracy performance, but CNN and XGBoost models can be tuned to achieve similar performance (recall Fig.6). On the other hand, when the number of classes is small ( $K \approx 50$ ), it is unreasonable to use an humongous number of weights to discriminate them (see Fig.5), particularly since it is possible to design parsimonious models with just hundreds of weights-per-class achieving same-or-better performance (i.e., bottom-right region of the scatter plot, which we further discuss in Sec.VII). As such, by neglecting model complexity, the risk is to propose solutions that are equivalent to “shoot a mosquito with a cannon”, which may hinder deployability.

## VI. DETECTION OF UNKNOWN APPLICATIONS

In this section we evaluate zero-day applications detection. We built upon the CNN models evaluated in the previous section by applying the zero-day detection techniques introduced in Sec.IV-B: input clustering, feature vectors clustering, SoftMax, OpenMax, and our novel GradBP method. We first introduce the experimental workflow (Sec.VI-A) to overview zero-day detection across the methods (Sec.VI-B). Then we elaborate on the reasons behind their performance differences (Sec.VI-C), and we conclude evaluating their computational complexity (Sec.VI-D).

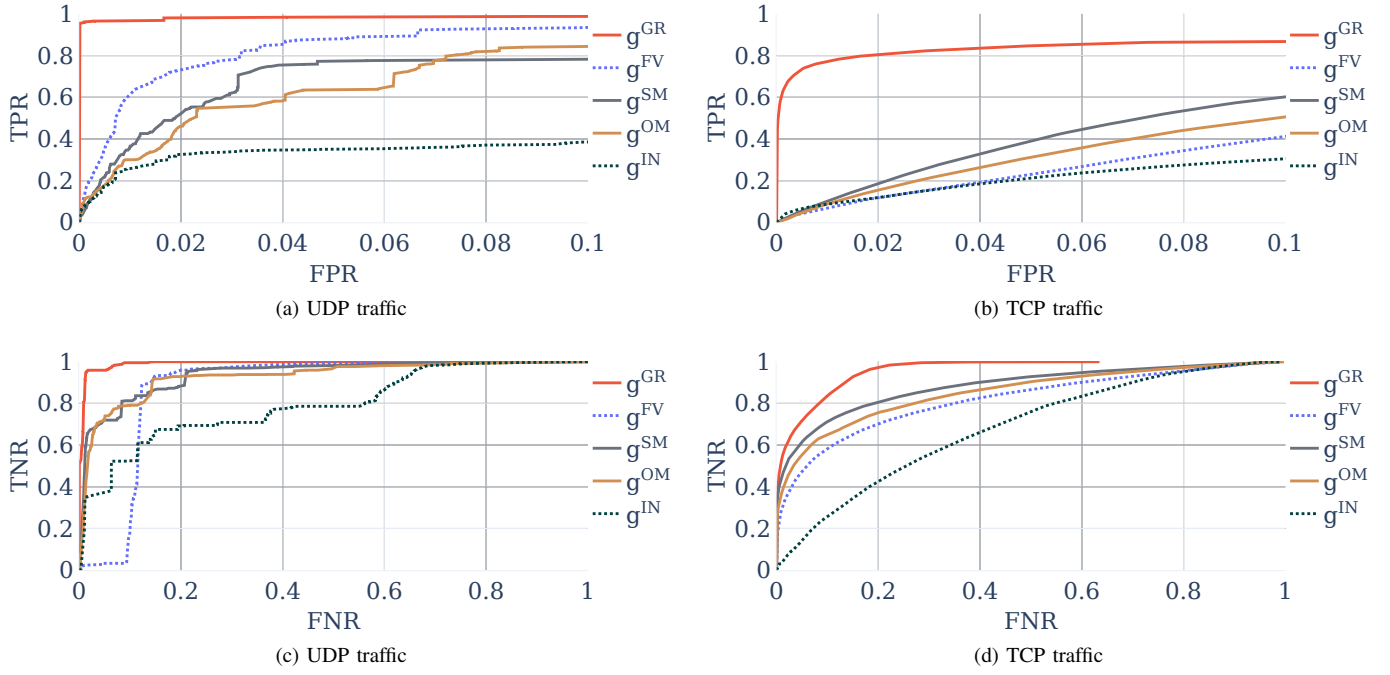


Fig. 10: Zero-day applications detection: ROC curves for UDP (left column) and TCP (right column) applications. Two families of ROC curves are reported: TPR vs FPR (top row) and TNR vs FNR (bottom row).

#### A. Workflow

1) *Overall comparison:* Fig.9 sketches our evaluation methodology. Given a trained model, we perform the classification of both known and unknown classes, and construct the Receiver Operating Characteristic (ROC) curve for all novelty detection methods. Usually, ROC depicts the achievable True Positive Rate (TPR) on the y-axis, i.e., how well the model correctly identifies the 635 zero-day applications, for a given False Positive Rate (FPR) on the x-axis, i.e., how often the model rejects a known top-200 application, wrongly relabeling it as zero-day application. We additionally build a ROC curve to depict the complementary True Negative Rate (TNR) on the y-axis, i.e., the fraction of correctly accepted classification of known traffic, for a given False Negative Rate (FNR), i.e., the fraction of undetected zero-day traffic that are left mislabeled.

The ROC curves depict the different (FPR, TPR) and (FNR, TNR) achievable operational points, by spanning over all possible values of the zero-day rejection thresholds  $\epsilon$  discussed in Sec. IV-B. Hence, the ROC curves are an intuitive way to compare different methods irrespective of the specific tuning: when the ROC curve for a method-A is higher than the ROC for method-B, method-A is better than method-B.

2) *Fixed tuning:* Any rejection method operates in fixed setting, i.e., it would be desirable to select a threshold  $\epsilon$  enabling the system to operate as close as possible to the ideal (FPR,TPR) and (FNR, TNR) operational points (top-left corner of a ROC curve). Therefore, we additionally compare methods at a *fixed operational point* by selecting, for each method, the value of  $\epsilon$  that corresponds to a *target performance* in a metric of interest. In detail, using the training dataset, we set the  $\epsilon_X$  for method  $g^X(\cdot)$  such that 99% of classifications are accepted as TNR, thus tolerating 1% of wrongly rejected

FPR classifications.

Despite individual  $\epsilon$  settings differ across methods, one can then relatively compare all methods fairly by constraining their TPR when FPR=1% (or their FNR when TNR=99%). Importantly, as the above TNR=99% and FPR=1% targets only depend on the known applications that the model has been trained with, setting  $\epsilon$  can be done without any prior knowledge of zero-day traffic.

#### B. High-level view

1) *Overall comparison:* Recalling Table IV, the TCP model can identify 162 known classes, and zero-day detection is tested against 162+500 known+unknown classes; the UDP model can identify 38 known classes, and zero-day detection is tested against 38+135 known+unknown classes. Notice that input and feature vectors clustering require a specific training step to apply K-means: we use  $C=1,000$  clusters (which is  $5\times$  the number of known applications, to account for multimodal classes) and gather the cluster centroids from the same training data used to train the supervised model  $f(x)$ . The ROC curves depicting zero-day detection capabilities for all methods are reported in Fig. 10, for both UDP and TCP, and for both types of ROC curves (notice that in the FRP-vs-TPR ROC, we visually limit FPR<10% as higher values would represent unreasonable operational points).

It is immediate to gather that GradBP ( $g^{GR}$ ) outclasses all alternatives by a large margin: GradBP performance is almost perfect for UDP (note that TPR=96% when FPR=0.1%) and remains good in the more challenging TCP scenario (TPR=78%, FPR=1%) especially compared to alternative approaches. We tuned OpenMax as recommended by the original authors [46], i.e., with a tail size of 20 when fitting the

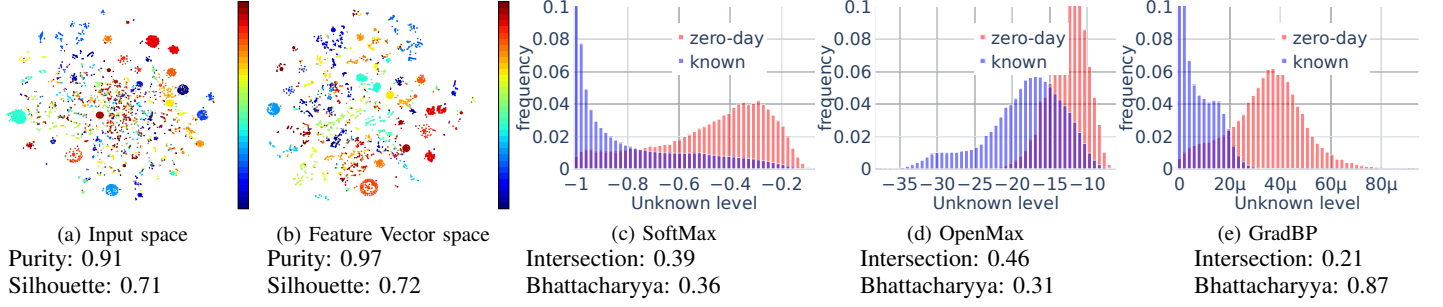


Fig. 11: *Deeper insight on zero-day application detection*: Qualitative visual differences of (a) Input-clustering and (b) FV-clustering via t-SNE projection (UDP application only), and histograms of unknown level for (c) SoftMax, (d) OpenMax and (e) GradBP methods. Relevant cluster (purity, silhouette) and distribution distance metrics (Intersection, Bhattacharyya distance) are used to quantitatively support the visual comparison.

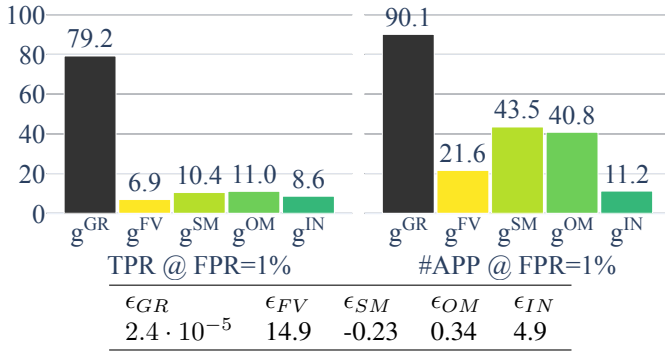


Fig. 12: *Zero-day applications detection*: Fraction TPR of correctly detected zero-day flows (left) and applications (right) when FPR=1% of correctly identified known applications are rejected (the specific  $\epsilon$  settings yielding the FPR=1% target are tabulated in the picture for each method).

$\omega_b$  Weibull distribution. Yet, differently from [46], SoftMax outperforms (or marginally differs from) OpenMax. Clustering the input data ( $g^{IN}$ ) is the least performing option. Instead, the performance for clustering the feature vectors ( $g^{FV}$ ) varies depending on the scenario: for UDP, it is the second best option; for TCP, is not better than alternatives. While increasing the number of clusters may provide benefits [32], as we shall see, the current setting already has a prohibitive computational cost (Sec. VI-D), which makes increasing the number of clusters hardly viable.

2) *Fixed tuning*: Fig.12 compares the different methods performance when focusing on the top-200 applications (i.e., combining the TCP and UDP models), and by settings  $\epsilon$  so to obtain FPR=1%. The left-side of the figure measures zero-day detection per-flow; the right-side measures zero-day classes detection, where a zero-day class is considered detected if at least one of its flows is detected. While GradBP correctly detects 79.2% of all zero-day flows (90.1% of the classes) the other approaches are only able to detect at most 11% of zero-day flows (43.5% of classes).

Additionally, Fig.10 and Fig.12 also suggest that tolerating a higher FPR would only have marginal improvements and that GradBP is less sensible to threshold tuning. In fact, while

for all methods the TPR varies depending on the FPR, for GradBP the TPR increases very quickly for small FPR, and grows slowly afterward, i.e., after the curve “knee” it is almost insensitive to the configured threshold.

### C. Deeper insights

Results discussed so far clearly show differences between methodologies, but the underlining reasons causing the performance gaps are not obvious. We therefore dig into the specific details of each method, starting from clustering. Recalling Fig.10, input-clustering is the worst performing method. This is expected since the non linear transformations operated by the CNN feature extraction indeed empowers better separation in the latent space than what could be achieved in the input space.

We can however visualize this effect by mean of a bi-dimensional t-SNE [61]. Fig.11-(a,b) t-SNE plots show the 38 UDP known applications, each associated to a different color. Both plots show sharp clusters, but notice how the center of the input space (Fig.11-a) is more confused than the FV space (Fig.11-b). We further quantify this effect by computing the *clusters purity* (the average fraction of the dominant class in each cluster) and the *clusters silhouette* — the metric captures how similar a class is to its own cluster (*cohesion*), compared to other clusters (*separation*). We find that the latent space has the highest purity (0.97-vs-0.91), but presents a similar silhouette with respect to the input space (0.71-vs-0.72).

To investigate the three remaining methods we measure the *unknown level* which quantifies predictions uncertainty: based on (3), the unknown level for SoftMax is the negative of the maximum probability  $UL^{SM} = -\max_{k \in [1, K]} P(y = k|x)$ ; following (6), the unknown level for OpenMax is the synthetic activation  $UL^{OM} = \hat{v}_0(x)$ ; finally, following (10), for GradBP the unknown level is the magnitude of the backpropagation update  $UL^{GR} = \|\delta^{L-1}\|_n$ .

Fig. 11-(c,d,e) show the distribution of the unknown level for known and unknown classes. To further quantify the overlap of the distributions, we annotate the figure with the *Intersection index I*

$$I = \sum \min(UL_{known}, UL_{zero-day}) \quad (14)$$



TABLE V: Zero-day application detection complexity:  $g(x)$  bootstrap and inference costs

Method	$g(x)$	Bootstrap cost (per dataset)	Inference cost (per sample)
Gradient	$g_{u,2}^{GR}$	-	0.01 ms
	$g_{u,1}^{GR}$	-	0.01 ms
SoftMax	$g^{SM}$	-	0.01 ms
OpenMax	$g^{OM}$	11s	0.13 ms
FV-clustering	$g^{FV}$	24m 36s	4.58 ms
Input-clustering	$g^{IN}$	20m 4s	4.46 ms

and the *Bhattacharyya distance*  $D$

$$D = -\ln \sum \sqrt{UL_{known} * UL_{zero-day}} \quad (15)$$

Ideally,  $g(x)$  should be able to well separate the unknown level distributions for known and zero-day classes, or equivalently  $g(x)$  should have a low intersection index  $I$ , and high distance  $D$ . Intuitively, known classes should provide a small unknown level. For instance, the unknown level for SoftMax is capped at -1 which expresses a strong confidence in the model prediction. Conversely, when processing traffic of a zero-day application, the confidence is expected to reduce. The same is true for OpenMax and Gradient-based, although the unknown levels are unbounded. We can see that distributions overlap for OpenMax, leading to an higher  $I$ , and a lower  $D$  compared to the other methods, which explains the lower performance. This large confusion hints that the activation vectors might have a large diversity for the known classes, hence a simple mean activation vector (6) is insufficient. Similarly, GradBP presents half the  $I$  of SoftMax, while  $D$  is over  $2\times$  larger, which explains the better performance of GradBP.

#### D. Complexity

We conclude comparing initialization costs and inference computational costs as summarized in Table V. Results are obtained using TensorFlow v1.9 on an NVIDIA P100 GPU. SoftMax and GradBP do not need any specific pre-processing, OpenMax introduces a small cost to fit the Weibull distribution, and clustering-based methods (as expected) have the largest initialization cost (recall that those use  $K$ -means with  $C=1,000$  clusters in our setup).

We report initialization costs just for reference. In fact, given their episodic and offline nature, they are less relevant from a practical perspective than inference computational costs. In particular, we find OpenMax being  $13\times$  slower than GradBP, and clustering-based methods being more than two orders of magnitude slower than our method. Indeed, they suffer from the computation of a large number (specifically  $C$ ) of pair-wise distances (involving power of two, and square root operations) in a fairly large space (considering Fig.4 which details a TCP model, the FV space has 128 dimensions, and input space has 100 dimensions). We point out that while distance computations can be reduced (e.g., smaller latent space, int8 quantization, Manhattan distance, etc.) the performance of clustering-based methods can be hardly expected to improve. In fact, while increasing the number of clusters

$C$  might improve accuracy performance, this would increase their computational cost too.

GradBP therefore brings the best of both worlds, as it is significantly more accurate and computationally lightweight. Specifically, for a TCP+UDP model with  $K=200$  classes, GradBP performs an Hadamard product between vectors of size 200 in (9), and a matrix multiplication between a  $128\times 200$  and  $200\times 1$  matrices, followed by a Hadamard product between vectors of size 128 in (10). Additionally, while the computational costs results are to be interpreted in a relative sense (i.e., across algorithms), we stress that the raw performance reported here comes from a non optimized implementation — yet, using TensorFlow v1.9 on a NVIDIA P100 GPU, we reached 100,000  $g(x)$  inferences per-second which is enough to sustain real system requirements [19].

## VII. SUMMARY AND DISCUSSION

In this paper, we focused on the design of a commercial-grade traffic classification engine capable of (i) fine-grained applications identification of hundreds of classes and (ii) zero-day detection of applications that were not part of models knowledge base. We tested the engine on a large dataset using the top-200 applications (covering 95% of the flows and bytes) for the identification of known classes, and top-835 applications (extending the flows and bytes coverage to over 99%) for zero-day detection.

Summarizing our main findings, we gathered that (i) ML and DL models are both equally able to provide satisfactory solutions for the classification of known traffic. In particular, the type of CNN architectures used in the literature are well suited as their accuracy exceeds 90% for the top-200 applications. At the same time, our results also pointed out that models complexity is commonly overlooked in literature. Given the relatively simple tasks, this generates unnecessarily complicated models which endangers the practical relevance of the research.

More interesting, we gathered that (ii) ML and DL differ in their abilities to detect zero-day traffic, as the non-linear feature extraction process of DL models backbone yields to sizable advantages over ML for this task. In particular, our main contribution was to provide a novel gradient-based technique for zero-day detection tailored for DL classifiers, that we found to outperform the related literature in both accuracy and computational costs.

Despite our efforts, we still identify a few key open questions, related to (i) open datasets, (ii) models complexity, (iii) novelty discovery techniques and (iv) models deployability, that needs further discussion.

1) *Open dataset — pooling efforts*: In reason of our results, constructing an open corpus with rich *class diversity* (and thus necessarily a large *class cardinality*  $K$ ) should be a priority goal to allow for meaningful and fair cross-comparison of research proposals. While this is within reach for large industrial players, legal and business aspects prevent them to share openly their datasets.

For instance, we are currently investigating the possibility to release a highly anonymized version of the dataset used



in this paper, as part of Huawei Rapid Analytics & Model Prototyping (RAMP) Data Challenges.<sup>2</sup> However, while the dataset we used does not contain information that constitutes a *privacy-related* risk (e.g., IP addresses, timestamps), it contains *business-sensitive* information (e.g., fine grained labels and packet size sequences). Such information requires further processing (label obfuscation, time series shuffling, etc.), to ensure that researchers can still carry meaningful experimental activities (e.g., use the dataset as benchmark for their neural architectures), while not endangering the business (e.g., if relevant labels were open-sourced, this would allow to train and openly release models of commercial value).

This said, we observe that while collecting large volumes of real labeled network data is a daunting effort for a *single academic* partner, a pooling effort across *multiple research groups*  $N$  in a coordinated manner can be an effective strategy to achieve this goal — for instance, each partner can gather  $K/N$  classes, coordinating to keep a null/low class overlap between groups. This is a commonplace in computer vision (e.g., ImageNet has 15 million images labeled in 20k classes), and the traffic classification community should take inspiration from those efforts. Indeed, different research groups are already doing active measurement collections for specific application types (video, games, etc.) and with a different goal than traffic classification (congestion control, QoE, etc.), so we conjecture that the true burden lies in the *coordination*. Breaking this barrier would allow to tackle more challenging classification problems, with a more significant application diversity and a set of fine-grained labels.

2) *Complexity — the need for models ecology*: For the classification of known traffic, in this paper we decided not to participate to the arm race against literature (as we admittedly did in the first wave), and we avoided to propose yet another “new” model for traffic classification. Quite the opposite [57], we decided to perform the simplest and meaningful selection of relevant DL and ML models (based on our own experience, as well as on consolidated results reported in the state of the art). This choice allowed us to clearly state a very simple yet important message — as far as classification of known traffic is concerned, ML and DL models are equally well fitted to support large-scale commercial scenarios.

Given our choice to limit the set of considered models, in this paper we intentionally avoided answering the question of which is the model that best fit traffic classification. We argue that while most of the research in DL focused on models *accuracy* comparison (recall Table I), we believe that it is now necessary to explore more systematically the *accuracy-vs-complexity* trade-off. In particular, we observe that the “Green networking” [62] research wave (i.e., the explicit consideration of energy expenditures in network protocols, algorithms, and architecture designs) predates by over a decade the corresponding “Green AI” wave [63]. To fully understand models complexity we need to go beyond the analysis reported in this paper, and in literature in general. In particular, as early

pointed out in Sec.V, models size does not directly translate into neither models execution time (as operations differ across models) nor models energy expenditure (which depends on the hardware). Therefore, we argue that further research is needed to relate models accuracy with their computational complexity — this requires to additionally consider metrics such as *classification/Joule* or *accuracy/Watt*.

At the same time, we identify many open research questions regarding the *design* of more efficient DL architectures (bi-dimensional convolutions [17], sparse-LSTM [49], point-wise convolution as in ShuffleNet [50], inverted residual as in MobileNet [64], etc.) and on how to improve their *implementation efficiency* considering traffic processing challenges (as in [65], [66] for the general case, and [19], [67] for traffic classification in particular). As exploring this space is time consuming, automated Neural Architecture Search (NAS) [68] techniques result appealing. However, NAS is guided by an indirect measure of computation complexity (i.e., FLOPs). Therefore, further research is needed to explicitly include in the NAS loop more direct metrics, such as speed or energy consumption, to perform *ecology* of models design space.

3) *Novelty discovery — the broader picture*: The main contribution of this paper is represented by GradBP, a novel DL-tailored technique for zero-day traffic detection. We found that, while DL and ML are equally good for the classification of known traffic, DL-gradient based techniques may give to DL methods a significant advantage over ML methods for the detection of unknown traffic.

We also believe that GradBP opens an interesting avenue for further research. For instance, in our initial design, we experimented by backpropagating to more than one layer. In particular, we found that running a *full backpropagation* renders GradBP significantly slower (about  $4\times$  in our 1-d CNN models) and it additionally reduces accuracy (by almost 10%). While the extra computational cost is somewhat expected, the accuracy drop is more intriguing, and we believe it relates to the known *vanishing gradient* problem: the higher the number of layers traversed with backpropagation, the smaller the gradients, thus the greater the difficulty to measure differences.

Whereas GradBP experimental results show that focusing on the last layer allows to catch large modifications and suffices for zero-day detection of known applications, we argue that the design space for alternatives is still open. For instance, *weighting layers* might reduce the vanishing gradient problem. However, more research is needed to devise simple and effective algorithms to precisely state how layers should be weighted. Additionally, the potential accuracy gains should be contrasted with the extra computational complexity (recall that zero-day detection is additional work to accumulate to the detection of known classes, thus should be minimal).

Finally, while in this paper we only applied our proposed technique for the purpose of zero-day detection, we argue that its application is more general. As such, it would be interesting to evaluate GradBP in the scope of open-set recognition, e.g., by casting its application into computer vision where DL models have initially flourished.

<sup>2</sup>These data challenges, hosted at <http://xianti.fr>, are a new Huawei initiative, to exactly facilitate this type of data sharing through the quite common format of data science challenges popularized by Kaggle <https://kaggle.com>.

4) *Deployability — the elephant in the room*: Given the maturity of traffic classification research, focusing only on raw classification performance, albeit of novel DL models, does not help academic models to step out of academic venues. Instead, it would be imperative to tackle other pressing problems that impact models deployability in the real world [20].

A key aspect of deployability is the *model training loop*. For instance, recalling Fig.1, it is clear that known applications will evolve and eventually some will become unpopular. Similarly, new (zero-day) applications will emerge, for which only a few samples will be initially available. In ML/DL terms, this maps to problems where learning *evolves over time*, such as in *continuous learning* (e.g., to tackle the knowledge drift of existing classes [69]) or in *incremental/decremental learning* (e.g., to add zero-day applications [70] or to remove “old” classes from an existing model). Due to the heavy-tailed nature of the applications popularity, samples of zero-day applications will be scarcer than for known classes. This calls for experimentation with *few-shot learning* [71] techniques to extend models beyond popular classes.

Additionally, deployability raises problems where learning needs to be *distributed over space*, as in *federated learning* (e.g., for privacy or business-sensitive constraints [55], [72], or due to applications heterogeneity [55] and multi-modality as shown in Sec.III-4).

Finally, beside models complexity, deployability also relates to *model inference latency*, especially critical to support line rate analytics, and the *auditing of classification decisions* for the network experts (unlikely experts in the ML/DL domain too). In fact, DL models are known to be harder to interpret [73], [74], which directly translates to the need for *models explainability* [75], an aspect not covered in this work.

In our experience, the ability to seamlessly and automatically incorporate new applications into a model, providing as much insights as possible to domain experts, is paramount for the successful transfer of research into products and deployments.

#### ACKNOWLEDGEMENT

We wish to thank the Editor and the anonymous Reviewers, whose feedback improved the quality of this paper.

#### REFERENCES

- [1] M. Roughan *et al.*, “Class-of-service mapping for qos: a statistical signature-based approach to ip traffic classification,” in *ACM IMC*, 2004.
- [2] A. W. Moore and K. Papagiannaki, “Toward the accurate identification of network applications,” in *Proc. PAM*, 2005.
- [3] L. Bernaille *et al.*, “Traffic classification on the fly,” *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 2, pp. 23–26, 2006.
- [4] M. Crotti *et al.*, “Traffic classification through simple statistical fingerprinting,” *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 1, pp. 5–16, 2007.
- [5] D. Bonfiglio *et al.*, “Revealing skype traffic: when randomness plays with you,” in *Proc. ACM SIGCOMM*, 2007.
- [6] H. Kim *et al.*, “Internet traffic classification demystified: myths, caveats, and the best practices,” in *Proc. ACM CoNEXT*, 2008.
- [7] T. T. Nguyen and G. J. Armitage, “A survey of techniques for internet traffic classification using machine learning,” *IEEE Communications Surveys and Tutorials*, vol. 10, no. 1-4, pp. 56–76, 2008.
- [8] A. Krizhevsky *et al.*, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [9] Z. Wang, “The applications of deep learning on traffic identification,” *BlackHat USA*, 2015.
- [10] V. F. Taylor *et al.*, “Appscanner: Automatic fingerprinting of smartphone apps from encrypted network traffic,” in *Proc. IEEE EuroS&P*, 2016.
- [11] W. Wang *et al.*, “Malware traffic classification using convolutional neural network for representation learning,” in *Proc. IEEE ICOIN*, 2017.
- [12] W. Wang *et al.*, “End-to-end encrypted traffic classification with one-dimensional convolution neural networks,” in *Proc. IEEE ISI*, 2017.
- [13] M. Lopez-Martin *et al.*, “Network traffic classifier with convolutional and recurrent neural networks for internet of things,” *IEEE Access*, vol. 5, pp. 18 042–18 050, 2017.
- [14] Z. Chen *et al.*, “Seq2img: A sequence-to-image based approach towards ip traffic classification using convolutional neural networks,” in *Proc. IEEE BigData*, 2017, pp. 1271–1276.
- [15] L. Vu *et al.*, “A deep learning based method for handling imbalanced problem in network traffic classification,” in *ACM International Symposium on Information and Communication Technology*, 2017.
- [16] G. Aceto *et al.*, “Mobile encrypted traffic classification using deep learning,” in *Proc. IEEE TMA*, 2018.
- [17] T. Shapira and Y. Shavitt, “Flowpic: Encrypted internet traffic classification is as easy as image recognition,” in *IEEE INFOCOM Workshops*, 2019.
- [18] M. Lotfollahi *et al.*, “Deep packet: A novel approach for encrypted traffic classification using deep learning,” *Soft Computing*, vol. 24, no. 3, 2020.
- [19] M. Gallo *et al.*, “Real-time deep learning based traffic analytics,” in *ACM SIGCOMM, Demo session*, Aug. 2020.
- [20] F. Pacheco *et al.*, “Towards the deployment of machine learning solutions in network traffic classification: A systematic survey,” *IEEE Communications Surveys and Tutorials*, pp. 1–1, 2018.
- [21] CISCO, “Cisco application visibility and control user guide,” 2018. [Online]. Available: [https://www.cisco.com/c/en/us/td/docs/ios/solutions\\_docs/avc/guide/avc-user-guide.pdf](https://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/avc/guide/avc-user-guide.pdf)
- [22] “Spectre-6: The system of control and analysis of traffic,” 2019. [Online]. Available: <https://spectredpi.com/wp-content/uploads/2019/07/Spectre-DPI-Brochure.pdf>
- [23] “Traffic perspective application intelligence,” 2020. [Online]. Available: <https://www.sandvine.com/hubfs/downloads/products/packetlogic/sandvine-ds-traffic-perspective.pdf>
- [24] “Service awareness,” 2021. [Online]. Available: <https://support.huawei.com/enterprise/de/doc/EDOC1000012889?section=j00c>
- [25] S. Rezaei *et al.*, “Large-scale mobile app identification using deep learning,” *CoRR*, vol. abs/1910.02350, 2019. [Online]. Available: <http://arxiv.org/abs/1910.02350>
- [26] R. Boutaba *et al.*, “A comprehensive survey on machine learning for networking: evolution, applications and research opportunities,” *Journal of Internet Services and Applications*, vol. 9, no. 1, p. 16, 2018.
- [27] A. W. Moore and D. Zuev, “Internet traffic classification using bayesian analysis techniques,” in *Proc. ACM SIGMETRICS*, 2005.
- [28] G. Vasiliadis *et al.*, “Midea: a multi-parallel intrusion detection architecture,” in *Proc. ACM CCS*, 2011.
- [29] P. S. del Rio *et al.*, “Wire-speed statistical classification of network traffic on commodity hardware,” in *Proc. ACM IMC*, 2012.
- [30] D. Eastlake *et al.*, “Transport layer security (tls) extensions: Extension definitions,” 2011.
- [31] J. Erman *et al.*, “Offline/realtime traffic classification using semi-supervised learning,” *Performance Evaluation*, 2007.
- [32] J. Zhang *et al.*, “Robust network traffic classification,” *IEEE/ACM Transactions on Networking*, vol. 23, no. 04, pp. 1257–1270, jul 2015.
- [33] S. Liang *et al.*, “Principled detection of out-of-distribution examples in neural networks,” *CoRR*, vol. abs/1706.02690, 2017. [Online]. Available: <http://arxiv.org/abs/1706.02690>
- [34] K. Lee *et al.*, “A simple unified framework for detecting out-of-distribution samples and adversarial attacks,” in *Proc. NIPS*, 2018, p. 7167–7177.
- [35] Z. Ge *et al.*, “Generative openmax for multi-class open set classification,” *CoRR*, vol. abs/1707.07418, 2017. [Online]. Available: <http://arxiv.org/abs/1707.07418>
- [36] L. Neal *et al.*, “Open set learning with counterfactual images,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [37] M. Hassen and P. K. Chan, “Learning a neural-network-based representation for open set recognition,” *CoRR*, vol. abs/1802.04365, 2018.
- [38] E. Aljalbout *et al.*, “Clustering with deep learning: Taxonomy and new methods,” *CoRR*, vol. abs/1801.07648, 2018.
- [39] K. G. Dizaji *et al.*, “Deep clustering via joint convolutional autoencoder embedding and relative entropy minimization,” *CoRR*, vol. abs/1704.06327, 2017.

- [40] J. Yang *et al.*, “Joint unsupervised learning of deep representations and image clusters,” *CoRR*, vol. abs/1604.03628, 2016.
- [41] T. DeVries and G. W. Taylor, “Learning confidence for out-of-distribution detection in neural networks,” 2018.
- [42] R. Yoshihashi *et al.*, “Classification-reconstruction learning for open-set recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [43] S. Zhao *et al.*, “Towards unknown traffic identification via embeddings and deep autoencoders,” in *2019 26th International Conference on Telecommunications (ICT)*, 2019, pp. 85–89.
- [44] J. Jang and C. Kim, “One-vs-rest network-based deep probability model for open set recognition,” *ArXiv*, vol. abs/2004.08067, 2020.
- [45] E. M. Rudd *et al.*, “The extreme value machine,” *CoRR*, vol. abs/1506.06112, 2015. [Online]. Available: <http://arxiv.org/abs/1506.06112>
- [46] A. Bendale and T. E. Boulton, “Towards open set deep networks,” *CoRR*, vol. abs/1511.06233, 2015. [Online]. Available: <http://arxiv.org/abs/1511.06233>
- [47] D. Hendrycks and K. Gimpel, “A baseline for detecting misclassified and out-of-distribution examples in neural networks,” *CoRR*, vol. abs/1610.02136, 2016.
- [48] J. Zhang *et al.*, “Autonomous unknown-application filtering and labeling for dl-based traffic classifier update,” in *INFOCOM 2020*, 2020.
- [49] Y. Hua *et al.*, “Deep learning with long short-term memory for time series prediction,” *IEEE Communications Magazine*, vol. 57, no. 6, pp. 114–119, 2019.
- [50] N. Ma *et al.*, “Shufflenet v2: Practical guidelines for efficient cnn architecture design,” in *Proc. of the European conference on computer vision (ECCV)*, 2018, pp. 116–131.
- [51] J. H. Friedman, “Greedy function approximation: a gradient boosting machine,” *Annals of statistics*, pp. 1189–1232, 2001.
- [52] T. Chen and C. Guestrin, “Xgboost: A scalable tree boosting system,” in *Proc. ACM KDD*, 2016.
- [53] D. Ulmer and G. Cinà, “Know your limits: Uncertainty estimation with relu classifiers fails at reliable ood detection,” 2021.
- [54] D. Wen *et al.*, “An overview of data-importance aware radio resource management for edge machine learning,” 2019.
- [55] C. B. Lixuan Yang and D. Rossi, “Heterogeneous data-aware federated learning,” in *Proc. IJCAI Federated Learning workshop*, 2020.
- [56] M. Nasr *et al.*, “Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning,” *IEEE Symposium on Security and Privacy*, 2019.
- [57] Z. C. Lipton and J. Steinhardt, “Troubling trends in machine learning scholarship: Some ml papers suffer from flaws that could mislead the public and stymie future research,” *ACM Queue*, vol. 17, pp. 45–77, 2019.
- [58] M. Hutson, “Eye-catching advances in some ai fields are not real,” *Science Magazine*, May 2020.
- [59] Y. Wang *et al.*, “Generalizing from a few examples: A survey on few-shot learning,” *ACM Computing Surveys (CSUR)*, vol. 53, no. 3, pp. 1–34, 2020.
- [60] G. Hinton *et al.*, “Distilling the knowledge in a neural network,” in *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [61] L. van der Maaten and G. Hinton, “Visualizing data using t-SNE,” *JMLR*, vol. 9, pp. 2579–2605, 2008.
- [62] A. P. Bianzino *et al.*, “A survey of green networking research,” *IEEE Communications Surveys & Tutorials*, vol. 14, no. 1, pp. 3–20, 2010.
- [63] R. Schwartz *et al.*, “Green AI,” *Communications of the ACM*, vol. 63, no. 12, pp. 54–63, 2020.
- [64] A. Howard *et al.*, “Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation,” in *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2018.
- [65] N. P. Jouppi *et al.*, “In-datacenter performance analysis of a tensor processing unit,” *SIGARCH Comput. Archit. News*, vol. 45, no. 2, p. 1–12, Jun. 2017.
- [66] R. David *et al.*, “Tensorflow lite micro: Embedded machine learning for tinyml systems,” *Proceedings of Machine Learning and Systems*, vol. 3, 2021.
- [67] M. Gallo *et al.*, “FENXI: Fast in-network analytics,” in *IEEE/ACM Symposium on Edge Computing (SEC)*, Dec. 2021.
- [68] T. Elsken *et al.*, “Neural architecture search: A survey,” *The Journal of Machine Learning Research*, vol. 20, no. 1, pp. 1997–2017, 2019.
- [69] V. Carela-Español *et al.*, “A streaming flow-based technique for traffic classification applied to 12+ 1 years of internet traffic,” *Telecommunication Systems*, vol. 63, no. 2, pp. 191–204, 2016.
- [70] G. Bovenzi *et al.*, “A first look at class incremental learning in deep learning mobile traffic,” in *IFIP Traffic Monitoring and Analysis (TMA)*, September 2021.
- [71] Y. Wang *et al.*, “Generalizing from a few examples: A survey on few-shot learning,” *ACM Computing Surveys (CSUR)*, vol. 53, no. 3, pp. 1–34, 2020.
- [72] J. Konečný *et al.*, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [73] C. Beliard *et al.*, “Opening the deep pandora box: Explainable traffic classification,” in *Proc. IEEE INFOCOM, Demo session*, 2020.
- [74] Z. Meng *et al.*, “Interpreting deep learning-based networking systems,” in *Proc. ACM SIGCOMM*, 2020.
- [75] A. B. Arrieta *et al.*, “Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai,” *Elsevier Information Fusion*, vol. 58, pp. 82–115, 2020.



**Lixuan Yang** is a Senior Engineer in the DataCom Lab at Huawei Technologies, France. She received her Ph.D. from the CNAM (2017) and MSc from Jean Monet University(2013). She worked at the XXII Group as an artificial Intelligence team leader on computer vision and natural language processing projects. Her current interests include traffic classification, federated learning, continual learning.



**Alessandro Finamore** is a Principal engineer in the DataCom Lab at Huawei Technologies, France. Before joining Huawei in 2019, he was a Principal engineer at Telefonica UK/O2 (London, United Kingdom), and an associate research at Telefonica Research (Barcelona, Spain). He received his Ph.D from Politecnico di Torino (2012). He has coauthored more than 80 papers in leading conferences and journals, including a best paper award at CoNEXT.



**Jun Feng** Jun is a data scientist and AI architect expert in Huawei Data Communication Product Line. He has over 20 years of experience in data analysis and AI deployment in the telecommunications industry.



**Dario Rossi** is Network AI CTO and Director of the DataCom Lab at Huawei Technologies, France. Before joining Huawei in 2018, he held Full Professor positions at Telecom Paris and Ecole Polytechnique and was holder of Cisco’s Chair NewNet Paris. He has coauthored 15 patents and over 200 papers in leading conferences and journals, that received 9 best paper awards, a Google Faculty Research Award (2015) and an IRTF Applied Network Research Prize (2016). He is a Senior Member of IEEE and ACM.