ML-based tunnel detection and tunneled application classification

Johan Mazel   Matthieu Saudrais   Antoine Hervieu
ANSSI

Abstract

Encrypted tunneling protocols are widely used. Beyond business and personal uses, malicious actors also deploy tunneling to hinder the detection of Command and Control and data exfiltration. A common approach to maintain visibility on tunneling is to rely on network traffic metadata and machine learning to analyze tunnel occurrence without actually decrypting data. Existing work that address tunneling protocols however exhibit several weaknesses: their goal is to detect application inside tunnels and not tunnel identification, they exhibit limited protocol coverage (e.g. OpenVPN and Wireguard are not addressed), and both inconsistent features and diverse machine learning techniques which makes performance comparison difficult.

Our work makes four contributions that address these limitations and provide further analysis. First, we address OpenVPN and Wireguard. Second, we propose a complete pipeline to detect and classify tunneling protocols and tunneled applications. Third, we present a thorough analysis of the performance of both network traffic metadata features and machine learning techniques. Fourth, we provide a novel analysis of domain generalization regarding background untunneled traffic, and, both domain generalization and adversarial learning regarding Maximum Transmission Unit (MTU).

1 Introduction

Off-the-shelf encrypted tunneling protocols such as SSH or OpenVPN provide a reliable way for users to protect their network traffic from passive analysis. Companies often deploy such tools between locations or to access remote infrastructure. General public also use tunneling to avoid passive monitoring or tampering [17]. These same tools are also used by attackers to hide C2 (Command and Control) and data exfiltration. Although many tools already exists, such as IPsec, OpenVPN or SSH, new ones also recently appeared (e.g. Wireguard). The encrypted tunneling protocol landscape thus keep getting more complex and makes malicious actors monitoring more difficult.

Port-based heuristic relying on tunnel default configuration are easy to circumvent by attackers. Supervised machine learning (ML) applied to network traffic metadata (usually byte- and time-based) is thus a natural port-agnostic approach to identify some protocols without actually inspecting encrypted data. Several work use this approach to identify applications inside tunneling tools such as SSL/TLS [5,6,8,11], both IPsec and PPTP [9,10,12], and SSH [5,7]. Their tunneling coverage is however limited: they do not address common protocols such as OpenVPN or Wireguard. These work also use non-overlapping network traffic metadata feature sets. Furthermore, they use different machine learning techniques. These last two aspects makes it difficult to compare their results. Although machine learning is now easier to use than ever thanks to libraries such as Scikit-learn [18], clear guidelines on how to use machine learning to analyze tunneling protocols are missing. Beyond application classification inside tunnel, up to our knowledge, the preliminary task of identifying tunnels remains unaddressed.

Our goal is to address these limitations and design an approach to detect and classify existing tunneling protocols. We also want to provide a detailed analysis of network traffic metadata features and machine learning techniques that provides clear and actionable guidelines for real-world deployment. Finally, we want to address domain generalization as it is paramount to assess the whether a previously learned model will be usable beyond its initial training context. We here address two aspects: background untunneled traffic and Maximum Transmission Unit (MTU).

Our work provide four main contributions. First, we include the two tunneling protocols that are not addressed in existing work: OpenVPN and Wireguard. Second, we design a complete pipeline to detect and classify tunneling protocols and tunneled applications. Third, we provide a thorough comparisons of network traffic metadata features and machine learning techniques in the context of tunneling traffic detection and classification. Fourth, we address domain generalization regarding background untunneled traffic, and then, both domain generalization and adversarial learning regarding

	Authors	Refs.	Year	Tunnels					MM methods							
				SSL/TLS	Tor	SSH	IPSec	PPTP	Decision Tree	K-Nearest Neighbor	Naive Bayes	Support Vector	Gaussian Mixture Model	K-Means	Markov Chains	Neural Networks
Protocol classification	Bernaille et al.	[1]	2006	-	-	-	-	-			x					
	Bernaille et al.	[2]	2006	-	-	-	-	-		x			x		x	
	Kim et al.	[3]	2008	-	-	-	-	-	x	x	x	x				x
	Jaber et al.	[4]	2011	-	-	-	-	-						x		
Protocol classification inside tunnel	Wright et al.	[5]	2006	x		x				x					x	
	Bernaille et al.	[6]	2007	x									x			
	Maiolini et al.	[7]	2009			x								x		
	Sun et al.	[8]	2010	x	x						x					
	Okada et al.	[9]	2011				x	x			x					
	Okada et al.	[10]	2011				x	x	x		x	x				
	Korczyński et al.	[11]	2014	x											x	
	Kumano et al.	[12]	2014				x	x	x			x				

Table 1: Tunnelling procotocols and machine learning techniques used in previous work.

MTU. Up to our knowledge, both of these aspects have never been addressed before.

Our paper is structured as follows. Section 2 present existing work on tunneling protocols analysis. Section 3 details our methodology. Section 4 presents our detection and classification results.

2 Related work

The idea of using network traffic metadata can be traced back to the original work of Wright et al. [5] who applied it to application classification inside SSL/TLS and SSH tunnels. Several works extend this initial contribution regarding other tunneling tools such as Tor [8] or IPsec and PPTP [9, 10, 12]. Some works revisit previously addressed protocols but are relevant considering the quick evolution of the tunneling and encryption landscapes [6, 7, 11].

Another use case of network traffic metadata that generated many contributions is protocol classification. Bernaille et al. [1] is the seminal work for this use case. Several other works then improve on this initial contribution [2–4].

Table 1 describes the targeted tunneling tools and ML techniques used. ML techniques and targeted tunneling protocols widely vary across studies which make their results extremely difficult to compare. These work also do not address tunneling tools that are common such OpenVPN or Wireguard. Table 2 describes features used in existing work and tools. They are very diverse. The first papers on protocol classification inside tunnel used N first packet size or IAT, while later work usually leverage five tuple flow-related feature such as total number of packet or byte exchanged. Features that were more recently proposed such as bursts [19] have also not been evaluated in the context of tunneling protocol analysis.

Up to our knowledge, there is no existing work on the preliminary steps to application classification inside tunnels which are tunnelling protocol detection and/or classification.

Our goal is thus to 1) provide a complete pipeline for the analysis of tunelling protocol (detection, classification, and application classification within tunnels), and 2) address the diversity of existing work by providing a unified point of view on performance of both network traffic features, and machine learning techniques.

3 Methodology

In this section, we present our approach. First, we address the general machine learning-based pipeline in Section 3.1, and
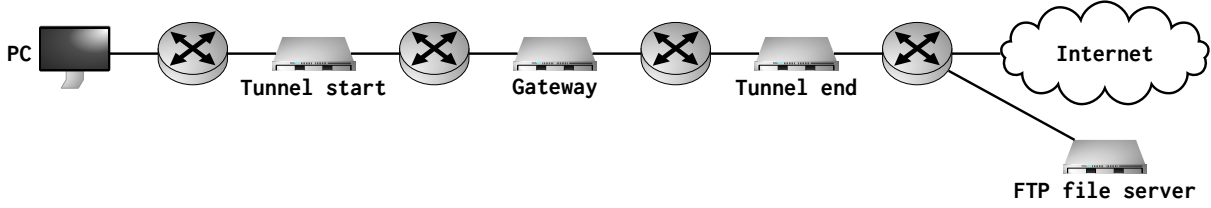
| Type | Authors | Ref. | Year | Flow | | | | | N first | | | | | |
				# packet	Packet size	Duration	IAT	TCP-based (eg. flag)	TLS type	Direction	Packet size	Elapsed time	IAT	Packet/byte burst
Protocol classification	Bernaille et al.	[1]	2006								x			
	Bernaille et al.	[2]	2006								x/$\mathcal{D}$			
	Kim et al.	[3]	2008	x	{T,m,M,μ,σ}/$\mathcal{B}$	x	{T,m,M,μ,σ}/$\mathcal{B}$	x			x			
	Jabber et al.	[4]	2011											x
Protocol classification inside tunnel	Wrigh et al.	[5]	2006								x/$\mathcal{D}$			x
	Bernaille et al.	[6]	2007								x/$\mathcal{D}$			
	Maiolini et al.	[7]	2009							x	x	x		
	Sun et al.	[8]	2010	x	μ,m,M		μ,m,M							
	Okada et al.	[9]	2011	x/$\mathcal{A}$	{T,m,M,$\mathcal{P}$,μ,σ^2}/$\mathcal{A}$	x	{m,M,$\mathcal{P}$,μ,σ^2}/$\mathcal{A}$							
	Okada et al.	[10]	2011	x/$\mathcal{A}$	{T,m,M,$\mathcal{P}$,μ,σ^2}/$\mathcal{A}$	x	{m,M,$\mathcal{P}$,μ,σ^2}/$\mathcal{A}$							
	Korczyński et al.	[11]	2014					x						
	Kumano et al.	[12]	2014	x/$\mathcal{A}$	{T,m,M,$\mathcal{P}$,μ,σ^2}/$\mathcal{A}$	x	{m,M,$\mathcal{P}$,μ,σ^2}/$\mathcal{A}$							
	Netflow v5	[13]		x	T,μ	x	(μ)							
	Netflow v9	[14]		x/$\mathcal{A}$	T/$\mathcal{A}$,{m,M}/I	x								
	Argus	[15]		x/$\mathcal{A}$	T/$\mathcal{A}$	x	m							
	Zeek (conn.log)	[16]		x/$\mathcal{A}$	T/$\mathcal{A}$	x	m							
	Our work			x/$\mathcal{A}$	{T,m,M,μ,σ}/$\mathcal{A}$	x	{m,M,μ,σ}/$\mathcal{A}$			x	x/$\mathcal{D}$	x	x	x/$\mathcal{D}$

Table 2: Network traffic features used for protocol classification and tunnel protocol identification. Abbreviations: x: used, $\mathcal{D}$: sign encodes packet direction ; μ: mean, σ: standard deviation, σ^2: variance ; $\mathcal{B}$: uni-directional ; I: incoming only ; $\mathcal{A}$: all direction ; T: total ; m: minimum ; M: maximum ; $\mathcal{P}$: 25^{th} centile, 50^{th} centile and 75^{th} centile.

Figure 1: Processing pipeline. Functional boxes have a purple background and rounded corners. Data boxes have a white background and square corners.

then, we detail the main steps: traffic generation in Section 3.2, used datasets in Section 3.3, feature extraction in Section 3.4, and, machine learning in Section 3.5.

3.1 General pipeline

Figure 1 picture the whole pipeline that we use to analyze tunneling protocols. First, we generate network traffic (see

Figure 2: Traffic generation topology.

Application	Tool	Target	Target nb	Max application run	Max target nb
Web	Firefox/browsh	Alexa Top	1M	20	1
wget	wget	Debian packages subset	42669	4	1
FTP get	lftp/vsftp	Debian packages subset	1000	20	10
FTP put	lftp/vsftp	Debian packages subset	1000	20	10

Table 3: Deployed applications. Max application run is the maximum number of times an application is run inside a tunnel. Max target nb is the maximum number of target (website or debian package) retrieved by each application run. In both cases, the actual number is randomly picked between 1 and this maximum value. Timewise, browsing a website lasts between 30 and 45s and file downloads are consecutive.
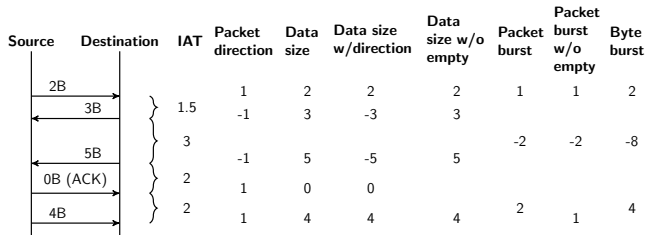
Figure 3: Example of used N first feature.

Section 3.2) to complete the external datasets that we obtained (see Section 3.3). Then, using generated PCAP files, we extract network traffic metadata feature (see Section 3.4). The next step is the use of ML to determine if a given flow is a tunnel. This task is named *tunnel detection* in the remainder of the paper. If it is the case, we then identify which tunnel is present. This step is called *tunnel classification*. If a tunnel is present, we classify the application used inside the tunnel. We call this phase *application classification inside tunnel*.

3.2 Traffic generation

Although some existing dataset provide some tunneling protocol network traffic (eg. SSH in UNIBS [20]), we could not find any data for IPsec, OpenVPN or Wireguard. We thus design a generic topology to collect tunnel traffic. Here, a host uses a tunnel to access some resources on Internet or on an FTP file server. This topology is presented on Figure 2. The tunnel is setup between two hosts (tunnel start and tunnel end) and is routed through a machine called gateway. This topology is setup using Vagrant with the libvirt provider. All

			Tunnels					
Source	Untunneled	SSH	OpenVPN-TCP	OpenVPN-UDP	IPSec	Wireguard		Total
UN-B [20]	66,585	27	-	-	-	-		66,612
UN-S [20]	0	27,978	-	-	-	-		27,978
UPC [21, 22]	525,537	4,093	-	-	-	-		529,630
This work	241,941	2,400	2,400	2,400	2,400	2,400		253,941

Table 4: Data used in our work. UN-B (resp. UN-S) is the UNIBS dataset with untunneled (resp. SSH) traffic.

hosts use Debian Testing with the version 20210228.1 [1] (this is equivalent to Debian Bullseye). We deactivate NIC offloading on all machines in order to avoid packet size alteration during capture.

We deploy each tunnel between the tunnel start and tunnel end machines. We setup the following tunnels: IPsec, OpenVPN TCP/UDP, SSH and Wireguard.

We use several application inside each tunnel: web browsing, FTP get/put and wget. We actually use a single application inside each tunnel instance. Table 3 provides additional details regarding applications used inside tunnels.

---

[1] https://app.vagrantup.com/debian/boxes/testing64

4

Algorithms	Scikit learn function	Short name	Parameter name	Parameter default value	Parameter values for grid search
AdaBoost	[23]	AB	learning rate	1.0	$10^{-3,-2,-1,0,1}$
Gradient boosting	[24]	GB	max depth	3	1, 2, 3, 4, 5, 6, 7, 10, 20, 50
Gaussian naive Bayes	[25]	GNB	-	-	-
K nearest neighbors	[26]	KNN	n neighbors	5	1, 5, 10, 50
Support vector with liblinear	[27]	SV LL	C	1.0	$10^{-1,0,1,2,3,4,5,6,7}$
Logistic regression	[28]	LR	solver C	saga 1.0	saga $10^{-1,0,1,2,3,4,5,6,7}$
Logistic regression with SGD	[29]	LR SGD	loss alpha	log 10^{-4}	log $10^{-8,-7,-6,-5,-4,-3,-2,-1,0,1}$
Support vector with SGD	[29]	SV SGD	loss alpha	hinge 10^{-4}	hinge $10^{-8,-7,-6,-5,-4,-3,-2,-1,0,1}$
Random forest and decision tree	[30] [31]	RF DT	criterion min samples split	gini 2	gini, entropy 2,3,4,5,10,50,100

Table 5: Machine learning algorithms used and parameters explored. Random forest and decision tree both use the criterion and min samples split parameters. Short names are used to identify algorithms in the rest of the paper.

3.3 Datasets

Table 4 present the data we use in this work. The UNIBS datasets [20] were collected in 2009 and contain both SSH tunneled and untunneled traffic. The UPC dataset [22] was gathered in 2014 and contains mostly untunneled traffic and SSH and Tor tunnels. In this work, we only use the untunneled and SSH tunneled traffic. The network traffic generation in Section 3.2 create 100 flows for six MTU values (see Section 4.5 for additional details on MTU choice) and four applications used with 5 tunnels and without tunnel. We thus obtain 2400 flows for each tunneling protocol. Untunneled traffic does not yield the same number of flows for each MTU due to varying interactions with third-party during website browsing.

3.4 Feature extraction

Once network traffic has been generated, we extract five tuple flow (source and destination IP, transport protocol, and source and destination port) and their associated features. We extract the following features about each five tuple flow: transport protocol (e.g. TCP, UDP), packet number, duration, and, mean, minimum, maximum, standard deviation in both direction and the total of both packet size and Inter-Arrival Times (IAT). We also extract the following N first feature (see example in Figure 3): elapsed time since the start of the flow, IAT, packet direction, packet size, packet size with direction in

sign (called packet size with direction in the remainder of the paper), packet burst and byte burst. Burst are build by grouping consecutive packet in the same direction. Packet bursts are the grouping sizes and byte bursts are the grouping total byte number. We also build packet size from source and destination only, and byte burst from source and destination only. The Table 2 presents our features and those of previous work.

3.5 Machine learning

We use algorithms available in scikit-learn [18] for multiclass classification: random forest, decision tree, AdaBoost, logistic regression and support vector with SGD, logistic regression with the saga solver, Gaussian naive Bayes, and K nearest neighbors. Some algorithms are omitted in the following experiments due to excessive running time (e.g. K nearest neighbors for tunnel detection).

In terms of performance metric, we use the F1 score for tunnel detection and the average F1 score across classes (macro in scikit learn) for tunnel classification and application classification inside tunnels.

Performance comparison between algorithms is performed using a nested cross-validation [32]. The inner loop uses a grid search to find the best parameters for each algorithm

(see Table 5). The best parameters are determined using F1 score as specified above. The outer loop then reuse these optimal parameters to provide a performance lower bound for comparison with other algorithms.

4 Results

In this section, we present the performance evaluation for all three steps defined in Section 3.1. We first present a study of the impact of N where N is the number N first feature used, e.g. packet size (see Section 4.1 and Figure 3). We then compare N first features inside feature families (here packet direction/size-related and burst-related) in Section 4.2.1. Next, we reuse the best N first features inside each family and compare their performance in Section 4.2.2. We then provide a performance evaluation of these best features used together along with feature importance analysis (see Section 4.3). Finally, we address domain generalization regarding change of background network traffic in Section 4.4, and both domain generalization and adversarial learning regarding MTU in Section 4.5.

We only present the results of application classification inside tunnel for the SSH tunnel due to the lack of space. This tunnel however exhibits the worst performance across all experiments. Presented results thus always provide a performance lower bound for application classification inside tunnels.

In the next sections, we often use a subset of all extracted features (see Section 3.4) to compare them. We however always use the one-hot encoded transport protocol (TCP or UDP) as a additional feature to any of these feature subsets.

4.1 N first feature

In this section we analyze the impact of N on performance, where is N the number of N first feature used. N first features are: packet direction, packet size, packet size with direction encoded in sign, packet burst, byte burst, and inter-arrival time. We do not present elapsed time since flow start due to the lack of space and because its performances are always lower than IAT. Due to the computational cost of testing mutliple values of N, we do not perform a grid search to find optimal parameters inside the inner loop, and simply use default parameters from Table 5. We address all pipeline steps: tunnel detection, tunnel classification, and application classification.

4.1.1 Tunnel detection and classification

Figure 4 (resp. Figure 5) presents the F1 score obtained with N values between 1 and 50 for tunnel detection (resp. classifi-

cation). For both tasks, we use N between 1 and 50 because N values above 20 do not provide significant performance improvement for the best algorithm (random forest).

Algorithms exhibits widely varying performance. Random forest is the best one. When N increases, decision tree and random forest exhibit increasing performance for all features. Other algorithms' performances sometimes decrease when N increases (e.g. N between 20 and 25 with packet direction and Gaussian naive Bayes for tunnel detection on Figure 4).

We then compare feature performance using the F1 score for N equal to 50 with the best performing algorithms (decision tree and random forest). Packet size, packet size with direction, byte burst and IAT provide very good overall results. Packet direction and burst exhibit slightly lower performance.

In terms of smallest N to reach maximum performance, ie smallest possible memory usage to reach optimal performance, byte burst is the best feature, followed by packet size with direction and IAT, then packet size, and finally packet burst and packet direction. Byte burst is the most efficient feature in terms of memory usage. It however does not mean that byte burst is quick time-wise to reach its best performance level because burst position in a flow is not directly linked to packet position. Indeed, the fourth byte burst may actually be located at the 30th packet which would make byte burst slower than packet size with direction.

When one compares tunnel detection and tunnel classification, the smallest N values to reach maximum performance are similar. Algorithm performances are however better and vary less across algorithms for tunnel classification than for tunnel detection.

4.1.2 Application classification

Figure 4 presents the F1 score obtained with N values between 1 and 150 for application classification inside SSH tunnels. We use N between 1 and 150 because N values above 120 do not provides significant performance improvement for the best algorithm (random forest). We do not present result for other tunnels due to the lack of space.

Results are similar to the tunnel detection and classification use cases. Algorithm performance are diverse. Random forest exhibits the best F1 scores, and its performances increase when N increases for all features. This is not the case for other algorithms (see N between 100 and 150 with packet direction and Gaussian naive Bayes). Overall, byte burst is the best feature. Packet size, packet size with direction, packet burst are close second. Packet direction exhibits the worst performance. In terms of smallest N to reach maximum performance, byte burst is the best feature, closely followed by packet burst, then packet direction, packet size with direction, packet direction and IAT, in that order.
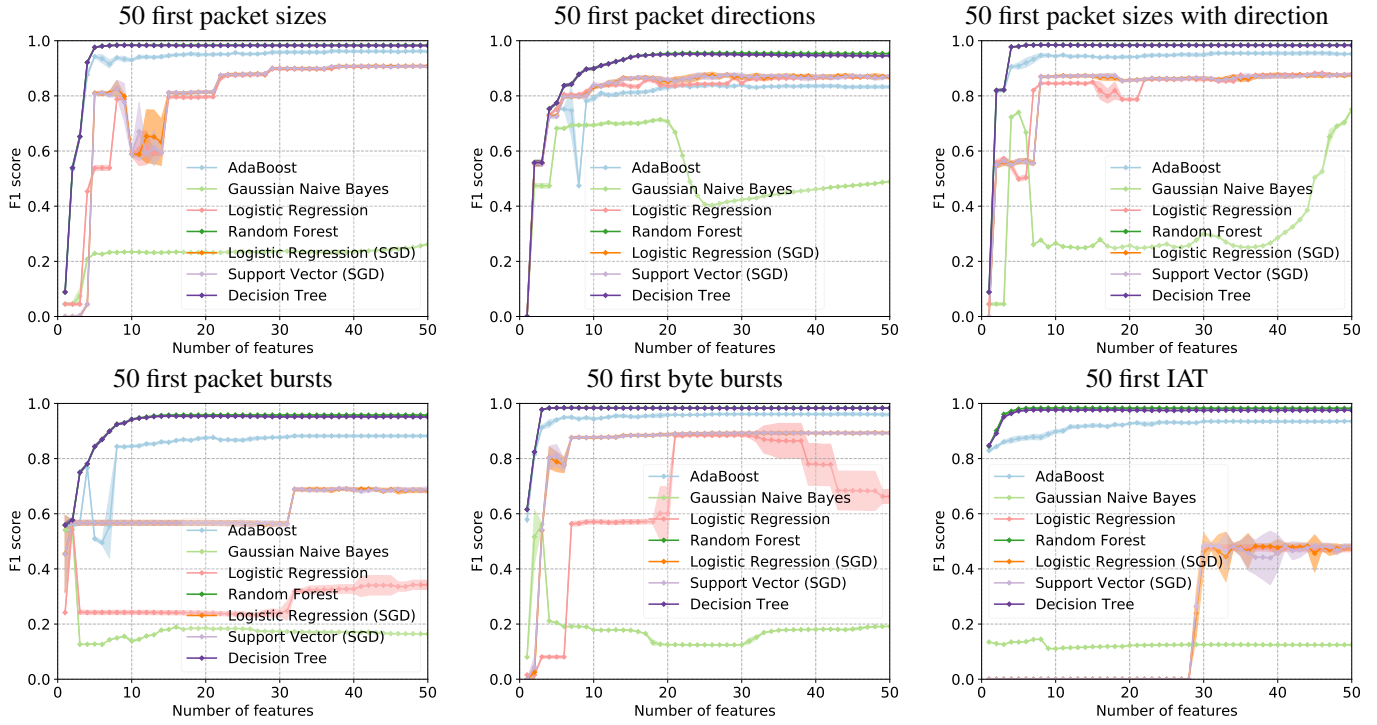
Figure 4: Tunnel detection using 50 first feature. Shaded areas below and above the curve represent a confidence interval with 99% confidence level.
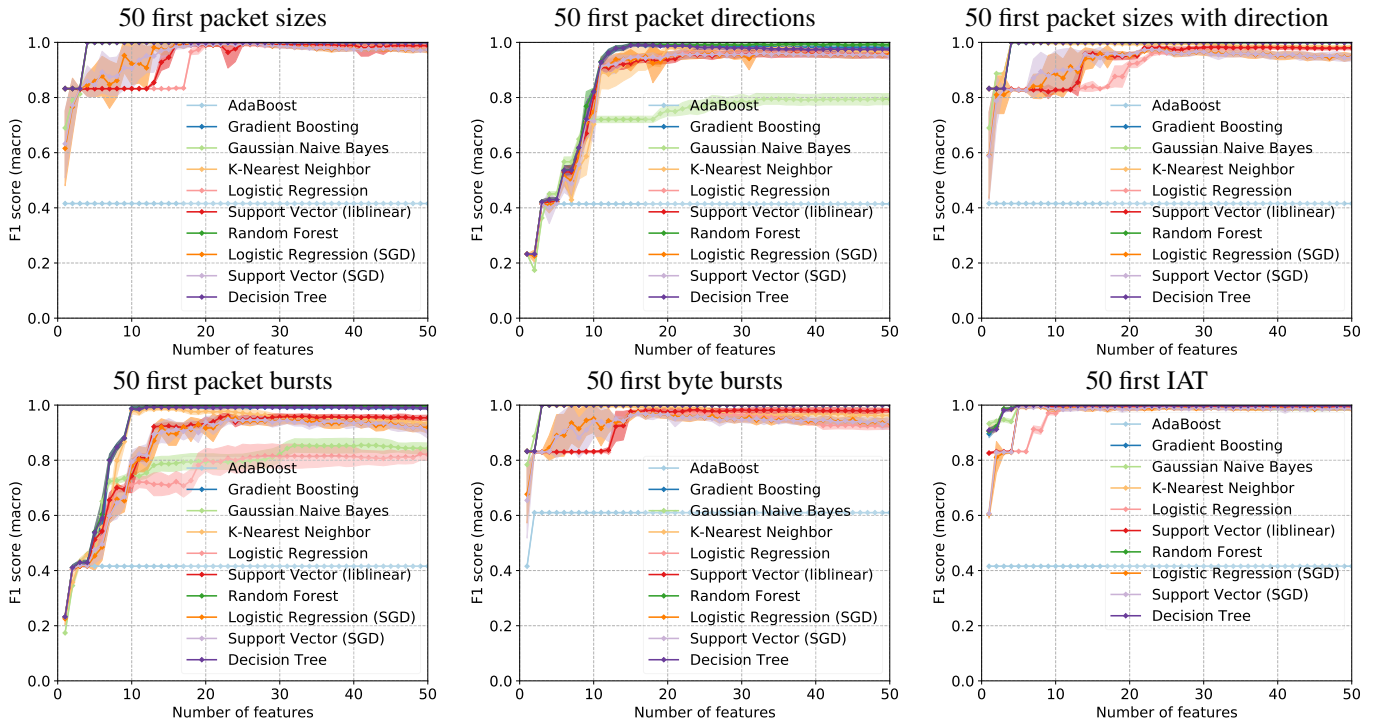

Figure 5: Tunnel classification using 50 first feature. Shaded areas below and above the curve represent a confidence interval with 99% confidence level.
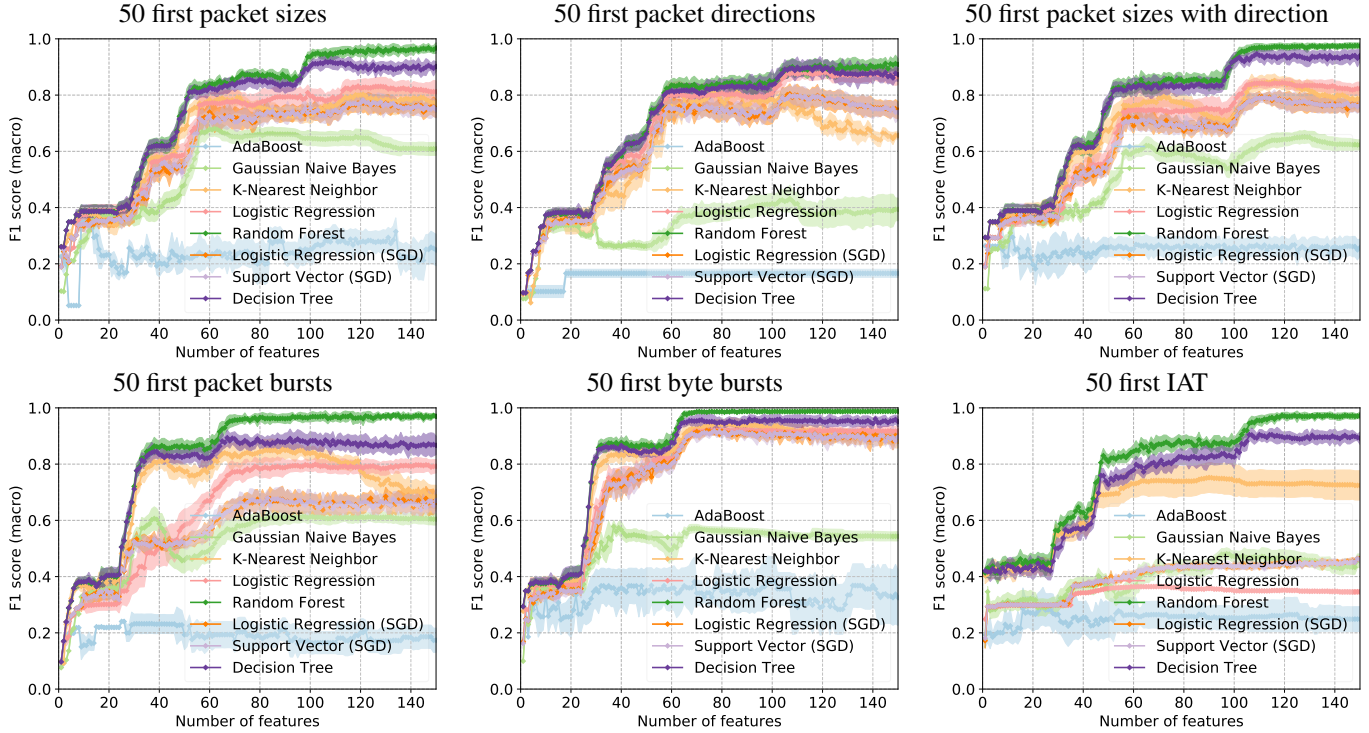
Figure 6: Application classification inside SSH using 150 first features. Shaded areas below and above the curve represent a confidence interval with 99% confidence level.

4.1.3 Summary

In this preliminary comparison, we determine that using 50 (resp. 150) features is enough for tunnel detection and tunnel classification (resp. application classification). Overall, random forest is the best algorithm. Byte burst is the best feature in terms of smallest N to reach the best performance.

4.2 Features

In this section, we compare features: first, among feature families (e.g. packet size vs packet size with direction or packet burst vs byte burst), then across families using the best features from each families (e.g. packet size with direction vs byte burst).

4.2.1 Preliminary feature comparison

We here analyze performance inside each feature families: packet-related (with packet size, packet direction, packet size from source, packet size from destination, packet size with direction in sign), and burst-related (packet burst, byte burst, byte burst from source, byte burst from destination).

Packet direction, packet size from source, from destination, from both source and destination, and packet size

with direction Figure 7 pictures the F1 scores obtained with ML algorithms for packet direction and feature related to packet size. F1 scores obtained with packet size from source are always better than packet size from destination. Packet size from source is also always better than packet size, except for tunnel detection with algorithms such as logistic regression, support vector and AdaBoost. Packet direction is almost always worse than packet size from source, except for application classification in SSH with AdaBoost. Packet size with direction and packet size from source exhibits very similar performances. In the remainder of the paper, we arbitrarily choose to only consider packet size with direction.

Packet burst and byte burst from source/destination/both Figure 8 pictures the F1 scores obtained with ML algorithms for features related to bursts. F1 scores obtained with packet burst are always worse than byte burst (except for SSH application classification with Gaussian naive Bayes). Byte burst from source and destination does not yield better results than byte burst.

Byte burst bigram and trigram We now explore the use of bigram and trigram on byte burst values. We do not use a figure due to the lack space but one is provided in appendix (see **??**).
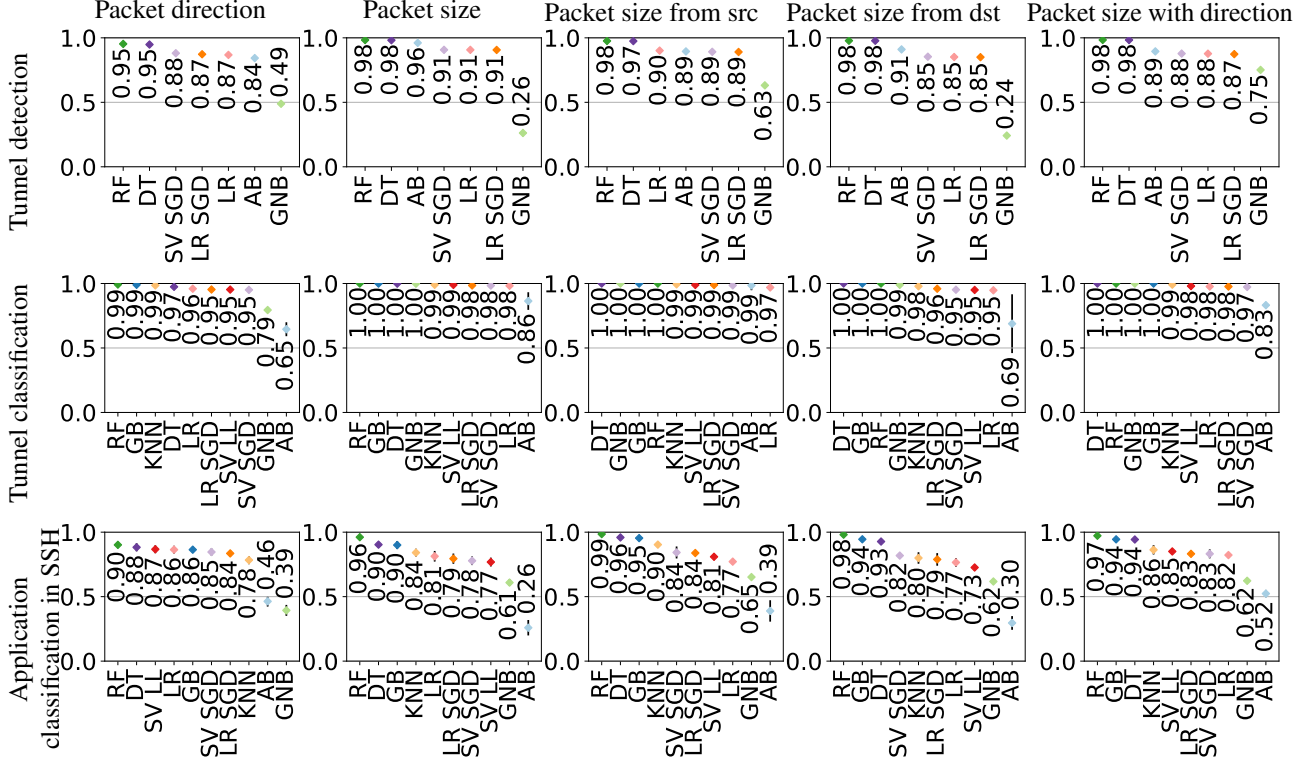
Figure 7: Comparison of packet size, packet size from source, packet size from destination, packet direction and packet size with direction for all pipeline steps. Error bars represent a confidence interval with 99% confidence level.

Byte burst bigram used alone and in combination with byte burst does not improve results, except for IPsec-ESP with Gaussian Naive Bayes and OpenVPN TCP with AdaBoost and Decision Tree. Byte burst trigram used alone and in combination) does not improve results (except for IPsec-ESP with all methods and OpenVPN TCP with AdaBoost and Decision Tree). Considering the limited performance improvements and costly feature generation, we do not use byte burst bigrams or trigrams in the remainder of the paper.

4.2.2 Feature comparison

Next we compare features selected from the first two preliminary comparisons, packet size with direction and byte burst, with others features such as N first IAT and Netflow features (v5 and v9) on Figure 9. Netflow features always yield worse results than any other features. IAT are worse than packet size with direction and byte burst, except for tunnel classification and some ML algorithms such as both logistic regressions. We do not present results for elapsed time since the start of the flow because its results are worse than IAT. Previous work [33] show that tampering with IAT and elapsed time since the flow start using a proxy yields a performance decrease. We thus only consider packet size with direction and

byte burst as relevant features for the remainder of the paper.

4.2.3 Summary

Overall, packet size with direction and byte burst offer the best performance. We emphasize that Netflow v5 and v9 feature sets do not provide good performances.

4.3 General results

We now provide a performance evaluation of packet size with direction and byte burst used together. Figure 10 pictures the F1 scores obtained for tunnel detection and classification. We observe very good performance for both tunnel detection and tunnel classification using the best algorithms (e.g. random forest or decision tree). Figure 11 displays the 10 most feature importance based on Mean Decreased Impurity (MDI) for tunnel detection. Overall, IP protocol field-related one-hot-encoded features are not relevant. This is consistent with the fact that tunnels here use both TCP and UDP. 16 out of the 20 most important feature are related to byte burst. Byte bursts are thus here more important than packet sizes with direction. Figure 12 displays the 10 most feature importance based on Mean Decreased Impurity (MDI) for tunnel classification. IP protocol field-related one-hot-encoded features
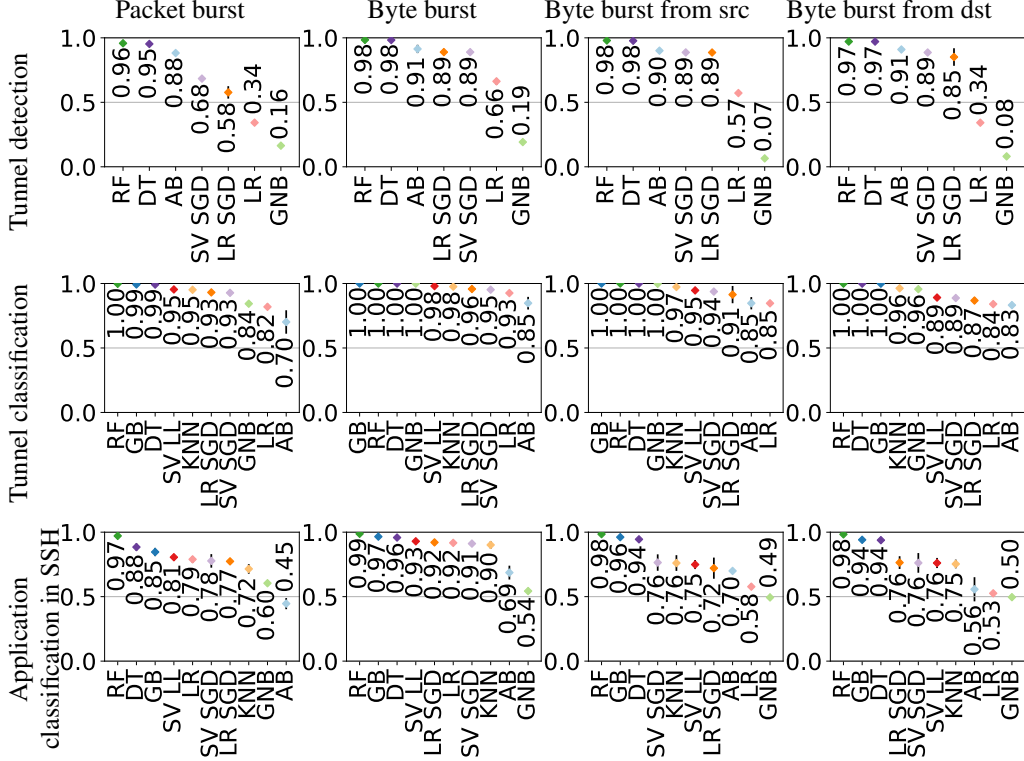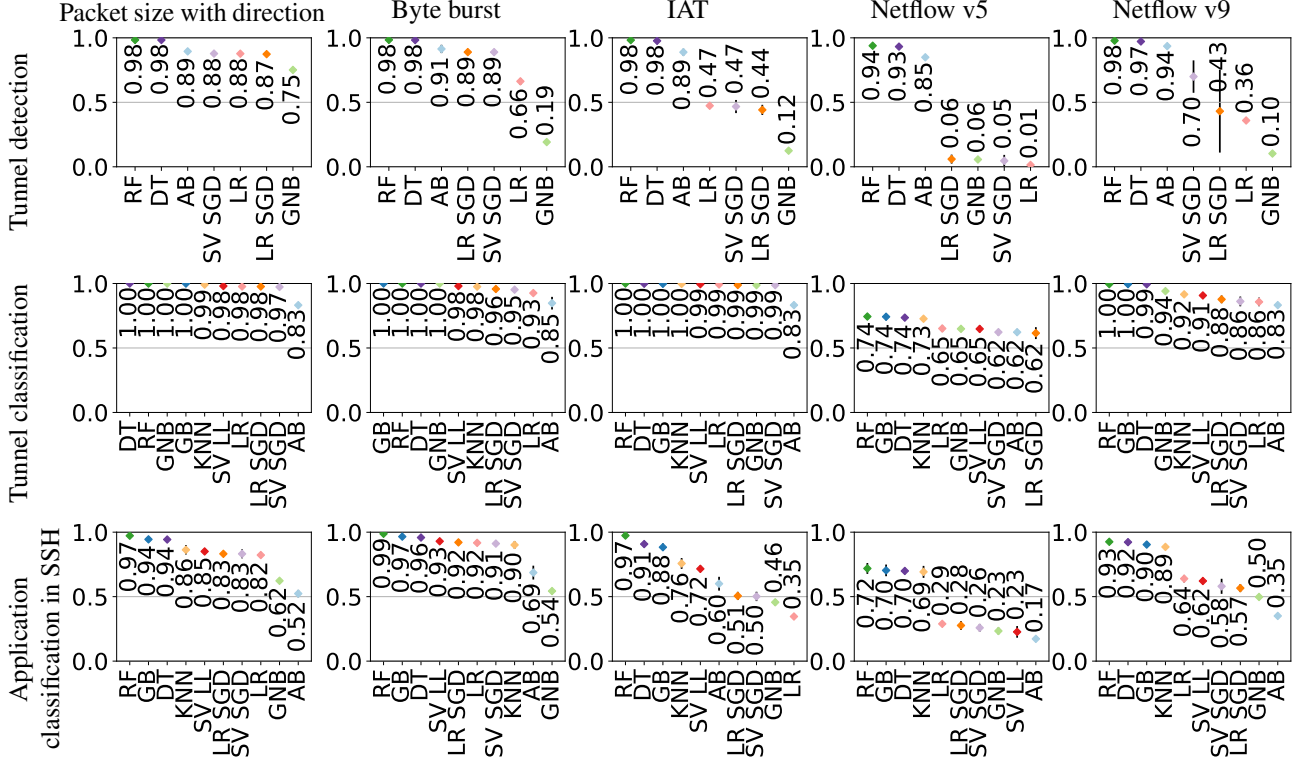
Figure 8: Comparison of packet burst, byte burst, byte burst from source, and byte burst from destination for all pipeline steps. Error bars represent a confidence interval with 99% confidence level.

are here the first and fourth most important features. Beyond these two attributes, packet size with direction and byte burst are similarly important.

Figure 13 pictures the F1 scores obtained for tunnel detection and classification. The most difficult tunneling protocol to classify application inside is SSH, followed by both OpenVPN tunnels, and finally, IPsec and Wireguard. Similarly to Figure 10, using the best algorithms, random forest or decision tree, provides very good performances.

4.4 Domain generalization regarding untunneled traffic

We here address the impact of the untunneled traffic on tunnel detection performance. We here only consider tunnel detection because untunneled traffic is only used in this step of the ML pipeline (see Figure 1). To this end, we use the UNIBS dataset for training and the UPC dataset for testing, and then, the other way around. We here do not use tunneling protocol-related flows from the UNIBS and UPC datasets. Table 6 present F1 scores for packet sizes with direction, byte bursts, and Netflow v5 and Netflow v9 features. Although we previously show in Section 4.2.2 that Netflow features exhibit poor performances, we still include them here as classic flow-based

Table 6: Domain generalization performance between the two background traffic datasets regarding F1 score using random forest for tunnel detection.

		Test	
	Feature	UNIBS	UPC
UNIBS	Netflow v5		0.28±0.00
	Netflow v9		0.37±0.02
	Byte burst		0.61±0.03
	Packet size w/ dir.		0.75±0.02
UPC	Netflow v5	0.54±0.01	
	Netflow v9	0.74±0.01	
	Byte burst	0.98±0.00	
	Packet size w/ dir.	0.99±0.00	

features. When we train a random forest model on the UPC (resp. UNIBS) dataset and test it on the UNIBS (resp. UPC) dataset using Netflow v5, we obtain a 0.54 (0.28) F1 score. These results are very poor, especially compared to what we previously obtained in Sections 4.2.2 and 4.3. Netflow v9 exhibit the second worst performance decrease, followed by byte burst and then, packet size with direction. We hypothesize that training with UPC yield better results than training with

10

Figure 9: Comparison of packet size with direction, byte burst and IAT for all pipeline steps. Error bars represent a confidence interval with 99% confidence level.
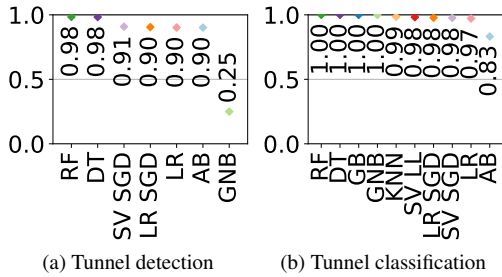

Figure 10: Tunnel detection and tunnel classification using selected features (open world). Error bars represent a confidence interval with 99% confidence level.
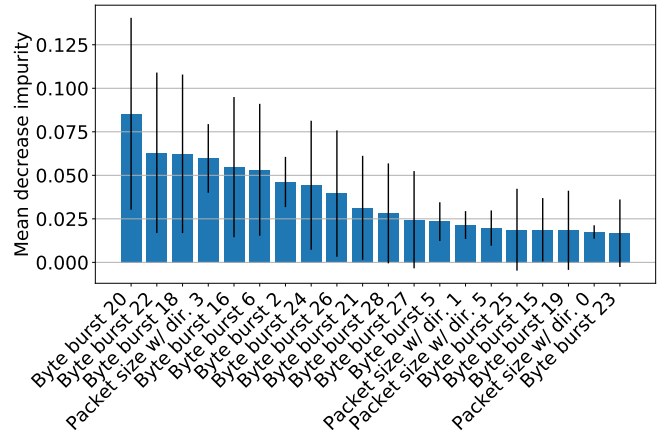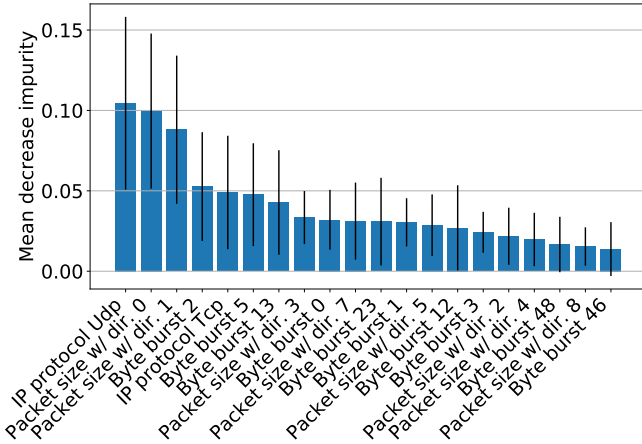

Figure 11: 20 most important feature regarding Mean Decreased Impurity (MDI) using Random Forest with 100 trees for tunnel detection. Error bars represent a confidence interval with 99% confidence level.

UNIBS because of two reasons. First, the number of instances in the UPC dataset is much bigger than in the UNIBS one (see Table 4) which helps generalization of models trained with UPC data. Second, UPC dataset is more recent than UNIBS (2014 vs 2009). This means that traffic categories present in UNIBS are probably also in UPC while new protocols, applications, or usages that appeared between 2009 and 2014 may be present in UPC but not in UNIBS. Although we do not actually try to detect these new applications, we hypothesize that their presence inside training data may be enough to

modify learned model.

The observed performance decrease is consistent with existing work targeting traffic classification [34]. This section shows that training data must be as diverse as possible to ensure a limited performance decrease when the trained model

11

Figure 12: 20 most important feature regarding Mean Decreased Impurity (MDI) using Random Forest with 100 trees for tunnel classification. Error bars represent a confidence interval with 99% confidence level.

is deployed in a new context.

4.5 Domain generalization and adversarial learning regarding MTU

The Maximum Transmission Unit (MTU) is the maximum amount of data that can be sent on a physical medium. Its value is 1500 for Ethernet and much of Internet [35]. This value is configurable by OSes on network interfaces. Reducing MTU decreases the size of sent and received packets. This in turns impacts both the number sent and received packets, and the total amount of data sent (because of the increasing overhead). We here address the impact of MTU modification on our three pipeline steps. We first consider common MTU used on Internet beyond the default value of 1500 which are, according to [35], 1472 and 1420. Then, we also consider an attacker that leverages a purposely lowered MTU to generate adversarial samples. We here use MTU values of 1400, 1300, and 1200. We thus obtain six MTU values (1500, 1472, 1420, 1400, 1300 and 1200) that we use to generate traffic (see Section 3.2) for the following experiments. We here do not use the SSH flows from UNIBS because we do not know which MTU was used.

Figure 14 presents our results for all three pipeline steps, packet size with direction, byte burst and Netflow v5 and v9 features when a 1500 bytes MTU is used for training. Results for MTU testing equal to 1500 are different than Section 4.2.2 because UNIBS SSH is not used. Similarly to Section 4.4, we include Netflow feature as classic flow-based features. Performances modifications regarding testing MTU change across pipeline steps and features: tunnel classification performances only decrease for Netflow v5, tunnel detection is impacted with both Netflow feature sets, while application classifica-

tion's F1 scores deteriorate for all features except packet sizes with direction. MTU lowering reduces the amount of data sent in a single packet. This, in turn, increases the packet number in a flow. Netflow v5 and v9 feature sets are impacted because they contains packet number-related features, either for all packet of flow or a subset (e.g. only from the source in Netflow v9). We observe that byte burst performances decrease as the testing MTU is smaller for application classification, while this does not happen with packet sizes with direction. We hypothesize that byte burst is actually able to look further in time into the flow, and is thus more sensitive to traffic changes caused by MTU reduction. Contrary to byte burst, we observe that performances for packet sizes with direction do not deteriorate as testing MTU decreases. Figure 15 displays boxplots and violinplots of packet sizes with direction which exhibit the ten biggest MDI feature importance when random forest is used. It shows that most feature values are far from the testing MTU. We thus hypothesize that packet size with direction performances do not decrease because feature values are not modified when testing MTU changes.

We do not display the results obtained with a training MTU of 1420 and 1470 due to the lack of space. The results are similar to what is obtained with a training MTU of 1500: best results are obtained when testing MTU is equal to the training one, and performances deteriorate otherwise.

5   Discussions and future work

We identify two main limitations to our work. The first limitation is related to tunneling protocols and application coverage. Although our tunneling protocol set is more complete than existing work, we did not address tools such as PPTP. We also did not deploy DNS-based tunelling which is commonly used for obfuscation. Regarding applications, we use web browsing and file transfer as applications inside tunnels. Common usage such as command line is not included in our work. We want to improve both deployed tunneling and application in future work. The other limitation is that we only use a single application inside each tunnel. Users however do not use tunelling this way, and often use several applications inside a tunnel. Our results regarding tunnel detection and classification using N first features such as packet size or byte burst are however not impacted by this limitation because they need a very small of features N to perform well (see Section 4.1).

Many recent on network traffic metadata analysis work use deep learning-related techniques. We however achieve very good performance with classical ML algorithms such as random forest or decision tree and thus decided to not use deep learning at all.

We also plan to extend this work along two other axes. First, we use each tunelling protocol with a single cipher suite. We want to extend this work by exploring other cipher suites (e.g. a different block size) which may alter network traffic metadata. Second, we plan to investigate prediction intervals
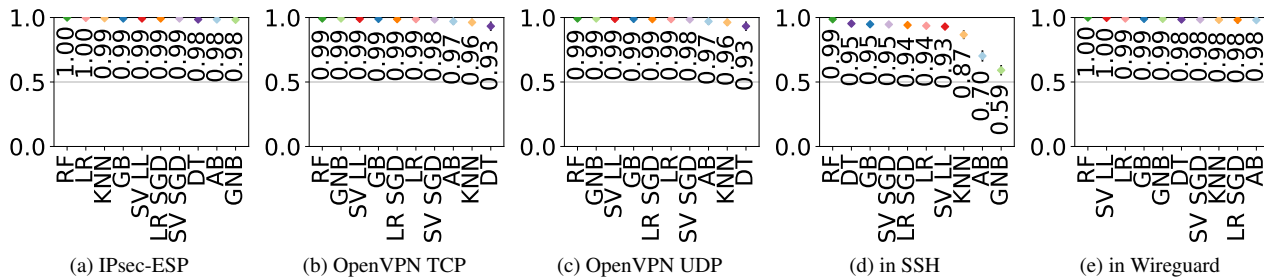
Figure 13: Application classification detection using selected features. Error bars represent a confidence interval with 99% confidence level.
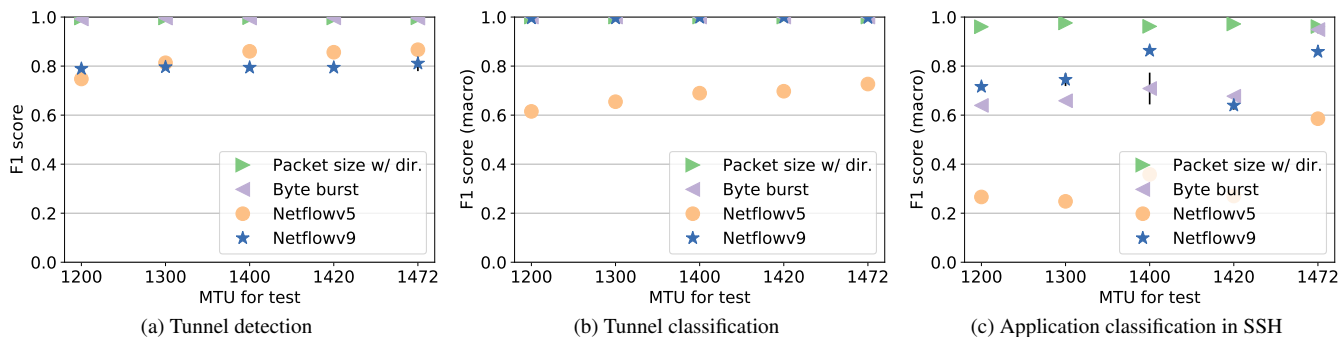

Figure 14: Random forest performance for training using MTU equals to 1500 and several MTU values for testing data. Error bars represent a confidence interval with 99% confidence level.
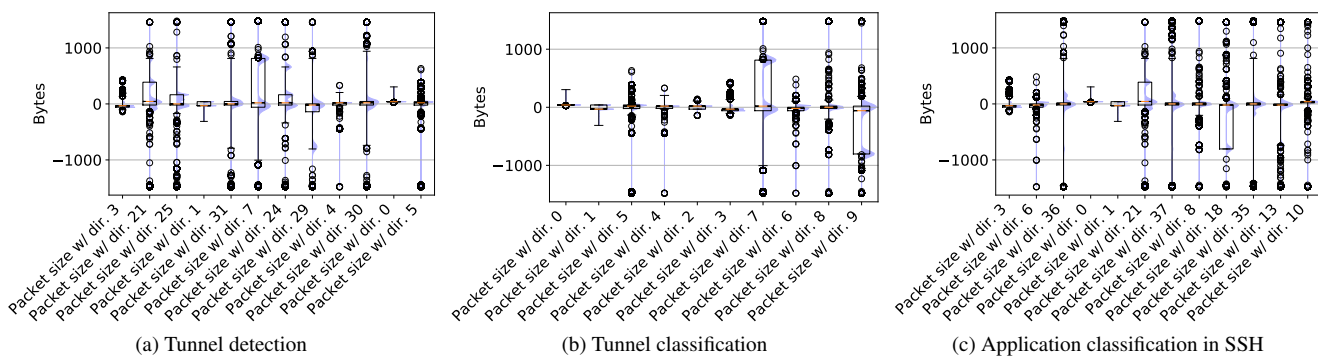

Figure 15: Packet sizes with direction in byte of the 10 most important feature regarding MDI using Random Forest with 100 trees for all pipeline steps. Boxplots encode median as the line inside the box, first and third quartile as upper and lower limit of the box, whiskers as 1st and 99th quantiles, and flier as points outside of these quantiles.

for both aleatoric and epistemic uncertainties using [36].

We think that our contribution on domain generalization and adversarial learning regarding MTU is very important for all use cases that leverage network traffic metadata. Regarding security and privacy, we hypothesize that use cases such as malware C2 detection, website fingerprinting, DoH fingerprinting, Tor detection are impacted by MTU manipulation.

6 Conclusion

In this work, we design a methodology to detect tunneling protocols, classify them once detected, and identify application used inside tunnels. We provide a thorough analysis of both features and machine learning algorithms. We also address domain generalization evaluation regarding the untunneled traffic for tunnel detection, and domain generalization and adversarial learning analysis regarding MTU for all three
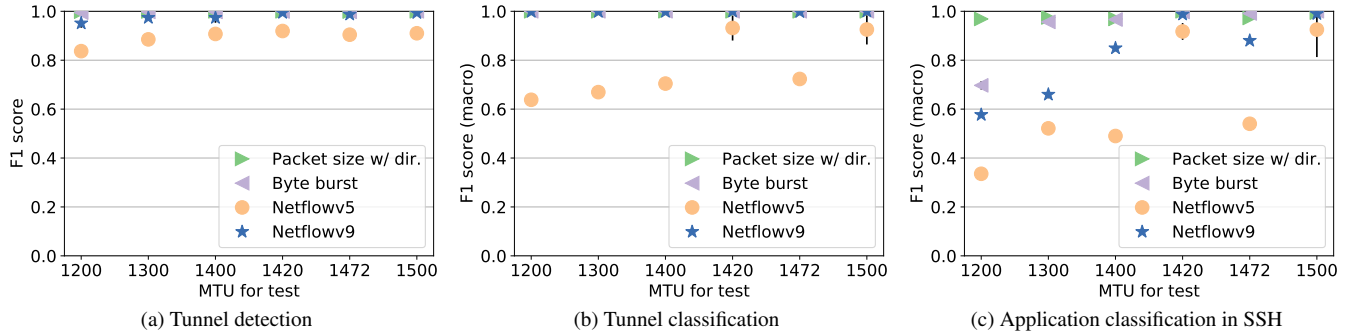
Figure 16: Random forest performance for training using MTU equals to 1500 and several MTU values for testing data. Error bars represent a confidence interval with 99% confidence level.

pipeline steps.

A Appendix

A.1 Learning curve

Figure 17 pictures the F1 scores obtained with random forest for several trainign dataset sizes. For all scenario, performance increase more important for smaller increase of training data size than This performance increase is negligible for the biggest increase in training data size (cf point to the right of each figure and log scale).

A.2 Netflow

A.3 Byte burst bigram and trigram

Figure 19 pictures the F1 scores obtained with ML algorithms for features from the packet burst family.

References

[1] L. Bernaille, R. Teixeira, I. Akodkenou, A. Soule, and K. Salamatian, "Traffic classification on the fly," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 2, pp. 23–26, 2006.

[2] L. Bernaille, R. Teixeira, and K. Salamatian, "Early application identification," in *Proceedings of the 2006 ACM CoNEXT conference*. ACM, 2006, p. 6.

[3] H. Kim, K. C. Claffy, M. Fomenkov, D. Barman, M. Faloutsos, and K. Lee, "Internet traffic classification demystified: myths, caveats, and the best practices," in *Proceedings of the 2008 ACM CoNEXT conference*. ACM, 2008, p. 11.

[4] M. Jaber, R. G. Cascella, and C. Barakat, "Can we trust the inter-packet time for traffic classification?" in *Communications (ICC), 2011 IEEE International Conference on*. IEEE, 2011, pp. 1–5.
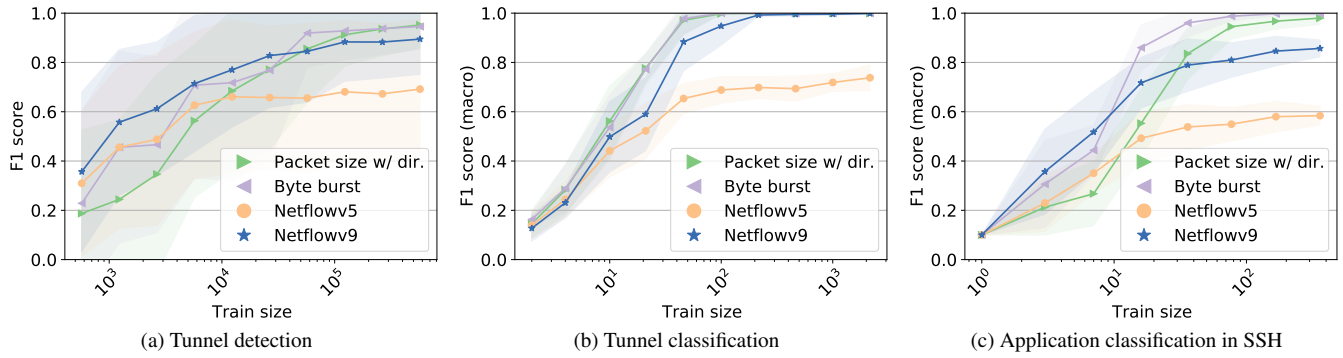
Figure 17: Random forest performance for several training dataset sizes. Error bars represent a confidence interval with 99% confidence level.
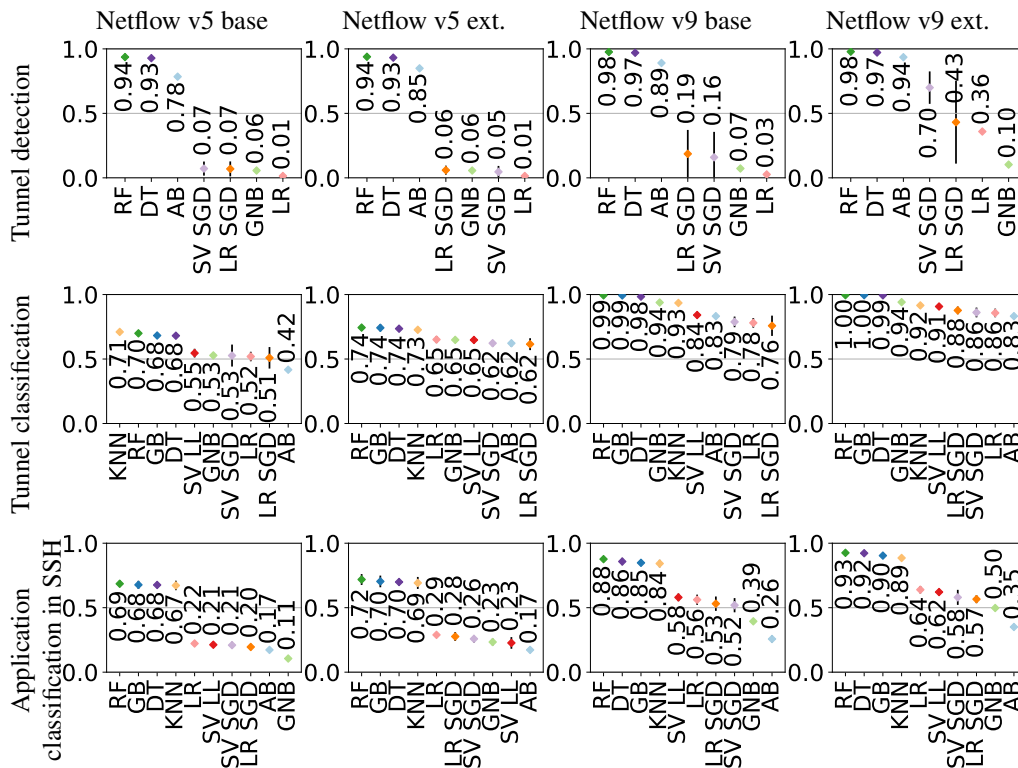

Figure 18: Comparison of packet size with direction, byte burst and IAT for all pipeline steps. Error bars represent a confidence interval with 99% confidence level.

[5] C. V. Wright, F. Monrose, and G. M. Masson, "On inferring application protocol behaviors in encrypted network traffic," *Journal of Machine Learning Research*, vol. 7, no. Dec, pp. 2745–2769, 2006.

[6] L. Bernaille and R. Teixeira, "Early recognition of encrypted applications," *Passive and Active Network Measurement*, pp. 165–175, 2007.

[7] G. Maiolini, A. Baiocchi, A. Iacovazzi, and A. Rizzi, "Real time identification of ssh encrypted application

flows by using cluster analysis techniques," *NETWORKING 2009*, pp. 182–194, 2009.

[8] G.-L. Sun, Y. Xue, Y. Dong, D. Wang, and C. Li, "An novel hybrid method for effectively classifying encrypted traffic," in *Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE*. IEEE, 2010, pp. 1–5.

[9] Y. Okada, S. Ata, N. Nakamura, Y. Nakahira, and I. Oka, "Application identification from encrypted traffic based
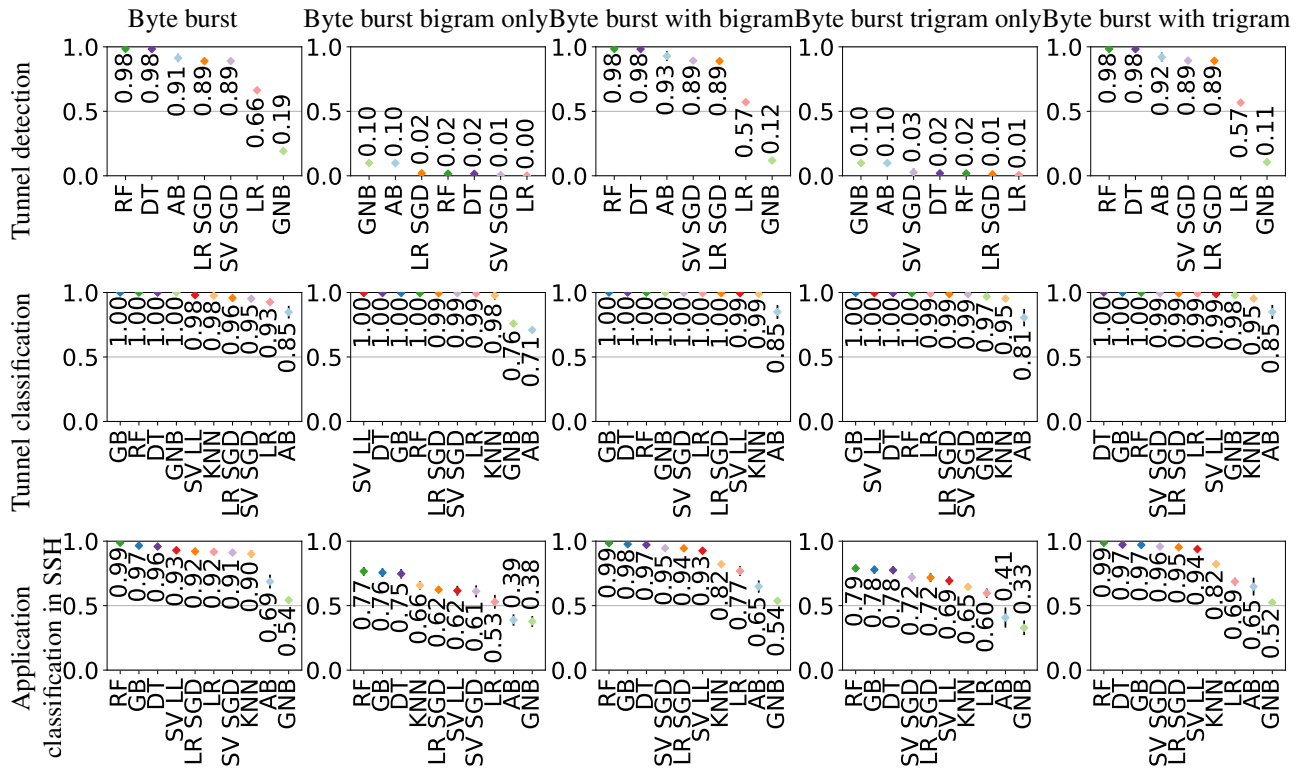
Figure 19: Comparison byte burst, byte burst bigram only, byte burst with bigram, byte burst trigram only, and byte burst with trigram for all pipeline steps. Error bars represent a confidence interval with 99% confidence level.

on characteristic changes by encryption," in *Communications Quality and Reliability (CQR), 2011 IEEE International Workshop Technical Committee on*. IEEE, 2011, pp. 1–6.

[10] ——, "Comparisons of machine learning algorithms for application identification of encrypted traffic," in *Machine Learning and Applications and Workshops (ICMLA), 2011 10th International Conference on*, vol. 2. IEEE, 2011, pp. 358–361.

[11] M. Korczyński and A. Duda, "Markov chain fingerprinting to classify encrypted traffic," in *Infocom, 2014 Proceedings IEEE*. IEEE, 2014, pp. 781–789.

[12] Y. Kumano, S. Ata, N. Nakamura, Y. Nakahira, and I. Oka, "Towards real-time processing for application identification of encrypted traffic," in *Computing, Networking and Communications (ICNC), 2014 International Conference on*. IEEE, 2014, pp. 136–140.

[13] "Cisco NetFlow v1, v5, v7 and v8," https://www.cisco.com/c/en/us/td/docs/net_mgmt/netflow_collection_engine/3-6/user/guide/format.html, accessed: 2017-01-31.

[14] "Cisco NetFlow v9," https://www.cisco.com/en/US/technologies/tk648/tk362/technologies_white_paper09186a00800a3db9.html, accessed: 2017-01-31.

[15] "Audit Record Generation and Utilization System," https://www.qosient.com/argus/index.shtml, accessed: 2017-01-31.

[16] "Zeek IDS," https://www.zeek.org/, accessed: 2017-01-31.

[17] "Verizon Injecting Perma-Cookies to Track Mobile Customers, Bypassing Privacy Controls," https://www.eff.org/deeplinks/2014/11/verizon-x-uidh, accessed: 2021-09-29.

[18] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[19] T. Wang, X. Cai, R. Nithyanand, R. Johnson, and I. Goldberg, "Effective attacks and provable defenses for website fingerprinting," in *23rd {USENIX} Security Symposium ({USENIX} Security 14)*, 2014, pp. 143–157.

[20] "UNIBS: Data sharing," http://netweb.ing.unibs.it/~ntw/tools/traces/, accessed: 2021-09-29.

[21] T. Bujlow, V. Carela-Español, and P. Barlet-Ros, "Independent comparison of popular dpi tools for traffic classification," *Computer Networks*, vol. 76, pp. 75–89, 2015.

[22] "UPC Traffic Classification ," https://cba.upc.edu/monitoring/traffic-classification, accessed: 2021-09-29.

[23] "Scikit-learn - AdaBoost classifier," https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.AdaBoostClassifier.html, accessed: 2021-09-29.

[24] "Scikit-learn - Gradient boosting classifier," https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.GradientBoostingClassifier.html, accessed: 2021-09-29.

[25] "Scikit-learn - Gausian naive Bayes," https://scikit-learn.org/stable/modules/generated/sklearn.naive_bayes.GaussianNB.html, accessed: 2021-09-29.

[26] "Scikit-learn - K nearest neighbors classifier," https://scikit-learn.org/stable/modules/generated/sklearn.neighbors.KNeighborsClassifier.html, accessed: 2021-09-29.

[27] "Scikit-learn - Linear SVC," https://scikit-learn.org/stable/modules/generated/sklearn.svm.LinearSVC.html, accessed: 2021-09-29.

[28] "Scikit-learn - Logistic regression classifier," https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.LogisticRegression.html, accessed: 2021-09-29.

[29] "Scikit-learn - SGD classifier," https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.SGDClassifier.html, accessed: 2021-09-29.

[30] "Scikit-learn - Random forest classifier," https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html, accessed: 2021-09-29.

[31] "Scikit-learn - Decision tree classifier," https://scikit-learn.org/stable/modules/generated/sklearn.tree.DecisionTreeClassifier.html, accessed: 2021-09-29.

[32] S. Raschka, "Model evaluation, model selection, and algorithm selection in machine learning," *arXiv preprint arXiv:1811.12808*, 2018.

[33] C. Novo and R. Morla, "Flow-based detection and proxy-based evasion of encrypted malware c2 traffic," in *Proceedings of the 13th ACM Workshop on Artificial Intelligence and Security*, 2020, pp. 83–91.

[34] M. Pietrzyk, J.-L. Costeux, G. Urvoy-Keller, and T. En-Najjary, "Challenging statistical classification for operational usage: the adsl case," in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*. ACM, 2009, pp. 122–135.

[35] A. Custura, G. Fairhurst, and I. Learmonth, "Exploring usable path mtu in the internet," in *Network Traffic Measurement and Analysis Conference*. IFIP Open Digital Library, 2018.

[36] "MAPIE - Model Agnostic Prediction Interval Estimator," https://mapie.readthedocs.io/en/latest/index.html, accessed: 2021-09-29.