

UC Irvine

UC Irvine Previously Published Works

Title

Efficacy of Full-Packet Encryption in Mitigating Protocol Detection for Evasive VPNs

Permalink

<https://escholarship.org/uc/item/2781r51d>

Journal

2025 IEEE Latin American Conference on Communications Latincom 2025, 00

Author

Parker, Amy I

Publication Date

2025-11-07

DOI

10.1109/latincom67778.2025.11345386

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Efficacy of Full-Packet Encryption in Mitigating Protocol Detection for Evasive VPNs

Amy I. Parker

Department of Computer Science

University of California, Irvine

Irvine, CA, USA

aiparker@uci.edu

Abstract—Full-packet encryption is a technique used by modern evasive Virtual Private Networks (VPNs) to avoid protocol-based flagging from censorship models by disguising their traffic as random noise on the network. Traditional methods for censoring full-packet-encryption based VPN protocols requires assuming a substantial amount of collateral damage, as other non-VPN network traffic that appears random will be blocked. We tested several machine learning-based classification models against the Aggressive Circumvention of Censorship (ACC) protocol, a fully-encrypted evasive VPN protocol which merges strategies from a wide variety of currently in-use evasive VPN protocols. Our testing found that while ACC was able to survive our models when compared to random noise, it was easily detectable with minimal collateral damage using several different machine learning models when within a stream of regular network traffic. While resistant to the current techniques deployed by nation-state censors, the ACC protocol and other evasive protocols are potentially subject to packet-based protocol identification utilizing similar classification models.

Index Terms—censorship circumvention, VPNs, protocol obfuscation, machine learning classification

I. INTRODUCTION

Internet censorship, particularly at the nation-state level, has become a pervasive means of restricting access to information and digital services. Virtual Private Networks (VPNs), which create encrypted tunnels between networks [1], have become a vital tool in evading these censorship regimes [2]. Users can send encrypted traffic — not demonstrably meeting topical censorship criteria — to a VPN provider outside the censorship regime, which will process the network requests and deliver encrypted responses back to the user.

Many censors attempt to combat the usage of VPNs in evading their censorship regimes by simply blocking all traffic utilizing VPN protocols. Whether through exposed header information or other details trivial to fingerprint, such as repeated values at a fixed or local offset, most VPN protocols are easy to detect; aside from typical header detection, Xue et al. studied methods of fingerprinting the OpenVPN protocol [3]. At the nation-state level, the Islamic Republic of Iran routinely deploys a protocol whitelister, prohibiting the usage on ports 53, 80, and 443 of any protocols other than DNS, HTTP, and HTTPS [4]. Other systems have employed active probing techniques to detect hosts running VPN servers and block their IP addresses. As it is generally uncommon to need to use a VPN across a censorship barrier for purposes not

related to evading censorship (such as accessing servers in a foreign office), and alternative methods exist for these use cases, the effects on non-evasive traffic are minimal.

VPN applications seeking to circumvent these protocol-level restrictions utilize a newer method of evasion: full-packet encryption (FPE). Instead of just encrypting the original datagrams (payloads), FPE-based protocols encrypt the entire packet, including protocol headers; through the properties of cryptographically secure encryption methods, the packets become indistinguishable from random noise [5]. Traditional protocol detection systems rely on identifiable headers in packets, but in FPE-based protocols, no such identifiable headers exist, making trivial protocol detection (also referred to as “deterministic” detection) impossible. Censors targeting FPE-based VPNs often instead use probabilistic, or “non-deterministic”, models which compute heuristics on packets and censor based on probabilities yielded by those heuristics; China’s modern approaches to censoring FPE utilize such nondeterministic methods [5].

In this paper, we present a detailed series of tests which utilize machine learning classification systems — a nondeterministic method — to attempt to discriminate between packets of a fully-encrypted VPN protocol and, depending on the test, either random noise or regular network traffic. The selected fully-encrypted VPN protocol is the Aggressive Circumvention of Censorship (ACC) protocol [6], which provides the `libacc` library [7] for encrypting and decrypting packets outside of a running ACC VPN tunnel. After training a given classification model on both types of packets, test packets are passed to the model for evaluation. We report back on the results of these tests, and identify future areas of research for FPE models.

The rest of the paper is organized as follows. §II reviews prior work in detecting and analyzing FPE models. §III describes our methodology, the machine learning classifiers, and the collection and encryption of the sample packets. §IV reviews the results of the tests and discusses these results. Finally, §V concludes.

II. PRIOR WORK

The most primitive form of censoring FPE-based protocols — protocol whitelisting — has been studied at the nation-state level primarily in relation to the Islamic Republic of Iran. The

first appearance occurred prior to Iran’s June 2013 elections, and was studied by Aryan et al. [8] as part of a broader review of the nation’s substantial Internet censorship regime. The censorship models would, upon detecting abnormal behavior of a whitelisted protocol or a non-whitelisted protocol, throttle the user’s connection, substantially limiting connectivity and causing larger traffic loads such as file downloads to fail.

Bock et al. [4] studied Iran’s newer protocol whitelist model in 2020. This model is inverted from typical protocol whitelists; rather than block all traffic on non-whitelisted ports and test for detectable protocols, Iran’s new regime allows all traffic not on ports 53, 80, and 443 (DNS, HTTP, and HTTPS, respectively), and only permits the three aforementioned protocols on those ports. This model blocks FPE models that run under another protocol’s port, a method which some FPE models use to evade traditional port-based protocol whitelist systems.

FPE-based protocols that have traditionally been able to circumvent these methods, such as Shadowsocks, VMess, Obfs4, and Outline, were temporarily blocked by the People’s Republic of China’s Great Firewall (GFW) in November 2021, as analyzed by Wu et al. [5]. In addition to checking for protocol fingerprints and sequences of readable characters, the 2021 GFW added a new criterion for blocking packets. If all other whitelisting tests failed, packets would be blocked if the average number of set bits per byte (popcount) was between 3.4 and 4.6. This method attempts to use the apparent randomness of full-packet encryption models against the models, as random packets are likely to have a popcount very close to 4.0.

As the 2013 and 2020 Iranian whitelist models can be overcome by adapting the design of a protocol, we do not include their models in tests. Likewise, strategies identified by Wu et al., as well as trivial strategies such as randomly varying nonces or pre-encryption padding to trigger a passing popcount, are able to overcome the 2021 GFW; we thus do not include these models in the tests.

Master’s dissertation [9] provides a key acknowledgement of the possibility of machine learning-based censorship methods in its analysis of nation-state-level deep packet inspection. At the time, products from Palo Alto and Fortinet — which conducted machine learning-supported analysis outside of the network path and then generated rules for blocking certain protocols, such as Tor, for in-path nodes [10] — were starting to see further development and deployment. These models primarily look for vulnerabilities in implementations of the protocol [11], but could be modified to create general detection rules. In-path machine learning models were still as computationally unrealistic then as now, but Master’s key insight — that “the potential for abuse of heuristic or machine learning-informed DPI of encrypted traffic remains” — is one of the few analyses that highlight the possibilities explored in this paper.

Master and Garuman also found that active probing remains a key threat to fully-encrypted VPN protocols [12]; the GFW’s early attacks on Tor relays relied on attempting to establish

connections with suspected relays, blocking any IPs it was able to successfully connect to [11]. Later protocols like Obfs3 were detected with similar active probing methods; while the initial detections became nondeterministic, the active probing results generally were deterministic, making it trivial to check if a server was running a fully-encrypted protocol [13]. These results make it paramount for fully-encrypted VPN protocols to resist active probing and avoid subtly giving out other information, such as detectable three-way handshakes [14] [15].

III. METHODOLOGY

A. Packet Collection

Packets were collected using `libpcap` from a Proxmox node operating a public-facing WordPress instance. The collection algorithm accepted all IPv4 and IPv6 packets except for ARP requests, IPv4 UDP packets (mostly used for node-to-node `corosync` within the Proxmox cluster and thus not representative of regular network traffic), and ICMPv6 Router Solicitations and Router Advertisements. Requests to local network resources, such as MariaDB instances and the instance’s NGINX reverse proxy, were included in the capture; while these are not traditionally accessed by end-users over a VPN, the protocols are not significantly different, and the packets overall make up a minimal percentage of captured traffic. Packets longer than the ACC VPN’s Maximum Transmission Unit (MTU) size of 1280 bytes were also dropped.

This collection returned approximately four million packets over three weeks. Due to the limits of available hardware research, this amount was too great to process, making data reduction necessary. 87.5 percent of the packets were randomly eliminated from the collection, reducing the dataset to 474,619 packets totaling approximately 69.6 MiB of data. As the deletion was random, this dataset reduction should not compromise the robustness or representativeness of the dataset and results.

B. ACC Encryption




Fig. 1. Structure of an ACC packet

Collected packets were then processed through a script utilizing `libacc` [7] to generate equivalent packets to those that would be sent over the network during an ACC connection. ACC uses a three-step encapsulation model to prepare packets for transmission over the network. The original IP packet is encapsulated into a Binary JavaScript Object Notation (BSON)

document alongside protocol header information. In the second phase, a variable amount of random padding is placed along with character delimiters (‘’ and ’’) on either side of the BSON document. Finally, the packet is encrypted with the GCM-SIV mode of AES, which increases resistance against nonce collisions [16]; the nonce is appended to the end of the packet. These packets are then written out to a separate directory from the original packets.

ACC was selected as the protocol for testing due to its theoretical equivalence in results with other protocols, the combination of many existing protocols’ strategies in the development of ACC, and the ease of integrating it into the testing pipeline. ACC is not dependent on pre-existing application states, other than the registration of an AES key with the target server, making encryption of individual packets easier. `libacc` also provides a much easier interface than other VPN systems (such as Shadowsocks and Obfs4) for encrypting the packets. Further, because ACC is indistinguishable from random, its results generalize to all other typical (instinguishable from random) FPE-based protocols.

C. Random Packet Generation

Random packets were generated for the tests comparing ACC packets to random data. To avoid packet-size-based influence on these tests, the distribution of packet size in the random packets was set to be the same as the post-processing ACC packets. A total of 474,619 random packets (107.7 MiB) were generated using Python’s `os.urandom()` interface, which pulls from the `/dev/urandom` character device on Linux systems.

D. Classification Models

All algorithms were implemented using either `scikit-learn` 1.5.2 or TensorFlow 2.7.10 on Python 3.11.2 as shipped with Debian GNU/Linux 12. For applicable models, all sixteen threads of the research server’s CPU were assigned a worker. Data was split 80%/20% between training and testing data. Eight machine learning classification algorithms were used:

- Quinlan’s **C4.5** [17] is a simple decision tree classification algorithm that runs in a single step with minimal parameters, implemented using `sklearn.tree.DecisionTreeClassifier` with no additional parameters.
- *k*-Nearest Neighbors (*k*-NN) classifies objects based on the classification of the *k* nearest neighbors to an object. `sklearn.neighbors.KNeighborsClassifier` was used with `n_neighbors` (*k*) set to $\sqrt{\text{len}(\text{data}_{\text{training}})}$.
- Multilayer Perceptron (**MLP**) is a standard feedforward neural network model. Two different MLP implementations were used:
 - ACC versus random tests used TensorFlow with a Keras model of form `[Dense(950), Dense(950; RELU), Dense(950; RELU), Dense(1; Sigmoid)]`.

- ACC versus unencrypted network traffic tests used `sklearn.neural_network.MLPClassifier` with hidden layer sizes of (950,950,950) and a tolerance of 10^{-9} .

- **Random Forest**, which was implemented through `sklearn.ensemble.RandomForestClassifier`, combines multiple decision trees for classification; the number of trees, `n_neighbors`, was set to 128.
- **Logistic Regression** attempts to map an input dataset to an output dataset by configuring parameters of a logistic function; while not as useful for these categories, it sets a minimum baseline for performance. Implementation was done via `sklearn.linear_model.LogisticRegression` with a maximum of 10,000 iterations and a tolerance of 10^{-6} .
- Convolutional Neural Networks (**CNNs**) transform a dataset with a series of convolutions which can be configured to develop a model before finally outputting a classification result. While typically used for images and multi-dimensional data, CNNs can be used for linear data, albeit uneffectively when such data is grouped into blocks such as with an AES encryption pass. A CNN model was implemented with Keras layers of `[Conv1D(1500, 32; RELU), MaxPooling1D((5,)), Conv1D(128, 32; RELU), MaxPooling1D((5,)), Conv1D(128, 32; RELU), MaxPooling1D((5,)), Flatten, Dense(128; RELU), Dense(1; Sigmoid)]`.
- Support Vector Machines (**SVMs**) use linear models along with tunable kernel functions to classify data. The SVM model was implemented with `sklearn.linear_model.SGDClassifier`, the hinge loss function (which implements a linear SVM), and a tolerance of 10^{-8} .
- Recurrent Neural Networks (**RNNs**) use recurrent units, which hold a hidden state during training and update based on both the previous hidden state and the current input. The RNN was implemented with a Keras model of form `[Embedding(1500, 1500), Bidirectional(LSTM(750)), Dense(1; Sigmoid)]`.

All eight algorithms were used for the ACC versus random data tests. Due to the incredibly significant results exhibited, only C4.5, *k*-NN, and MLP were used for testing ACC packets against regular network traffic; CNN, RNN, and SVM were in particular only included in the random tests for completeness and due to their presence in other literature, where they had been shown to generally be ineffective. Input data was provided to the models in the form of NumPy arrays of shape (1500,1,*dtype* = *uint8*). TensorFlow models used the adam optimizer and BinaryCrossentropy loss function with `from_logits=True`.

For the ACC versus real network data tests in particular, IPv4 and IPv6 headers, as well as TCP and UDP headers, were stripped from the packets when comparing them to the ACC packets. In transit, the ACC packets would have IP

and transport headers, but `libacc` does not yield these; as such, a model could easily detect the presence of network and transport layer headers, ruining the tests. An alternative and nearly equivalent method would have been to add IP and TCP/UDP headers to the ACC packets; eliminating them reduces testing variability.

E. Metrics

Both `scikit-learn` and TensorFlow models output five primary metrics used for evaluating and comparing the efficacy of models:

- **Accuracy** measures the proportion of test results which are correct, being defined as

$$A = \frac{TP + TN}{TP + FP + TN + FN}$$

- **Precision** measures the proportion of true positives among all positives, being the true positive rate; it is defined as

$$P = \frac{TP}{TP + FP}$$

- **Recall** measures the ability to detect positives in the input by measuring the proportion of ACC packets that were detected, defined as

$$R = \frac{TP}{TP + FN}$$

- **Collateral Damage** measures the percentage of packets which were unintentionally censored. For censors, keeping collateral damage low is important to ensure that networks are still reasonably accessible by compliant users [5]. Unlike all other metrics measured here, a *lower* collateral damage score is better. It is defined as

$$C = \frac{FP}{FP + TN}$$

- **F_1 -score** provides a generalized metric for comparing classifiers. It is the harmonic mean of precision and recall, defined as:

$$F_1 = \frac{2}{R^{-1} + P^{-1}} = \frac{2(TP)}{2(TP) + FP + FN}$$

Note that TP refers to the count of true positives, TN true negatives, FP false positives, and FN false negatives. Metrics were manually calculated from the raw output counts for `scikit-learn`; for TensorFlow, F_1 -score and collateral damage were calculated using the precision and recall results, as Keras' F_1 metrics were inaccurate and there is no native support for collateral damage.

A censorship model is deemed *effective* at blocking ACC traffic in a test if the following conditions are true:

- $F_1 > 0.95$; this is a common criterion for any classification algorithm, as it accounts for the effects on both ACC and non-ACC traffic better than accuracy and other "complete" measures.
- $C < 0.01$; if more than 1% of data is falsely censored, the impact on the network will be significant, and the model




Fig. 2. Metrics from ACC versus Random testing

is generally infeasible to deploy. China's temporary block on FPE-based models had a collateral damage rate of 0.6%, which was only feasible for a short amount of time and later discontinued [5].

IV. RESULTS

A. ACC versus Random Packets

The results of these tests are displayed in Fig. 2. All models performed equivalently in the accuracy and precision benchmarks. The CNN-based model overperformed in the F_1 -score and recall metrics, suffering later in collateral damage. The SVM-based and RNN-based models did the reverse, having low recall and F_1 -scores but minimal collateral damage; this appears to be a result of the CNN model being too eager (more positive guesses than negative guesses) while the SVM and RNN models were too hesitant (more negative guesses than positive guesses). As expected, the CNN, RNN, and SVM models demonstrated their inappropriateness for VPN detection tasks. All other models performed nearly equally ($\sim 50\%$) on all other benchmarks.

None of the models were effective in classifying packets as ACC versus non-ACC. Taking into account the eagerness of the CNN-based model and the hesitant nature of the SVM-based and RNN-based models, the classifications returned by each model were essentially random.

B. ACC versus Network Packets

The results of these tests are displayed in Fig. 3; note that C has been scaled by a factor of 10 for ease of visibility.

The C4.5 model performed almost perfectly. The iteration showed in the graph had only *one* false negative, with every other prediction being correct; the testing set size was 78,925.




Fig. 3. Metrics from ACC vs Network testing

Re-running this test on the side several times produced similar results, with no false positives and the count of false negatives being within $[0, 7]$. While this model is still nondeterministic, as C4.5 is a probabilistic classification algorithm, this behavior makes it nearly equivalent to a deterministic algorithm. Given that it meets the effectiveness criteria ($F_1 \approx 0.9999873 > 0.95$, $C = 0.0 < 0.01$), C4.5 can be clearly classified as effective in identifying ACC packets within a network stream. Being a decision-tree-based algorithm, C4.5 is more naturally suited to these types of binary detection tests as it creates a direct set of tests on the data; it is thus more similar to deterministic classifiers in approach than the other models tested here.

k -NN did not perform as well, but still exceedingly better than in the ACC-versus-Random tests. It turned out to be very hesitant, having 200 times the number of false negatives as false positives; this was heavily reflected in its F_1 -score of approximately 0.91239, making it not effective as a censorship model. However, this hesitancy was able to keep its collateral damage score low at approximately 0.079%. k -NN is very sensitive to training data as it is a region-based classification model and is thus susceptible to regions which lacked entries in the dataset, creating a blurrier and poorly-defined boundary. While other data sets may be able to generate a better model, the effectiveness of C4.5 and its lower overhead relative to k -NN makes this moot.

MLP performed very well, reaching a satisfactory F_1 -score of approximately 0.9863. While not nearly as effective as C4.5, it achieved impressive results, even with a relatively small model. The particular model trained is not quite effective, as the collateral damage of 1.13% is above the efficacy threshold of $C < 0.01$; however, alterations to the model or an increase in training data quantities could yield better results.

C. Applicability to Censorship Models

Current individual-packet censorship models utilize deterministic or lightweight nondeterministic tests due to the computational overhead and latency added by censorship models to packet transmission. While checking header information or calculating a packet's popcount requires a minimal amount of additional computation relative to route calculation, deploying a machine learning model, even if pre-trained, adds a significant amount of overhead which current censorship and routing equipment may not be prepared for, requiring expensive investment in network infrastructure that may not be feasible at the nation-state level.

Despite this, the high performance of C4.5 indicates that these models may be deployable now or in the near future. Of the eight algorithms tested, C4.5 is the least computationally intensive; a single thread on a Ryzen 7 3700X (an AMD Zen 2 CPU) was able to achieve 548044.96 packets classified per second using C4.5 in preliminary testing, which would allow servicing of roughly 8.8 million packets/second across all threads; specialized hardware would likely be able to achieve much higher performance with at reasonably low electrical expenditure and initial cost. Generating new C4.5 models is itself incredibly efficient; C4.5 trains in $O(mn \log_2 n)$ time [18], where m is the size of the input packet and n is the number of samples.

The impact of a deployable model based on this framework would be substantial. As mentioned previously, machine learning models have been generally absent from real-time packet-level censorship evasion efforts, and especially efforts to detect fully-encrypted VPNs, due to the extremely high overhead posed [9] and concerns regarding collateral damage [5]. Deployment of a system efficient enough to run on network hardware in the near future with low collateral damage gives censors a significant advantage in blocking VPN traffic and poses a critical threat to censorship evasion.

V. CONCLUSIONS AND FUTURE WORK

In this work, we developed methods for testing the efficacy of the ACC protocol, which uses the full-packet encryption model, against machine learning classification-based censorship models. While ACC was able to withstand tests against random data, ACC packets were easily identified when in a stream of regular network packets. C4.5 was identified as the most effective model for detecting ACC experimentally, and the ability for C4.5 models to be integrated into censorship middleware was explored.

These tests were only run on the ACC protocol; while all FPE-based protocols should have similar results to ACC, future research should verify that these results apply to other protocols such as Shadowsocks and Obfs4. Additional research can be conducted on identifying which FPE-based protocol is in use, which may identify other detection strategies for these protocols.

Lastly, the efficacy of masquerade modes (such as sending VPN traffic over HTTPS) should be further investigated, as

such methods have previously been the most effective in evading protocol-targeting censorship. TLS-based masquerading remains the most promising; for lower-throughput connections on networks which do not enforce single-DNS servers or block ICMP, tunneling over DNS or ICMP data may introduce new effective strategies.

ACKNOWLEDGMENTS

I thank my research advisor, Marc Velasco, for the incredible guidance and advice he has provided in bringing this research to conclusion. I also thank the members of the ACC development team [6] for developing the FPE-based VPN protocol used in this research.

REFERENCES

- [1] T. Berger. “Analysis of current VPN technologies”. In: *First International Conference on Availability, Reliability and Security (ARES’06)*. 2006, 8 pp.–115. DOI: 10.1109/ARES.2006.30.
- [2] Yuzhi Wang et al. “GoHop: Personal VPN to defend from censorship”. In: *16th International Conference on Advanced Communication Technology*. 2014, pp. 27–33. DOI: 10.1109/ICACT.2014.6778916.
- [3] Diwen Xue et al. “OpenVPN is Open to VPN Fingerprinting”. In: *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, 2022, pp. 483–500. ISBN: 978-1-939133-31-1. URL: <https://www.usenix.org/conference/usenixsecurity22/presentation/xue-diwen>.
- [4] Kevin Bock et al. “Detecting and Evading Censorship-in-Depth: A Case Study of Iran’s Protocol Whitelister”. In: *10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20)*. Online: USENIX Association, 2020. URL: <https://www.usenix.org/conference/foci20/presentation/bock>.
- [5] Mingshi Wu et al. “How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic”. In: *32nd USENIX Security Symposium (USENIX Security 23)*. Anaheim, CA: USENIX Association, 2023, pp. 2653–2670. ISBN: 978-1-939133-37-3. URL: <https://www.usenix.org/conference/usenixsecurity23/presentation/wu-mingshi>.
- [6] Amy Iris Parker et al. *ACC: An Evasive Protocol for Secure Communications and Virtual Private Networks*. Dec. 3, 2023. URL: https://www.academia.edu/114397270/ACC_An_Evasive_Protocol_for_Secure_Communications_and_Virtual_Private_Networks.
- [7] Kevin Liu, Owen de Vita, and Amy Iris Parker. *libacc: library for the ACC protocol*. Feb. 21, 2024. URL: <https://github.com/amyipdev/libacc>.
- [8] Simurgh Aryan, Homa Aryan, and J. Alex Halderman. “Internet Censorship in Iran: A First Look”. In: *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*. Washington, D.C.: USENIX Association, 2013. URL: <https://www.usenix.org/conference/foci13/workshop-program/presentation/aryan>.
- [9] Alexander Master. “Modeling and Characterization of Internet Censorship Technologies”. PhD dissertation. West Lafayette, IN: Purdue University, 2023. URL: <https://www.proquest.com/docview/2890696337>.
- [10] Palo Alto Networks. *How to Block Tor (The Onion Router)*. 2018. URL: <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIRtCA> (visited on 12/25/2023).
- [11] Philipp Winter and Stefan Lindskog. “How the Great Firewall of China is Blocking Tor”. In: *2nd USENIX Workshop on Free and Open Communications on the Internet (FOCI ’12)*. Bellevue, WA: USENIX Association, 2012. URL: <https://www.usenix.org/system/files/conference/foci12/foci12-final2.pdf>.
- [12] Alexander Master and Christina Garuman. “A Worldwide View of Nation-state Internet Censorship”. In: *Free and Open Communications on the Internet 2023*. Lausanne, Switzerland: Censored Planet, 2023. URL: <https://petsymposium.org/foci/2023/foci-2023-0008.pdf>.
- [13] Roya Ensafi et al. “Examining How the Great Firewall Discovers Hidden Circumvention Servers”. In: *Proceedings of the 2015 Internet Measurement Conference*. Tokyo, Japan: Association for Computing Machinery, 2015, pp. 445–458. ISBN: 978-1-450-33848-6. DOI: 10.1145/2815675.2815690.
- [14] Michelina Hanlon et al. “Detecting VPN Traffic through Encapsulated TCP Behavior”. In: *Free and Open Communications on the Internet 2024*. Bristol, United Kingdom: Censored Planet, 2024. URL: <https://petsymposium.org/foci/2024/foci-2024-0016.pdf>.
- [15] Christian P. Wessel. “Detecting TCP Flows Inside Encrypted VPN Tunnels”. Master’s thesis. Karlstad, Sweden: Karlstad University, 2019.
- [16] Priyanka Bose, Viet Tung Hoang, and Stefano Tessaro. “Revisiting AES-GCM-SIV: Multi-user Security, Faster Key Derivation, and Better Bounds”. In: *Advances in Cryptology – EUROCRYPT 2018*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Cham: Springer International Publishing, 2018, pp. 468–499. ISBN: 978-3-319-78381-9.
- [17] J.R. Quinlan. *C4.5: Programs for Machine Learning*. Ebrary online. San Mateo, CA: Morgan Kaufmann, 2014. ISBN: 9780080500584. URL: <https://books.google.com/books?id=b3ujBQAAQBAJ>.
- [18] Habiba Muhammad Sani, Ci Lei, and Daniel Neagu. “Computational Complexity Analysis of Decision Tree Algorithms”. In: *Artificial Intelligence XXXV (2018)*. Ed. by Max Bramer and Miltos Petridis, pp. 191–197.