Fingerprinting VPN Traffic: An Evaluation of Website Fingerprinting Attacks on Modern Virtual Private Network Applications

Author: Oufan Hai

Advisors: Prof. Adam Purtee, Prof. John Criswell

University of Rochester

Background

VPNs are increasingly promoted as a privacy-enhancing technology and a solution to protecting users' privacy from surveillance and cyber attacks. While most VPN protocols encrypt users' browsing traffic, the research community has repeatedly demonstrated that encryption algorithms need not be broken for malicious agents with knowledge of users' encrypted traffic to fingerprint the websites they visit. We investigate whether newer VPN technologies (protocols and privacy-enhancing features) make VPN traffic harder to fingerprint.

Our Contribution

- New datasets containing **643 GB** of VPN traffic.
- We show that when performing website fingerprinting on VPN traffic, simple but effective machine learning classifiers achieved performance ($F1 \approx 0.95$) comparable to classifiers using **deep neural networks**.
- We show that VPN traffic is often as easy to fingerprint as **bare HTTPS traffic**.
- Website Fingerprinting classifiers trained on one dataset **remain effective** when evaluated on different datasets.

Website Fingerprinting Threat Model

We assume the attacker has direct access to users' Internet traffic. Some entities with such capabilities are the state (especially authoritarian ones), Internet Service Providers (ISPs), and compromised machines on the same network.

We assume the attacker is able to purchase all the commercial VPNs available on the market. The attacker collects their own datasets to train custom machine learning models. Similar to Sirinam et al.'s threat model, we further assume that the attacker has a list of "monitored" websites and the goal of website fingerprinting attacks is to find out if the user visits a website in the monitored set [2].

Dataset Creation

1. **Packet Capturing**: Capturing '.pcap' files using Selenium and tcpdump.
2. **Feature Extraction**: We extract three types of features [1, 3] from the network traffic captured.
   - Length + Direction (e.g. [4, -54, 60, -60, 54])
   - Direction (e.g. [1, -1, 1, -1, 1])
   - Unique Packet Sizes (e.g. [4, 54, 60])

Experiments

Performance of Models

Deep CNN-based website fingerprinting models (e.g., DF with 8 Convolutional Layers [2]) do not perform significantly better than simple models (e.g., single-layered CNN) when the number of monitored website is less than 100.

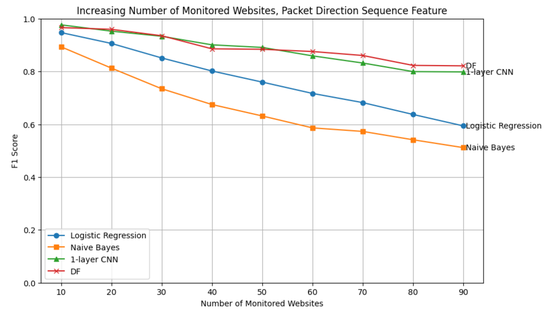

Performance of Features

	Length + Direction	Direction	Unique Sizes
HTTPS	**0.8995**	0.7211	**0.7470**
OpenVPN	0.8854	0.9156	0.7439
WireGuard	0.8954	**0.9500**	0.7437
Average	0.8935	0.8622	0.7448

VPN traffic is at least as easy to fingerprint as bare HTTPS traffic when the right feature is used. 1-dimensional metadata (packet direction sequence) makes VPN traffic more fingerprintable than HTTPS.

Cross-VPN Fingerprinting

We found that classifiers trained on one dataset remain effective when evaluated on different datasets. For example, machine learning models trained on ExpressVPN traffic accurately classify the NordVPN traffic.

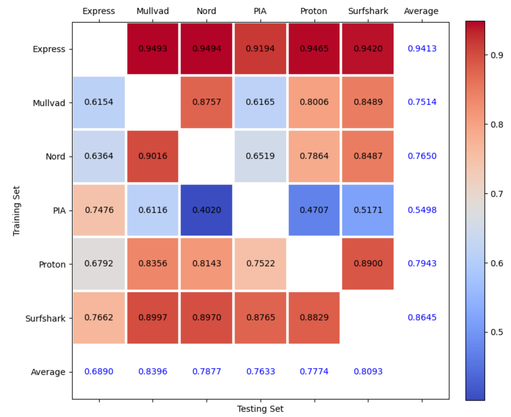

Conclusion

Our results challenge commercial VPN providers' claim that their services could serve as an effective deterrence against surveillance and cyber attacks. With limited data (our largest training set contains 20,000 traces), we trained classifiers that are highly effective at fingerprinting websites. We conclude that website fingerprinting continues to be a challenge to privacy.

References

[1] Dominik Herrmann, Rolf Wendolsky, and Hannes Federrath. Website fingerprinting: Attacking popular privacy enhancing technologies with the multinomial naïve-bayes classifier. In *Proceedings of the 2009 ACM Workshop on Cloud Computing Security*, CCSW '09, New York, NY, USA, 2009. Association for Computing Machinery.

[2] Payap Sirinam, Mohsen Imani, Marc Juarez, and Matthew Wright. Deep fingerprinting: Undermining website fingerprinting defenses with deep learning. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018.

[3] Junhua Yan and Jasleen Kaur. Feature selection for website fingerprinting. *Proceedings on Privacy Enhancing Technologies*, 2018.