DOI:10.1145/3556640

To view the accompanying paper, visit doi.acm.org/10.1145/3559439 rh

Technical Perspective
Traffic Classification in the Era of Deep Learning

By Athina Markopoulou

NETWORK TRAFFIC CLASSIFICATION is a fundamental problem in networking. Given observations of network traffic, the goal is to infer properties of interest, such as what application generated the traffic. This enables network operators to monitor and optimize performance, detect anomalies or malware, block unwanted traffic, inform capacity planning, and so on.

The problem has been extensively studied for more than 20 years, using a combination of heuristics, based on domain expertise, and automated methodologies. Some techniques rely on hard-coded rules, such as the use of well-known ports or servers. For example, a DNS request, the HTTP Host field, or the SNI field in TLS, may all reveal the name of the server contacted (for example, server.netflix.com), which may in turn be indicative of the service itself. Other techniques rely on behavioral characteristics, such as flow statistics, communication patterns, or traffic volume time series. For example, voice-over-IP applications generate small, evenly spaced packets; Web applications produce bursty traffic; and smart home devices exchange occasional status updates and commands with the cloud.

Two trends have heavily influenced modern network traffic classification, including this paper. First, network traffic is increasingly encrypted, which impedes techniques that rely on payload inspection. This leaves only with the option of behavioral characteristics, which can still be observed on encrypted traffic and are also difficult to obfuscate. For example, a voice-over-IP application maintains its characteristic traffic shape, independently of server or encryption. Second, recent advances in data mining and machine learning provide a powerful toolbox to "throw" at the traffic classification problem. Network traffic measurement typically collects large datasets with complex patterns that lend themselves naturally to learning techniques.

The combination of these trends has led to a plethora of papers over the past decade that apply machine learning to encrypted traffic. The following paper does a great job in reviewing related work in this space. Key questions and design choices include the following:

- *Classification task:* Is the goal to classify packets, flows, or groups of flows? Infer the application or service class? Detect malware or anomalies?
- *Feature engineering:* What features should be used as input to the machine learning model?
- *Model architecture:* What is the right model to capture traffic characteristics relevant to the classification task? Is the model treated as a black box, or is it interpretable and reflects domain-expertise?
- *Generalization:* Does the model overfit to the training data or does it generalize? How robust is it to variations in the protocols and server names involved? How easy is it to evade?

The authors seek to classify encrypted traffic flows according to their service class, and *not* to the particular application that generated them. For example, Netflix and Hulu traffic are classified as "video streaming," while Facebook and Twitter traffic are classified as "social media." Somewhat surprisingly, the paper argues that service classification is more challenging than application classification, because it requires identifying common traffic characteristics across similar applications, without relying on simple rules or training on all applications. A key observation is that when deep learning models are used as black boxes on the full raw traffic trace, they typically overfit and end up learning simple, deterministic rules, such as looking up the server name. As a result, these models are unnecessarily complex and expensive, yet they do not learn interesting, generalizable behaviors.

This paper leverages domain knowledge in network measurement to select network traffic features and a neural network architecture tailored specifically to them. The goal is to learn traffic patterns that are inherent in the service class and apply across protocols, server names, and applications in that class. More specifically, the paper extracts three types of features from a traffic flow: raw bytes from packet headers in the TLS handshake, flow statistics (pertaining to packet lengths, inter-arrival-times, number of bytes, among others), and flow time series (of packet size, timing, and direction). It then uses a convolutional neural network (CNN), a stacked long short-term memory (LSTM) and additional dense layers to learn spatial, temporal, and other types of traffic characteristics, respectively. It eventually combines the three parts to predict the service class of a traffic flow. Evaluation on real-world traffic datasets performs well on HTTPS, and across a range of Web and transport protocols and applications.

Overall, this paper provides great insights into protocol-agnostic classification of encrypted traffic. It also proposes a principled and intuitive solution using a set of features and a neural network architecture that capture domain expertise. Directions for future work include: refining the service class definition (for example, TikTok traffic may be considered as both "social network" and "video"); consider changing network conditions that may affect traffic characteristics; and a better understanding of the minimum set of features and the smallest model architecture necessary for a particular traffic classification task. C

Athina Markopoulou (athina@uci.edu) is a professor in the Department of Electrical Engineering and Computer Science at the University of California, Irvine, USA.