Operating system fingerprinting with Artificial Intelligence

Rubén Pérez-Jove, Cristian R. Munteanu, José M. Vázquez-Naya

In the field of computer security, the possibility of knowing which specific version of an operating system is running behind a machine can be useful in order to assist in a penetration test or to monitor the devices connected to a specific network. This task can be done interacting directly with the systems you want to fingerprint, performing an active scan, or just sniffing the normal traffic produced by the targets and analyze it, following a passive approach. One of the most widespread tools that better provides the first functionality is Nmap, which follows a rule-based approach for its active scan.

In this context, applying Machine Learning techniques seems to be a good option for representing the knowledge kept in the Nmap database and extract the main features which identify an operating system family. The first research line on this topic would be the exploration of the strengths of different Machine Learning algorithms to perform operating system fingerprinting. Furthermore, some clustering analysis could be interesting in order to understand which of the tests performed by Nmap are more decisive when carrying on this task.