Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection

Researchers

Yisroel Mirsky
yisroel@post.bgu.ac.il

Tomer Doitshman
tomerdoi@post.bgu.ac.il

Prof. Yuval Elovici
elovici@inter.net.il

Dr. Asaf Shabtai
shabtaia@bgu.ac.il

Publications

Y. Mirsky, T. Doitshman, Y. Elovici, and A. Shabtai, "Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection," Network and Distributed System Security Symposium (NDSS'18), 2018.

Goals

Neural networks have become an increasingly popular solution for network intrusion detection systems (NIDS). Their capability to learn complex patterns and behaviors make them a suitable solution for differentiating between normal traffic and network attacks. However, a major drawback of neural networks is the amount of resources needed to train them. Many network gateways and routers devices, which could potentially host an NIDS, simply do not have the memory or processing power to train and sometimes even execute such models. More importantly, the existing neural network solutions are trained in a supervised manner, meaning that an expert must label the network traffic and update the model manually. In response to this problem, the researchers have developed Kitsune: a plug and play NIDS which can learn to detect attacks on the local network, without supervision, and in an efficient online manner.

Description

Kitsune's core algorithm (KitNET) uses an ensemble of neural networks called autoencoders to collectively differentiate between normal and abnormal traffic patterns. KitNET is supported by a feature extraction framework which efficiently tracks the patterns of every network channel. The evaluations show that Kitsune can detect various attacks with a performance comparable to offline anomaly detectors, even on a Raspberry PI. This demonstrates that Kitsune can be a practical and economic NIDS.