Unsupervised intrusion detection through skip-gram models of network behavior

Rafael San Miguel Carrasco , Miguel-Angel Sicilia

PII: S0167-4048(18)30270-0
DOI: 10.1016/j.cose.2018.07.003
Reference: COSE 1362

To appear in: *Computers & Security*

Received date: 25 March 2018
Revised date: 7 June 2018
Accepted date: 2 July 2018

Please cite this article as: Rafael San Miguel Carrasco , Miguel-Angel Sicilia , Unsupervised intrusion detection through skip-gram models of network behavior, *Computers & Security* (2018), doi: 10.1016/j.cose.2018.07.003

Unsupervised intrusion detection through skip-gram models of network behavior

Rafael San Miguel Carrasco, Miguel-Angel Sicilia

University of Alcalá, Madrid, Spain

Keywords

Skip-gram modeling, Neural networks, Anomaly detection, Intrusion detection, Unsupervised learning, Word2vec, Unknown attacks

Abstract

Detecting intrusions is one of the main objectives of computer security. Attacks have become overly sophisticated over the years in order to remain effective and stealthy. Major breaches are typically perpetrated using techniques that are polymorphic, multi-vector, multi-stage and targeted, that is, adopting forms that were never seen before. Anomaly detection, which doesn't make any assumption about the shape of a potential attack but instead on legitimate behavior, seems to be a suitable approach in order to defeat sophisticated intrusions. Skip-gram modeling, a word2vec algorithm variant, was leveraged to model systems' legitimate network behavior. The resulting model was then used to spot intrusions in a test dataset. The optimal configuration led to 99.20% precision, 82.07% recall, and 91.02% accuracy, with a false positive rate of 0.61%, which is significantly lower than most state-of-the-art methods. These metrics were achieved under a fully unsupervised setting, that is, without any prior knowledge of what constitutes an attack. Furthermore, the approach provides benefits in terms of interpretability and log storage requirements, as it requires a small amount of input features. It also produces information about systems behavior and their relationships, that can be reused by other analysis techniques to obtain further insights.

1. Introduction

Computer and network intrusion detection is a popular research field. One of their objectives is to design and implement intrusion detection techniques with high detection rates and low false positive rates [1]. As cyber attacks become more sophisticated, the main challenge becomes deploying techniques that can recognize unknown attacks [2], that is, those for which an statistical or signature-based pattern doesn't exist. Another related issue is the need to handle huge amounts of data that are collected [3] from multiple data sources, and the fact that not all data might be relevant and relevant data might not be present. Multiple general-purpose algorithms have been re-engineered for intrusion detection. They fall into one of these categories: misuse detection and anomaly detection.

Misuse detection detects attacks by comparing current activity with the expected actions of an attacker [4]. Anomaly detection builds a normal activity profile for a system [5] instead. Misuse detection requires training observations to be labeled as normal or malicious activity. In anomaly detection, only legitimate behavior is considered during training. Behavior that significantly deviates from modeled legitimate behavior in the test data is then flagged as an anomaly. Labeled data for misuse detection might not be easy to obtain for the following reasons:

- A record might or might not constitute an attack depending on the context. Likewise, an attack might relate to multiple records. These relationships cannot be captured by record-level labels.
- Related records required to recognize an attack might have been triggered by different data sources.
- No datasets containing real-life traffic are generally available. Moreover, cyberattacks rates and impact metrics remain unavailable to researchers [6].