

Metrics for community dynamics applied to unsupervised attacks detection

J. Michel^{1,2}, P. Parrend^{1,2}

¹ Laboratoire des sciences de l'ingénieur, de l'informatique et de l'imagerie, Icube

² Laboratoire de Recherche de L'EPITA , LRE

May 4th 2023

Abstract

Attack detection in big networks has become a necessity. Yet, with the ever changing threat landscape and massive amount of data to handle, network intrusion detection systems (NIDS) end up being obsolete. Different machine-learning-based solutions have been developed to answer the detection problem for data with evolving statistical distributions. However, no approach has proved to be both scalable and robust to passing time. In this paper, we propose a scalable and unsupervised approach to detect behavioral patterns without prior knowledge on the nature of attacks. For this purpose, we define novel metrics for graph community dynamics and use them as feature with unsupervised detection algorithm on the UGR'16 dataset. The proposed approach improves existing detection algorithms by 285.56% in precision and 222.82% in recall when compared to usual feature extraction (FE) using isolation forest.

Keywords

Features Engineering, Graph community metrics, Scalability, Graph representation, Unsupervised detection approach, Dynamic graphs, Attacks detection

1 Introduction

Attack detection in big networks requires processing an increasing amount of data. Furthermore, the behaviour of the data changes over time in ways that cannot be predicted. This phenomenon called concept drift renders prior existing models invalid. Thus, it has come to light that there is a need for scalable and adaptable solutions. Due to the evolving nature of attacks to detect at any given time, methods which do not use knowledge of specific attacks have gained the attention of the community [1, 5]. To solve this kind of problems, unsupervised approaches for anomaly detection have been studied [8]. But unsupervised approaches, especially outlier detection algorithms have shortcomings in their inability to choose the right criteria for anomalies if the configuration scope is unconstrained. This configuration scope is controlled through hyperparameter tuning, but especially through feature engineering [3]. Only features relevant for detecting attacks of interest should be fed to the anomaly detection algorithm.

Our contributions In this paper we present graph community metrics used as features for anomaly detection and specifically applied to the detection of behaviour patterns for attack detection. *Density* and *Externality* show remarkable results, but they do not take into account the evolution of data over time. We therefore additionally define *Local* and *Global Stability* values as metrics for graph community dynamics. They are included in the set of candidate features for anomaly detection, and fed to anomaly detection models such as Isolation Forest. The evaluation of these metrics shows that they have a high correlation rate with specific attacks such as port scan and dos. Therefore, they are highly relevant for enhancing detection capabilities. As a first evaluation of our approach, we apply our pipeline on the UGR'16 data set [4].

2 Preliminaries

2.1 Problem definition

In this paper, we address a specific attack detection problem, namely the detection using unsupervised detection models. The objective is to loosen the dependency on labeled historical data with attacks and to better perceive novel, abnormal behaviours. We more specifically focus on the relation between anomaly detection and the characterization of these anomalies as actual attacks [6]. Anomalies are defined as rare data points or outliers in the data set, and the features of the data points are the mean to highlight those outliers. Most of the time, the available features of the data are not sufficient to properly discriminate attacks from benign data [2]. This can be explained either by the fact that the features of the dataset do not discriminate attacks more than benign data or that the attacks are better characterized not by a single feature, but by a relationship between two features : identical values (as in IP loops where IP sources and destination are identical), differences between scalars (differences of throughput), etc. Another reason for the bad discrimination of attacks is that the amount of data is too high and then most outliers are in fact only statistical anomaly.

To address this problem, we rely on machine learning outlier detection algorithms, but we make use of graph representation and community detection as a mean to extract new

features. Those features are used to better discriminate attacks from other data, and thus to increase the detection rate for considered attacks. We then evaluate which graph community metrics are features relevant to attack detection.

2.2 Metrics for graph community analysis

The following metrics have shown different definitions in the literature [7] or none. We therefore provide here explicit definitions for those metrics :

Density in community i is the chance for any node inside i to be adjacent to another given node in i . It is defined as :

c_i : Number of connections in i

C_i : Maximum number of connections in i (if i were a clique)

$$Dens_i = \frac{c_i}{C_i} \quad (1)$$

Externality is not defined in the literature. It is similar to the "expansion" found in [7] and in community i is the proportion of communication of between i and others communities compared to any communication involving nodes in i as defined as :

M : The number of edge with at least one vertex in i

Me : The number of edge with exactly one vertex in i

$$Ext_i = \frac{Me}{M} \quad (2)$$

In addition to those metrics, we define local and global stability of community in dynamic graph as a contribution :

Local stability is a ratio of similarity between the state of a community C at time t : C_t , and $t + 1$: C_{t+1} defined as :

V_t : Set of nodes belonging to C_t

V_{t+1} : Set of nodes belonging to C_{t+1}

$$Ls_i = \frac{|V_t \cap V_{t+1}| - |(V_t \cap \bar{V}_{t+1}) \cup (V_{t+1} \cap \bar{V}_t)|}{|V_t \cup V_{t+1}|} \quad (3)$$

Global stability of a community C at time n is the mean of all local stabilities between time 0 and n defined as :

$$Gs_i = \frac{\sum_{i=1}^n Ls_i}{n - 1} \quad (4)$$

2.3 Performance of metrics extraction

Graph community metrics for dynamic graphs are extracted using sliding windows for a given time interval. Our approach shall be applicable to a real time data stream, and therefore should be scalable. In order to determine the performance of our community metrics extraction algorithm, we measure the time spent on its application on a time windowed graph and reproduce the process for increasing amounts of data (Figure 2).

3 Evaluation

3.1 Relevance of graph community to detection

Performance metrics are compared for different models using the Isolation Forest algorithm which shows the best performances on the same sample of the UGR'16 dataset [4] in Table 1. Except for the baseline which does not make use of feature selection (FS) and Hyper-parameter tuning (HPT), the same process is applied for the different models. The process is repeated 10 times and the average of each performance metrics are reported.

In addition to the comparison between the different models including those taking dynamic graphs community metrics such as stability, the attack distribution in relation to those metrics is observed.

3.2 Scalability of metrics extraction

In order to evaluate the scalability of our algorithms, we test them on samples of different size of a same week of the dataset UGR'16. The biggest sample entails about 539 millions data points extracted for one complete week on the target system. Every other sample is then an evenly distributed proportion of the complete week. The algorithm is applied on those samples and the time spent for the extraction of graph community metrics by our algorithms on each time windows of the dataset is observed.

4 Results

4.1 Relevance of graph community to detection

TABLE 1 – Performance evaluation of the chosen approach

Isolation Forest on UGR'16	F-Score	Precision	Recall
Baseline	0.00049984	0.00103121	0.00032987
Feature extraction(FE)	0.05641072	0.03454118	0.15377554
FE + Feature Selection(FS)	0.30761931	0.32660442	0.29073454
FE + FS + Hyper Parameter Tuning(HPT)	0.32155489	0.30152777	0.34444449
FE+FS+HPT+ IP Graph(5 min)	0.48283249	0.46699468	0.49990232
FE+FS+HPT+ Ip&(Ip,port) Graph (5, 10 & 20 min)	0.80421724	0.85484417	0.75928304
Previous+local stability	0.81157849	0.86104321	0.76750087

The results in Table 1 show that graph community metrics as features can significantly improve detection performances. We observe a F-score up to **0.804** with graph community metrics against **0.322** for common feature extraction. Moreover, we show that using different features, or combinations of features for node as well as extracting the community metrics for different time intervals lead to a significant improvement of performance. A **0.804 F-score** is observed, using both **IP** and **couple of IP and port** as node with **5, 10 and 20 minutes time interval** against a **0.423 F-score** using IP as node and 5 minutes time interval.

4.2 Relevance of graph community dynamics




FIGURE 1 – Distribution of positives and negatives in detection with regard to global stability of communities in dynamic graphs using IP and port as node. The green box shows where most attacks are located

In Table 1, we additionally observe that our model with dynamic graph community metrics is slightly better than the one with only static graphs metrics for every performance metrics, with in particular an F-score of **0.812**. However, this model only uses local stability. Nevertheless we can observe in Figure 1 that there are still unnoticed interesting behaviours in regard to global stability

4.3 Scalability of metrics extraction

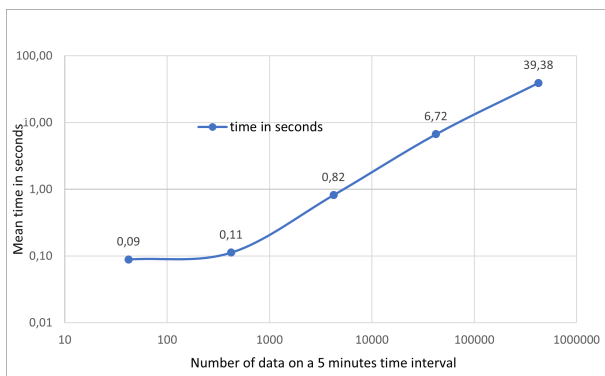


FIGURE 2 – Graph community metrics extraction time depending on the amount of data

As can be observed on figure 2, the graph community metrics extraction is sublinear in time complexity. The tendency of the time spent curve seems stable, with about 717% time increase for 900% data increase on average if we do not take the smallest sample in account. The speed-up observed of **25.52%** is assumed to be due to the amount of nodes in the graph not scaling as fast as the number of edges.

5 Conclusions

In this paper, we discussed on the importance of FE and FS when working with unsupervised detection algorithms. Graph based approaches are well suited for attack detection problems. As such, we propose two new metrics for dynamic graph communities with the local and global stability

and obtained a F-score of 0.812. Along others graph community metrics, in particular density and externality, we used them as input features for anomaly detection using the Isolation Forest algorithm and obtained encouraging results using a scalable approach. However while the approach is able to detect most of scan and dos attacks, it is unable to detect the other types of attacks (nerisbotnet, spam) in the dataset. Due to the ever-changeability and diversity of attacks in the data of our application case, we do not think it is possible to be able to detect every current type of attacks or new ones that will come to be. However we hope to be able to enhance the trustability in the positive detection in the future.

Références

- [1] A. Abou Rida, R. Amhaz, and P. Parrend. *Anomaly Detection on Static and Dynamic Graphs using Graph Convolutional Neural Networks*, chapter -, page 23. Studies in Computational Intelligence Series. Springer, 2022.
- [2] Siddharth Bhatia, Bryan Hooi, Minji Yoon, Kijung Shin, and Christos Faloutsos. Midas : Microcluster-based detector of anomalies in edge streams. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3242–3249, 2020.
- [3] Xavier Larriva-Novo, Víctor A. Villagrà, Mario Vega-Barbas, Diego Rivera, and Mario Sanz Rodrigo. An iot-focused intrusion detection system approach based on preprocessing characterization for cybersecurity datasets. *Sensors*, 21(2), 2021.
- [4] Gabriel Macià-Fernández, José Camacho, Roberto Magán-Carrión, Pedro García-Teodoro, and Roberto Therón. Ugr’16 : A new dataset for the evaluation of cyclostationarity-based network idss. *Computers & Security*, 73 :411–424, 2018.
- [5] J. Navarro, A. Deruyver, and P. Parrend. A systematic survey on multi-step attack detection. *Computers and Security*, page 102, 2018.
- [6] William Robertson, Giovanni Vigna, Christopher Krügel, and Richard Kemmerer. Using generalization and characterization techniques in the anomaly-based detection of web attacks. In *NDSS*, 01 2006.
- [7] Jaewon Yang and Jure Leskovec. Defining and evaluating network communities based on ground-truth. In *Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics*, MDS ’12, New York, NY, USA, 2012. Association for Computing Machinery.
- [8] Tommaso Zoppi, Andrea Ceccarelli, Tommaso Capecci, and Andrea Bondavalli. Unsupervised anomaly detectors to detect intrusions in the current threat landscape. *ACM/IMS Trans. Data Sci.*, 2(2), apr 2021.