RawPower: Deep Learning based Anomaly Detection from Raw Network Traffic Measurements

Gonzalo Marín
AIT Austrian Institute of Technology
IIE-FING, UDELAR, Uruguay
gonzalo.marin@fing.edu.uy

Pedro Casas
AIT Austrian Institute of Technology
pedro.casas@ait.ac.at

Germán Capdehourat
IIE-FING, UDELAR, Uruguay
gcapde@fing.edu.uy

ABSTRACT

Machine learning models using deep architectures (i.e., deep learning) have gained path in recent years and have become state-of-the-art in many fields, including computer vision, speech recognition and natural language processing. However, when it comes to network measurement and analysis, classic machine learning approaches are commonly used, heavily relying on domain expert knowledge. In this work, we explore the power of deep learning models to perform anomaly detection on network traffic data, taking as input *raw* measurements coming directly from the stream of monitored bytes. Our initial results suggest that deep learning can enhance anomaly detection without requiring expert domain knowledge to handcraft input features.

CCS CONCEPTS

• Security and privacy → Intrusion/anomaly detection and malware mitigation; • Computing methodologies → Machine learning;

KEYWORDS

Deep Learning; Anomaly Detection; Network Traffic Measurements and Analysis

ACM Reference Format:

Gonzalo Marín, Pedro Casas, and Germán Capdehourat. 2018. RawPower: Deep Learning based Anomaly Detection from Raw Network Traffic Measurements. In *ACM SIGCOMM 2018 Conference Posters and Demos (SIGCOMM Posters and Demos '18)*, August 20–25, 2018, Budapest, Hungary. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3234200.3234238>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCOMM Posters and Demos '18, August 20–25, 2018, Budapest, Hungary

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5915-3/18/08...\$15.00

<https://doi.org/10.1145/3234200.3234238>

1 INTRODUCTION

Deep learning (DL) models are today widely used in multiple signal processing problems, particularly in image processing, where they have shown an outstanding performance. The popularity of DL has seen tremendous growth due in large part to the increasing computational power, larger datasets (i.e., more training data available) and improvements in the techniques associated to train DL models. One of its most powerful characteristics is their ability to learn feature representations from input raw or basic, non-processed data. For example, a deep convolutional neural network (CNN) trained for image classification can learn to recognize edges and more complex structures along the sequence of neural layers, using as input only the image RGB pixel values [1].

However, neural-networks based models, and particularly DL models, have an inherent problem linked to model visibility and interpretation: a DL model is a black-box which can automatically create features from input raw data and provide highly accurate predictions, but it is difficult to understand its functioning. Indeed, it becomes challenging to understand the input features leading to a particular result, as these can be in general meaningless to a domain expert. This is one of the reasons why their application to networking problems is so far quite limited as compared to other domains. As a consequence, classic machine learning (ML) models are the most commonly applied in the network measurement domain. When building these models, a feature vector of expert-handcrafted features is usually built to achieve the best results. This is actually the critical step on which the success of the model depends.

There are nevertheless two problems when addressing network measurement problems using classic ML approaches. First, the lack of a consensual set of input features to tackle specific targets, such as network security, anomaly detection, traffic classification, etc.; second, the continuous changing in network measurements statistics that may cause static handcrafted features to fail. To improve these limitations, we explore in this paper the end-to-end application of DL models to complement traditional approaches for network measurement analysis.

Some DL approaches have been recently shown good performance, mainly associated to traffic classification tasks. In

Figure 1: The *RawPower* architecture.

2015, Z. Wang *et al.* [2] presented a deep neural network (DNN) for feature learning using feed-forward networks and Stacked Auto-Encoders (SAE) to perform network protocol recognition. In 2017, W. Wang *et al.* presented two models based upon CNNs [3, 4], transforming network flows into images to work as an input for the models. Also in 2017, Lotfollahi *et al.* [5] presented an approach for encrypted traffic classification using SAE and 1D-CNN at the packet level.

The main limitation of these proposals is that they all require some sort of data pre-processing as first step to extract meaningful input for the DNN models. Our goal is to train a DNN model with the stream of incoming bytes without requiring any pre-processing step or domain expert intervention, to make the approach generic and flexible. In particular, we build *RawPower*, a multi-layer DNN model for network traffic analysis, and apply it to the problem of malware detection in network data.

2 THE RAWPOWER DNN

As a first approach to the problem, we build a 1D-CNN based model to perform malware detection at the packet level using the USTCTFC2016 dataset [4]. This dataset consists of two groups of labeled pcap files from malware and benign applications. Since our main goal is to build a DNN classifier from raw data, we take every packet from each labeled pcap as part of the dataset, using decimal normalized representation of every byte as a different feature. Since different packets can have different sizes, we set a fixed threshold T_s to trim each incoming packet to the first T_s bytes, after removing irrelevant byte information such as MAC and IP addresses. In our study we set $T_s = 1300$ bytes; all packets with size larger than 1300 bytes are trimmed, and packets with smaller size are zero-padded. Each packet is finally labeled as either benign or malware, i.e., we consider a binary classification problem.

Fig. 1 depicts the architecture of the *RawPower* DNN. The model consists of two 1D-CNN layers with ReLU and max-pooling as activation functions, plus two feed-forward layers

Figure 2: Malware detection performance.

of 200 units each. At the end, a binary logistic regression classifier is used. Each 1D-CNN consists of 32 1D-filters of size 5. Spatial and normal batch normalization layers as well as dropout were added to the model in order to ease the training process and add regularization. We use the Keras framework running on top of TensorFlow to build the model.

Fig. 2 presents the initial results obtained by the *RawPower* model in the detection of malware packets, in the form of a ROC curve. The model is compared to a Random Forest (RF) one, using exactly the same input features and an internal architecture of 100 trees. We choose a RF model based on the good detection performance shown by the model in our previous work [6]. *RawPower* can detect more than 70% of the malware instances with a false alarm rate below 3%, largely outperforming the RF model. Indeed, when applying the RF model with the same set of raw, non processed input features, the obtained results are very poor. These results are highly encouraging, as they point to the ability of the *RawPower* model to better capture the underlying statistics of the malware, without requiring any specific handcrafted feature set. We are currently exploring the application of *RawPower* to raw inputs derived from network flows, as well as investigating other network architectures.

REFERENCES

- [1] K. Simonyan et al, "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps," in <http://arxiv.org/abs/1312.6034>, 2013.
- [2] Z. Wang, "The Applications of Deep Learning on Traffic Identification," *Black Hat USA, Las Vegas*, 2015.
- [3] W. Wang et al, "End-to-end encrypted traffic classification with one-dimensional convolution neural networks," in *IEEE ISI*, 2017.
- [4] W. Wang et al, "Malware traffic classification using convolutional neural network for representation learning," in *ICOD*, 2017.
- [5] M. Lotfollahi et al, "Deep Packet: A Novel Approach For Encrypted Traffic Classification Using Deep Learning," in <http://arxiv.org/abs/1709.02656>, 2017.
- [6] P. Casas et al, "(Semi)-Supervised Machine Learning Approaches for Network Security in High-Dimensional Network Data," in *ACM CCS*, 2016.