RADemics

Autoencoders for Unsupervised Intrusion Detection in High Dimensional Security Data

K.Saranya, B.Persis Urbana Ivy, N.Kavitha

Rajalakshmi Institute of Technology, Mother Teresa Institute of Engg and Technology,Melumoi (Post), Vels Institute of Science, Technology & Advanced Studies (VISTAS),

Autoencoders for Unsupervised Intrusion Detection in High Dimensional Security Data

[1]K. Saranya, Assistant Professor, CSE, Rajalakshmi Institute of Technology, Chennai, pksaranya@gmail.com,

[2]B. Persis Urbana Ivy, Dean (CSE & Allied branches), Mother Teresa Institute of Engg and Technology, Melumoi (Post), Palamaner – 517408. Mail ID: sperurban.mti@gmail.com

[3]N. Kavitha, Assistant professor, Artificial intelligence and machine learning, Vels Institute of Science, Technology & Advanced Studies (VISTAS), kavitha.sec@vistas.ac.in

Abstract

Intrusion detection systems (IDS) play a crucial role in safeguarding modern network environments from a wide array of cyber threats. As networks evolve, traditional detection methods based on signature matching or rule-based techniques have proven inadequate in identifying novel and sophisticated attacks. This chapter explores the application of autoencoders for unsupervised anomaly detection, focusing on their effectiveness in high-dimensional security data. The chapter delves into threshold selection strategies, which are vital in distinguishing between normal and anomalous network behavior. A particular emphasis is placed on the challenges of balancing false positives and false negatives, which significantly impact detection performance. Additionally, the chapter examines dynamic thresholding techniques, including incremental learning and adaptive calibration, which enable real-time detection and response to evolving network conditions. The integration of multi-feature data and multi-dimensional thresholding strategies is explored, demonstrating how autoencoders can capture complex patterns and enhance detection accuracy in diverse cybersecurity scenarios. Finally, the chapter provides a comprehensive analysis of the practical implementation of these techniques in real-world IDS, highlighting their potential to significantly improve intrusion detection and reduce detection latency. This work contributes to advancing the field of IDS, offering a robust framework for designing adaptive, scalable, and efficient security systems.

Keywords: Intrusion Detection Systems, Autoencoders, Anomaly Detection, Threshold Selection, Dynamic Thresholding, Incremental Learning.

Introduction

In the contemporary landscape of network security, the need for efficient and scalable Intrusion Detection Systems (IDS) is more critical than ever [1]. Cyber threats have evolved significantly, becoming more sophisticated and diverse, often bypassing traditional security mechanisms [2]. Signature-based detection, which has long been a staple of IDS, struggles to keep up with this ever-changing threat landscape, as it is reliant on predefined patterns and does not account for novel or previously unseen attacks [3]. As cybercriminals increasingly employ complex attack strategies, traditional methods often fail to detect emerging threats in real-time, leaving organizations vulnerable. In response to these challenges, there is a growing interest in leveraging

unsupervised machine learning techniques, particularly autoencoders, for anomaly detection in network traffic [4]. Autoencoders offer a unique advantage in that they do not require labeled datasets to function, making them particularly effective in dynamic environments where labeled data may be limited or absent. Their ability to model normal network behavior and detect deviations has made them a promising tool in the development of next-generation IDS [5].

One of the central challenges in applying autoencoders for intrusion detection is selecting an appropriate threshold that balances sensitivity and specificity [6]. The reconstruction error, which represents the difference between the original input and the reconstructed output of the autoencoder, serves as an indicator of whether an observed behavior is anomalous [7]. A threshold must be set to determine when this error is significant enough to flag a potential intrusion [8]. However, this threshold must be carefully calibrated, as selecting a threshold that is too low can lead to an overwhelming number of false positives, where normal network activities are incorrectly flagged as anomalous [9]. Conversely, setting the threshold too high can result in false negatives, where actual intrusions are not detected because they do not cause enough deviation from normal behavior. Achieving the right balance in threshold selection is critical for ensuring the practical viability of autoencoder-based IDS [10].

The complexity of real-world network environments further complicates the threshold selection process [11]. Networks today are not only vast and highly dynamic but also multifaceted, with different types of traffic, protocols, and user behaviors [12]. As a result, the data generated by these networks is multi-dimensional and can include a wide range of features, such as traffic volume, latency, and application-specific metrics [13]. In multi-feature environments, the challenge of threshold selection becomes more pronounced, as each feature can have its own distribution and variability. For instance, network traffic can vary greatly based on the time of day, user activities, and network conditions [14]. To handle this complexity, advanced multi-dimensional thresholding techniques are needed to integrate the behavior of multiple features into a unified decision-making process. These techniques help ensure that the system can adapt to changes in network conditions while accurately identifying anomalies across different features of the data [15].