E-ISBN: 978-93-7020-394-5
IIP Series, Chapter 7

ROLE OF MACHINE LEARNING IN ANOMALY DETECTION AND INCIDENT RESPONSE

Abstract

One area of cybersecurity called "AI-powered incident response automation" makes use of machine learning (ML) and artificial intelligence (AI) to speed and simplify the process of responding to security issues. AI's sophisticated threat detection, analysis, and reaction automation capabilities can greatly improve incident response. AI-powered solutions may quickly identify security events by continually monitoring networks, systems, and endpoints to identify anomalous activity or possible attacks in real time. AI can also evaluate enormous volumes of security data to find trends, patterns, and abnormalities that can point to malicious activity, which aids security teams in efficiently prioritizing and countering threats. Additionally, by automatically carrying out predetermined tasks like separating compromised systems, preventing malicious traffic, and applying patches or updates, AI-driven automation can expedite incident response procedures, cutting down on response times and lessening the impact of security. One essential component of data science is anomaly detection, sometimes referred to as outlier detection, which focuses on finding odd patterns that deviate from expected behaviour. By evaluating and contrasting data points within a collection, an anomaly detection system can identify those that deviate from the typical trend. Finding statistical oddities is only one aspect of AI's importance in anomaly detection; another is revealing important insights, underlying issues, or possibilities that could otherwise go overlooked. The main aim of this topic is to explain the benefits of Machine learning in case of anomaly detection and incident response.

Keywords: Machine Learning, Anomaly, incident, detection.

Authors

Dr. Sukhdev Singh
Assistant Professor
Department of Computer Science
D.A.V. College, Lahore, Ambala City
Haryana.
sukhdev_kuk@rediffmail.com

Namit
D.A.V. College, Lahore, Ambala City
Haryana.
ndogra757@yahoo.com

Rahul Gupta
D.A.V. College, Lahore, Ambala City
Haryana.
rgaeron@gmail.com

I. INTRODUCTION

The abilities of anomaly detection systems are enhanced by the use of ML and AI approaches, which allow them to identify dangers that were earlier undiscovered and adjust to developing attack patterns. Unsupervised machine learning has become one of the most extensively utilized AI processes in case of identification of anomaly. Unsupervised machine learning algorithms use a variety of methods, including density estimation, dimensionality reduction, and clustering, to find patterns or behaviours that differ from a system's typical or expected behaviours.

AI incident response is the process of finding, mitigating, and resolving security events by automating, improving, and streamlining incident response procedures with AI technologies. It makes use of machine learning and other cutting-edge technology to find the source of problems, detect anomalies, and automatically take action or give security teams insights.

II. ANOMALY DETECTION

Finding patterns or behaviours that differ from a system's typical or anticipated behaviour is known as anomaly detection. By identifying possible security concerns, it plays a crucial part in cybersecurity. In order to discover anomalies, profiles of typical user behaviour are created, actual user activity is compared to those profiles, and departures from the norm are flagged. The idea behind anomaly detection is that unusual patterns of behaviour point to system abuse. A collection of metrics is used to define profiles. Metrics are measurements of specific user behaviour characteristics. Every metric has a threshold or range of values attached to it. The premise behind anomaly detection is that people use the system in predictable, consistent ways. Additionally, the method allows for adjustments in response to evolving user behaviour. Since no one is certain that a particular collection of metrics is comprehensive enough to capture every aberrant behaviour, the completeness of anomaly detection has not yet been confirmed. Therefore, more research is needed to determine if anomaly detection—a powerful safeguard for systems—will ever be able to identify all cases of relevance.

A theoretical definition of an anomaly is a pattern that diverges through the distinctive anticipated behaviour. Three primary classifications are used to classify anomalies.

1. **Point Anomalies:** Point anomalies are thought to be the most fundamental category of anomaly where a single data illustration can be considered abnormal for the rest of the data. When defining the non-contextual aspects of an illustration, the second property is counted as an attribute of behaviour. A behavioural characteristic would be, for illustration, the quantity of rainfall that falls at any given place in a spatial dataset that characterizes the average rainfall worldwide. The significance of contextual anomalies in the target area influences the selection for adopting the contextual anomaly detection technique. Another significant constituent is the availability of qualitative qualities. It makes sense to utilize a technique of contextual detection in situations where identifying a context is simple. In other cases, it is impractical for creating an impression that some techniques are challenging to employ.

2. **Contextual Anomalies:** Such anomaly occurs when a data illustration is abnormal in one context but not in another. The two types of contextual anomalies are contextual and behavioural characteristics. The context (or neighbourhood) of an illustration is ascertained by applying the first feature. In geographic datasets, for instance, a location's longitude and latitude are contextual properties. Furthermore, the position of an instance on the entire sequence is determined by time, a contextual property in time series data.

3. **Collective Anomalies:** It occurs whenever a group of linked data illustrations is unusual for the dataset as a whole.

Among the earliest methods for identifying abnormalities are approaches of statistical anomaly detection. A statistical model representing the distinctive behaviour of the given data is produced utilizing statistical methods. To find out if an illustration fits into this model, a statistical conclusion test may then be performed. Statistical anomaly detection is done in a number of ways. This contains parametric, non-parametric, semi-parametric, and proximity-based approaches. One procedure of recognizing abnormalities is the rising use of machine learning (ML) techniques. The endeavour to "automate the procedure of awareness attainment through instances" is known as machine learning. Using this method, a model that distinguishes among normal and abnormal classes is constructed.

Thus, on the basis of the training data function that was used to construct the model, anomaly detection can be split into three major groups. The three major categories are:

1. **Supervised Identification of Anomalies:** Labelled cases can be found in both the normal and anomalous training datasets for this class. The method used in this model is to construct a forecasting model for both the normal and anomalous classes, after which the two models are compared. But there are two issues in this mode. Firstly, in contrast to distinctive cases, there are far less abnormalities in the training set. Second, it is hard to find accurate and representative labels, mostly for the anomaly class.

2. **Semi-Supervised Anomaly Detection:** Distinctive class cases are only utilized for training. Anything that cannot be categorized as normal is therefore labelled as abnormal. Semi-supervised methods presume that only the normal class has labelled instances in the training data. They are more popular compared to supervised approaches since they do not require anomalous class labels.

3. **Unsupervised Anomaly Detection:** In this situation, the techniques do not need training datasets. As a consequence, those approaches propose that in test datasets, regular instances are far more prevalent than abnormalities. On the other hand, this technique has a significant false alarm rate if the supposition is erroneous.

Making use of unlabelled dataset samples as training data, many semi-supervised approaches can be altered to perform in an unsupervised mode. This kind of adaptation is postulated on the test data having very few anomalies and those abnormalities being resilient to the model's learning during training.

III. MACHINE LEARNING ANOMALY DETECTION: IDENTIFYING OUTLIERS TO IMPROVE BUSINESS OPERATIONS

1. **Supervised Learning:** Real-world input and output data are used by supervised learning algorithms to identify abnormalities. Data analyst has to categorize data points as normal or abnormal to be utilized as training data for these categories of anomaly detection systems. On the basis of instances issued, a machine learning model trained with labelled data can recognize outliers. This kind of machine learning performs well to recognize known outliers, but it cannot recognize novel abnormalities or predict issues in the future.

   The following are typical machine learning algorithms for supervised learning:

   - **KNN Algorithm:** The K-nearest neighbour (KNN) algorithm is a regression modelling or density-based classifier that is used to identify anomalies. One statistical method for determining the link between labelled and variable data is regression modelling. It works on the premise that comparable data points will be located close to one another. A data point is deemed anomalous if it appears farther distant from a dense region of points.

   - **LIF:** LIF, or local outlier factor: As a density-based approach, local outlier factor is similar to KNN. The foremost difference is that LOF bases its findings on the data points that are farthest apart, whereas KNN bases its presumptions on the data points that are near to one another.

2. **Unsupervised Learning:** Unsupervised learning methods can handle more complicated data sets and don't need labelled data. Deep learning, neural networks, or auto encoders—which imitate the signals sent by organic neurons to one another—power unsupervised learning. These potent technologies are able to identify patterns in incoming data and draw conclusions about what constitutes normal data. These methods can significantly lessen the effort required to manually sort through big data sets and find unforeseen anomalies. Data scientists should, however, keep an eye on the outcomes of unsupervised learning. These methods have the potential to inaccurately diagnose anomalies since they are based on assumptions about the data being entered.

   Among the machine learning techniques for unstructured data are:

   - **K-Means:** The goal of the K-means method, a data visualization approach, is to cluster comparable data points by processing them through a mathematical equation. The points in the cluster's centre to which all other data is related are referred to as "means," or average data. These clusters can be utilized to identify patterns and draw conclusions about data that is discovered to be unusual through data analysis.

   - **Isolation Forest:** Unsupervised data is used in this kind of anomaly detection technique. This method aims to isolate anomalies as the initial step, in contrast to supervised anomaly identification methods that operate from labelled normal data points. It generates "decision trees," which map out the data points and choose a region at random for analysis, much like a "random forest." According to its location in relation to the other points, each point is given an anomaly score between 0 and 1,

with values below. This process is repeated, those above 5 are more likely to be unusual, while those below that threshold are typically regarded as normal. Scikit-learn, a free Python machine learning library, contains isolation forest models.

- **One-Class Support Vector Machine (SVM):** This anomaly detection method draws boundaries around what is deemed normal using training data. Points that are clustered inside the designated borders are referred to as normal, whereas those that are outside are called anomalies.

3. **Semi-Supervised Learning:** The benefits of the first two approaches are combined in semi-supervised anomaly detection techniques. Engineers can work with unstructured data and automate feature learning making use of unsupervised learning techniques. They may, however, keep an eye on and supervise the kinds of patterns the model picks up by merging it with human oversight. Generally, this enhances correctness of the model's forecasting.

- **Linear Regression:** Both dependent and independent variables are used in linear regression, a predictive machine learning tool. Utilizing a set of statistical equations, the independent variable serves as the basis for computing the value of the dependent variable. When only a portion of the data is known, these equations use both labelled and unlabelled data to predict future events.

IV. USE SCENARIOS FOR ANOMALY DETECTION

In many different industries, anomaly detection is a crucial tool for sustaining corporate operations. The kind of data being gathered and the operational problem being resolved will determine whether supervised, unsupervised, or semi-supervised learning methods are used. Use cases for anomaly detection include, for example:

1. Use Cases for Supervised Learning

- **Shops:** Future sales targets can be predicted with the use of labelled data from the sales totals of the prior year. Setting goals for certain salespeople based on their prior success and the demands of the business as a whole can also be beneficial. All sales data is known, so trends can be examined to gain knowledge about seasonality, marketing, and products.

- **Forecasting the Weather:** Supervised learning algorithms can help anticipate weather patterns by utilizing previous data. Forecasters can make more accurate predictions that account for changing conditions by analysing recent data on temperature, wind speed, and barometric pressure.

2. Use scenarios for Unsupervised Learning

- **Intrusion Detection System:** These devices, which can be either hardware or software, keep an eye on network traffic for indications of criminal behaviour or security breaches. In order to safeguard user data and system operations, machine learning algorithms can be trained to identify possible network assaults in real time.

These algorithms, which examine data points at predetermined intervals over an extended period of time, can produce a depiction of typical performance based on time series data. Unexpected patterns or spikes in network traffic can be identified and investigated as possible security breaches.

- **Producing:** Maintaining supply chains, maximizing quality assurance, and producing goods all depend on correctly operating machinery. By using unlabelled data from sensors mounted on machinery, unsupervised learning algorithms can be utilized for predictive maintenance to forecast probable faults or failures. This minimizes machine downtime by enabling businesses to make repairs prior to a serious malfunction**.**

3. Use Cases for Semi-Supervised Learning

- **Medical:** Images containing recognized diseases or disorders can be labelled by medical practitioners using machine learning techniques. However, it is impossible to identify every probable cause for concern because visuals will differ from person to person. These algorithms can interpret patient data, draw conclusions from unlabelled photos, and identify possible causes for concern once they have been taught.

- **Fraud Detection:** Semi-supervised learning, which requires both labelled and unlabelled data, can be used by predictive algorithms to identify fraud. The labelling of a user's credit card activity makes it possible to identify odd spending trends. Fraud detection systems, however, are not limited to transactions that have already been flagged as fraudulent; they can also draw conclusions from user activity, such as current location, log-in device, and other elements that call for unlabelled data.

V. INCIDENT RESPONSE

Attacks on an organization are becoming more frequent by the day. Installing an IDS or IPS without any objectives or response strategy is as dangerous as not installing any systems at all. Monitoring and reacting to network intrusions is a highly difficult undertaking that requires preparation. You must understand how the attack happened, when it happened, what the attacker did, and how you should react. Establishing an incident-response procedure is essential. IDSs can be deployed in two methods to identify incidents: intrusion detection and attack detection. Installing a sensor outside the firewall to log attack attempts is known as *attack detection*. Monitoring the quantity and kinds of assaults on your network may benefit from this. Attack detection will record the attack for use in identifying security requirements and analysing attack types, assuming that the network perimeter is safe. This kind of arrangement has the drawback of creating a large number of log files that are frequently disregarded and not utilized for their intended purpose.

Conversely, *intrusion detection* uses a sensor that is positioned inside the network. Any attacks that are detected could occur inside the perimeter that is protected.

VI. WORKING OF AI IN INCIDENT RESPONSE

AI in incident response works by detecting, analysing, and reacting to security events instantly using automation and machine learning. AI systems, in contrast to conventional

techniques, continuously learn from data to increase accuracy and speed up reaction times. This is an explanation of how AI works in this situation:

1. **Data Ingestion and Normalization:** Logs, network traffic, and threat intelligence feeds are just a few of the types of data that AI systems gather. After that, this data is standardized to guarantee analytical consistency.

2. **Anomaly Detection:** AI detects departures from typical behaviour, indicating possible security threats, using pattern recognition and machine learning techniques. This aids in incident detection early on.

3. **Event Correlation:** AI analyses data from various sources to find trends that point to intricate multi-phase attacks. This makes it possible to detect threats more precisely.

4. **Automated Incident Triage:** By classifying incidents according to their level of severity, AI streamlines the triage process, relieving security professionals of some of their workload and guaranteeing prompt answers.

5. **Root Cause Analysis:** By examining data and identifying the issue's origin, artificial intelligence (AI) speeds up the process of determining the underlying cause of occurrences. This shortens the time needed to resolve issues and helps stop them from happening again.

6. **Response Automation:** AI solutions increase response efficiency by automatically carrying out predetermined incident response activities, such as blocking hostile IPs or isolating compromised computers.

7. **Continuous improvement:** AI models are always learning from previous events to enhance detection capabilities and optimize reaction procedures, which leads to more proactive security measures.

VII. APPLICATIONS FOR AI-POWERED INCIDENT RESPONSE

Several phases of managing security events can benefit from the use of AI-driven incident response. These applications show how AI improves incident management's detection, response, and learning procedures.

1. **Identifying and Warning:** AI systems provide real-time identification of anomalies and possible threats by continually monitoring network traffic, records, and user activity. Early detection of anomalous patterns using machine learning algorithms sets up automated alarms that alert security professionals to potential problems.

2. **Analysis of Root Causes (RCA):** By swiftly evaluating data and connecting events across systems, artificial intelligence speeds up the root cause investigation process. With the use of this automated method, security teams may more quickly and accurately determine the root causes of occurrences, facilitating prompt remediation and incident prevention.

3. **Automation and Incident Resolution:** By carrying out preset tasks like patching, blocking malicious traffic, or isolating affected systems, AI-driven technologies simplify incident resolution. By utilizing incident response automation with Blink and Panther, you may drastically cut down on reaction times and the requirement for human participation during crucial stages of incident handling.

4. **Analysis and Learning after an Incident:** Security teams can learn from the past thanks to AI systems' ongoing post-event analysis. AI helps organizations proactively enhance their security posture and continuously improve incident response techniques by analyzing event data to find patterns and weaknesses in current defenses.

VIII. AI-POWERED INCIDENT RESPONSE'S ADVANTAGES

AI-powered incident response has several benefits that improve security operations' efficacy. The most significant advantages are shown below.

1. **Enhanced Incident Detection and Response:** By instantly evaluating vast amounts of security data, artificial intelligence (AI) greatly speeds up detection and response times and can spot dangers that conventional techniques might overlook.

2. **Faster Reaction Times:** By automating incident response procedures, AI technologies enable businesses to respond to security events quickly and reduce possible harm.

3. **Faster Root Cause Analysis:** AI speeds up the process of determining the underlying cause of security events, enabling quicker resolution and lowering the possibility that such incidents may recur in the future.

4. **Increased Accuracy:** By precisely examining patterns and anomalies, AI reduces false positives and negatives, allowing security professionals to concentrate on actual risks.

5. **Scalability:** AI-powered systems are perfect for growing infrastructures since they can grow to handle massive data volumes and an increasing number of occurrences without requiring more human resources.

6. **Cost Savings and Resource Optimization:** AI helps lower operating costs and maximize the use of existing resources by automating operations and increasing productivity. This frees up security teams to concentrate on more important responsibilities.

REFERENCES

[1] Ali Bou Nassif; Manar Abu Talib; Qassim Nasir; Fatima Mohamad Dakalbab,” Machine Learning for Anomaly Detection: A Systematic Review”, IEEE Access (Volume 9), DOI: 10.1109/ACCESS.2021.3083060, ISSN: 2169-3536
[2] Nachaat Mohamed, "Current trends in AI and ML for cybersecurity: A state-of-the-art survey", Article: 2272358,https://doi.org/10.1080/23311916.2023.2272358
[3] https://www.ibm.com/think/topics/machine-learning-for-anomaly-detection
[4] https://www.blinkops.com/blog/ai-incident-response
[5] Carl Endorf, Eugene Schultz, and Jim Mellander, “Intrusion Detection and Prevention”, McGraw-Hill/Osborne, ISBN: 0-07-222954-3

