BLUVECTOR®
A COMCAST COMPANY

How to Get the Right Machine Learning for Your Organization

Machine learning (ML) algorithms process massive amounts of data to extract patterns and build models. These models can then be used to rapidly classify new data based on the features and patterns extracted from the original dataset.

These capabilities mean that ML is uniquely suited to addressing the core challenge of cybersecurity: identifying malicious content and potential attacks within massive, noisy datasets. The use of ML in cybersecurity is a rapidly growing market, with **69% of organizations** claiming that ML and artificial intelligence (AI) will be essential for responding to cyberattacks in the future.

ML can be an invaluable tool; however, when evaluating ML-driven cybersecurity solutions, it is important to keep in mind that no ML algorithm can provide 100% perfect threat detection. False positives and false negatives are an intrinsic part of how these solutions operate and are impossible to eliminate completely. When tuning machine learning tools, organizations should focus on how to best balance false positive and false negative rates to ensure that a potential solution meets their business needs.

What's the Problem?

Machine learning and artificial intelligence are buzzwords. Everyone claims to be using them, but definitions and quality vary greatly. Customers looking for ML or AI may have limited ability to determine whether or not they are actually getting it.

Lack of Common Standards

A major challenge is that the market lacks a robust framework for evaluating machine learning and weighing its costs and benefits. Such a framework should lay out standardized test cases that evaluate various solutions' ability to detect certain types of threats.

In general, marketing for ML and AI solutions is based on cherry-picked numbers drawn from carefully controlled experiments. Often, these experiments have limited real-world value since an algorithm trained on a particular data set will be very effective at classifying a very similar one. If the training and validation datasets are collected from the same source via the same methods, then the results of the evaluation are unlikely to generalize well.

The Challenges of Evaluating ML Solutions

Customers looking to evaluate an ML-based solution want to know that it can detect the threats that they are facing. To do so, they want to test AI/ML engines based on their ability to detect provided, known-bad samples.

However, this is not the best way to test the effectiveness of an AI/ML solution. ML is designed to identify novel threats, and any signature-based detection model will likely be able to identify these known malware samples.

A better test is to evaluate detection engines using novel malware samples. However, this is more challenging because it requires finding malware samples with no known signatures. A good (but uncommon) approach to accomplishing this is to use malware with known signatures that were created after the model was finalized.

A Need for Understanding

Even for black-box algorithms, the inputs and outputs must be clearly understood. This is true for several reasons, including:

- Maximizing value for purchasers
- Ensuring that the right tools are applied to the right problems
- Applying the correct ML approach to different environments
- Placing the expectation on vendors to provide transparency and not cherry-picked numbers

BLUVECTOR®
A COMCAST COMPANY

How to Judge Machine Learning Tools?

Evaluating machine learning tools is possible. However, it is rarely done correctly because it is difficult to generalize results across customers. Marketing is easier when you can point to a single number.

The first step is to define the inputs and outputs of the algorithm, which sets the rules for the evaluation and provides a framework for interpreting the results. This is relatively simple, and the market knows to do this.

Properly defining the metrics for the test is less common and more difficult. An effective test of a machine learning engine measures both:

- False Positives: A benign event is incorrectly labeled as a threat
- False Negatives: A threat is incorrectly labeled as a benign event

Measuring both false positive and false negative rates is equally important because both create costs for an organization. A high false positive rate contributes to wasted time and the alert overload faced by many corporate security teams. On the other hand, a high false negative rate can result in undetected breaches and a false sense of security because the organization isn't detecting actual attacks against its systems. Moreover, false negatives can cause organizations to misallocate security resources and funds instead of directing them to authentic threats to their systems.

A Good Tool Balances False Positives and False Negatives

It is uncommon for evaluations of machine learning algorithms to measure and report both false positive and false negative rates. This is because it is difficult to achieve a simultaneously low false positive and false negative rate since these two measurements are inherently in conflict with one another.

The false-positive and false negative rates of an algorithm measure its sensitivity to threats. It is very easy to develop an algorithm that provides a perfect score for one or the other independently:

- False Positives: False positives can be eliminated by flagging nothing as a threat.
- False Negatives: By labeling everything as a threat, an algorithm can guarantee a zero false-negative rate.

An algorithm with zero false positives and zero false negatives is probably over-trained on a particular dataset. A good, generalized ML algorithm will balance false positives and false negatives in some way, accepting some false positives in exchange for the ability to detect most real threats or decreasing alert volume but acknowledging that some incidents will be missed.

The "right" balance between false positives and false negatives depends on an organization's mission and risk appetite. This is true for both ML-based and signature-based algorithms.

For example, a system protecting mission-critical systems might consider a higher false positive rate acceptable if all real attacks are detected and stopped. In contrast, threat detection engines evaluating requests to a web application may need to let some threats slip by to keep alert volumes manageable.

Effectively Testing False Positive and False Negative Rates

A good evaluation of a threat detection solution measures both false positive and false negative rates. However, these two metrics need to be measured in different tests.

Measuring False Positives

A false positive rate can only effectively be evaluated in a real-world scenario. A lab environment is unable to accurately emulate how a solution would behave with real samples. The only way to accurately measure this is by deploying the solution at scale with known-good traffic.

This type of evaluation is difficult to generalize because each organization's network and traffic profile is different. This means that each organization may experience a slightly different false positive rate for the same solution.

Without a clear portrayal of the test environment, data on false positive rates or other numbers in marketing literature should be taken with a grain of salt. Instead, look for testimonials and proof of value from existing customers that show that a solution generalizes well and has a low false positive rate across multiple different environments.

BLUVECTOR®
A COMCAST COMPANY

Measuring False Negatives

False negative rates measure how frequently a solution misses a real attack. The best approach to judge false negative rates is to use known-bad content that accurately simulates the threats that an organization is likely to face.

False negative testing is easier to perform and more generalizable than false positive testing. Still, it is a good idea to look for an independent third-party evaluation of a solution's false negative rate. As mentioned before, evaluating a machine learning approach against cherry-picked samples provides little value since a signature-based algorithm can detect these as well.

In fact, ML solution vendors can influence the results of a false negative evaluation by selecting a dataset that contains only known threats that their engine is effective at detecting. However, there is an underserved market for providing inexpensive validation of ML approaches against both benign and malicious activities.

Choosing the Right ML Solution

When evaluating potential ML solutions, it is essential to know what to look for. Marketing that promises 100% detection with zero false positive detections should be considered with suspicion.
Often "false positive" is considered a "scary word" because security teams are accustomed to solutions that create a lot of noise while providing minimal value. However, it is important to keep in mind that no solution can deliver a perfect result. A solution with no false positives is either tested only on carefully curated data sets or has a high false negative rate instead.

False positives and false negatives are a balance, and the right solution has a performance profile that meets the unique needs of your organization and the solutions that it will be deployed to protect. It may even offer the ability to tune this balance in operation. When considering cyber threat detection solutions, weigh the costs of false positives and false negatives and choose accordingly. Both can result in missed detections, whether by overwhelmed security teams ignoring alerts or a lack of alerts altogether due to a missed detection.

Evaluating a machine learning solution requires going beyond the marketing. Look for testimonials, proof of value from customers, and independent evaluations of false negative rates. Also, ask the solution provider about the configurability of their solutions. The ability to tune the detection sensitivity - and thus the false positive and false negative rates - is an important feature that can help you choose the best machine learning solution for your organization.

All too often, companies have buyer's regret after picking an ML solution that doesn't meet their needs. Don't choose a solution until you've asked all of the right questions and received the right answers.