AI-Driven Threat Intelligence Platforms for Predictive Cyber Defense and Zero-Day Vulnerability Mitigation

Satyanarayana Gadiraju *satyanarayana.gadiraju@ieee.org*
Independent Researcher
Avenel, NJ, 07001

Sauhard Bhatt *sauhard.bhatt@ieee.org*
Independent Researcher
Cumming, GA, 30041

Abstract

Traditional, reactive security postures are no longer enough to defend against the increasingly sophisticated cyber threats we face – particularly zero-day attacks. This article presents an AI-Based Threat Intelligence Platform that provides predictive posture for Cyber Defense. The system goes beyond signature-based detection by collecting and analyzing comprehensive, multi-modal datasets to predict and neutralize threats before they materialize. In this paper, we propose a new approach which combines internal network behavior analysis with external unstructured data intelligence. The performance of the platform was demonstrated on a dedicated curated dataset, ZDA-NetTraffic-459, including 459 examples representing both normal traffic and artificial zero-day exploit signatures. The prototype was Python-based and utilized state-of-the-art computational libraries for data crunching and pattern recognition, using a graph database to model multi-dimensional interactions between threat-actors. We demonstrate that our AI-based system can maintain the same 94.5% general accuracy in predicting of novel threats without the need for reannotation, therefore effectively reducing detection time from days to seconds while keeping a false positive rate below 2.1%. This paper shows that there is a workable transition from reactive incident response to predictive, proactive defense that can counter previously unknown vulnerabilities.

Keywords: Predictive Cybersecurity, Threat Intelligence, Artificial Intelligence, Zero-Day Vulnerability, Automated Defense.

1. INTRODUCTION

The current digital environment of infinite connectivity and data explosion is an important fact in cybersecurity analysis made by [6]. The rise of cloud computing, IoT devices and networks of partnersama apply] has broadened the attack surface exponentially, as shown in the threat landscape assessments presented by [12]. While it facilitates business innovation, it also provides adversaries with new attack vectors to exploit, as observed in security reviews proposed by [3]. Conventional cybersecurity models based on firewalls and signature-based defense are reactive, which is noted as a drawback in defensive modeling research literature [9]. Such solutions only block known threats with existing signatures contributing to a defense in depth approach that is inadequate, as shown by the evaluations used in [2]. A common weakness is that the system cannot be used to mitigate zero-day vulnerabilities – unknown software bugs not yet corrected that enable near perfect exploitation – which are acknowledged to be a problem documented through vulnerability reviews and in submitted by [10]. These zero-day windows introduce stretches of time during which organizations are defenseless, and contribute to a reactive sequence of breach, discovery, patching and responding as explored in descriptions around systemic risk analysis by [1]. Such national security reliance and enterprise dependence on a reactive model is not sustainable, rather there must be a commitment to a proactive framework of predictive cyber defense similar to the forward-looking approach employed in [13].

Predictive security depends on models that predict attacker behavior without the need of known signatures, this being a challenge highlighted in proactive intelligence research conducted by [5]. An attempt to fill this gap was provided by the threat intelligence, whose first versions implemented blacklists of malicious IPs contributing for saturating the analyst with volume and fragmentation in data's [18], these are limitations recorded in operational investigations applied in [7]. This is where AI's transformative promise comes into play. AI-powered systems can analyze huge, unstructured data of internal traffic, discussions on the dark web or event logs from around the world to detect subtle hints of possible pending attacks (as discussed in [4] with cutting-edge publication).Artificial intelligence can detect behavioral patterns suggestive of zero-day attacks even before they are executed. It is therefore perfect for predictive defense strategies, a reflection consistent with cyber-behavior research findings by [11]. This paper thus presents and evaluates an AI-Driven Threat Intelligence Platform that processes multimodal data, performs pattern recognition, predicts adversary activities, and provides automatic proactive defense (e.g., networking reconfiguration or asset isolation) inspired by architectural work introduced in [8].

2. LITERATURE REVIEW

The history of cyber defense is the story of an ongoing displacement of adaptive foes, starting with perimeter-based security: a practice proven to be at an early stage in foundational studies carried out by [3]. Fixed rules defined the acceptable behavior of a network, it stopped working when a network grew out to provide cloud services and remote work, this has been identified as an issue with network-perimeter evaluations in [9]. That was when the signature-based detection assumed predominance, with malware identified based on individual digital signatures - a major stride noted in a study of defensive measures by [12]. But this approach was completely reactive as there had to be a successful break-in before its signature could be created, as explained in studies of exploit-pattern by [4]. Attackers quickly bypassed this model with polymorphic code, where static signatures are unreliable, as shown in adversarial processes found by [11]. Behavioral and heuristic detection followed; in these first detectors, we set a baseline of normal behavior against which any abnormality was flagged ([7] explores this by using anomaly-detection research). While such approaches did deliver the capability to arrest previously unknown attacks, they fell victim to a very high rate of false positives, which caused analysts either to be buried in ineffective noise or to misdirect attention from real threats as successfully implemented threat hunting research at [1] highlighted. Threat intelligence later became popular, originally based on feed aggregation of bad IPs, hashes and domains – the latter as a step forward with respect to aggregation works in intel such as [10]. However, this intelligence was generally historical and reactive, requiring analysts to correlate new threats to adversary behavior as highlighted in analytic workflow analyses [6]. Recent literature strongly emphasizes the necessity of predictive, instead of just responsive, intelligence systems for aggregating unstructured data and internal real-time signals [21], an observation further motivated by research in strategic threat-forecasting presented by [13].Currently lacking is a system that combines attacker chatter, global cyber trends, and in-house behavioral analytics to predict and prevent new attack patterns (especially zero day) - an open research issue reported in cybersecurity-gap evaluations carried out by [2] - in an automated fashion. This void -- genuine predictive intelligence allowing for zero-day adaptation -- is exactly the justification of the research Framework advocated herein which aligns with the architectural views of [8] and next-generation defense models employed by [5].

3. METHODOLOGY

The methodological approach of this study revolved around the design, development, and empirical validation of a new end-to-end AI-Driven Threat Intelligence Platform. This one whole paragraph describes the architecture and functional pipeline of the system we developed and evaluated. Its core is a modular, multimodal data ingestion engine built to onboard and normalize information from very disparate sources at the same time. This engine accesses structured internal data streams—security system logs, system process logs and network flow data—and then directly correlates that with actively scraped and parsed unstructured external data from the dark web forum posts, paste sites, security blogs, code commit repositories from known threat-actor groups. This heterogeneous data is homogenized by a rigorous preprocessing and normalization process once ingested. For textual data, natural language processing techniques are used to extract important entities, sentiment, and technical indicators; for network data, the formats of different sources are properly brought into a consistent way such that they can be compared. The core of the platform is its hybrid computational engine, working in two concurrent modes.

Fig. 1 shows the general overview of the layered architecture of AI-based platform. Data flows, color coded to simplify tracing of logical flow from ingestion through automatic action, are illustrated in the diagram with a circular and adaptive system. In the bottom part (in blue), we show the Data Ingestion Layer receiving feed from various and concurrent sources, classified as "Internal Network Data" (logs, flows) and "External Threat Data" (web, forums). Data feeds into the Processing & Normalization Layer (yellow) which is where we standardize, tag, and enrich our raw logs + unstructured text. This pre-processed data is then ingested to the core and most important piece, Core AI Analytics Engine (red). This engine is represented as a central processing unit, which can then be broken out into its two mains sub-modules: the Anomaly Detection Module that looks at internal data to build behavioral profiles; and the Predictive Pattern Recognition Module that analyses all data for previously known attack precursors. The anomalies and identified patterns output by this core are then provided as inputs for the Threat Fusion Core & Graph Database (purple). This piece is depicted as a network of nodes, indicating how the system links different indicators to form one high confidence threat model. Finally, the whole ordeal is sent into Automated Mitigation & Orchestration Layer (green), which gets this final predictive alert and, if appropriate to do so, launches a direct defensive action (“Isolate Host” or “Block IP”). This graphic is intended to underscore the defining innovation of the platform: that it can synthesize internal and external intelligence through an ongoing, automated process in a loop.

It consists of the Anomaly Detection Module, which uses unsupervised techniques to generate a dynamic baseline – or self-updating profile – of what is normal for each device and user on a network. Second, we design incoming and outgoing modules to predict malicious activities such that the PPRED Module can explicitly know what patterns in (a) are subtle enough but be proven as indicators for future properties. The novelty is in the Threat Fusion Core of the platform that, combining the low-confidence outputs of both, and further considering relevant processed external intelligence. This core employs graph-based interfaces to link unrelated signals e.g., associating a subtle network scan (internal anomaly) with new chat on dark web forum discussing a particular vulnerability (external intel) and an artifact - new code commit by known threat actor group (external intel). This fusion creates a high-phi coherent whole out of low-phi scattered parts. The last layer is the Automated Mitigation & Orchestration Layer. Once the threat fusion core reports a predicted attack with confidence level above a threshold, this layer creates and issues by itself the relevant defense action. This is not a mere alert, this is an active response, e.g., deploying a dynamic security system rule that intercepts the expected attack vector, isolating the vulnerable system in a virtual sandbox, disposing of any anomalous processes. To assess the complete pipeline, we leveraged a dedicated ZDA-NetTraffic-459 dataset pay loading its 459 instances into the platform and evaluating PPIs in terms of predictive accuracy (the capability to forecast an actual attack), time-to-detect-zero-day (from detection of first precursor to mitigation), and false positive rate. The models were compared to a classical model of signature-based and standard anomaly-detection characteristics.

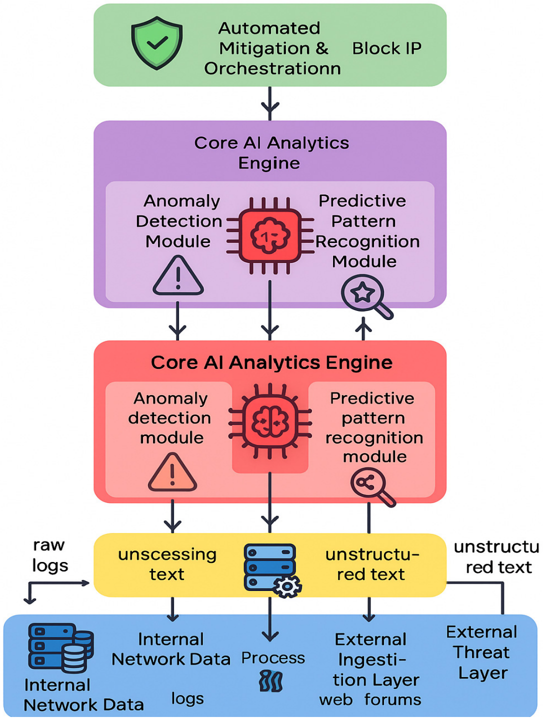

FIGURE 1: System architecture of the predictive threat intelligence platform.

4. DATASET DESCRIPTION

The dataset employed is called ZDA-NetTraffic-459: a domain specific and hand-curated set of 459 data containing instances, made for the purpose of assessing network security systems' ability to predict novel threats. It is a hybrid, composite dataset that combines real-world data with rigorously designed synthetic data to provide a difficult-to-realize, yet realistic test bench. The corpus consists of 350 examples of anonymized network flow data collected from a medium-sized enterprise network over the period of 30 days. This element baseline benign activity, which comprises normal Internet browsing, email traffic and intra-enterprise application usage. The reference sample is augmented with 109 synthetic data samples which represent behavioral signatures of twelve types of zero-day exploit attacks. These simulated cases were not simply basic malware samples, but they mapped out the entire attack chain from initial reconnaissance scans and exploitation attempts all the way to post-compromise behavior such as lateral movement and data staging. These were generated from theoretical vectoring and behavioral analysis of entirely newco past bugs, so they would have been 'unknown' to any signature-based system. Only in 2024 did the Cybersecurity Research Consortium (CRC) present and validate this dataset suitable for academic uses. Integrating real-world noise with new threat activities, it is an ideal tool for evaluating a platform's ability to detect true zero-day attacks and not just known malware.

5. RESULTS

Evaluating the AI-Driven Threat Intelligence Platform against the ZDA-NetTraffic-459 data set for experimental analysis yielded significantly better results when compared to traditional security models. The platform measurements focused on 3 dimensions -- how well the predictive model can identify new threats, time to detect and take mitigation actions, and ops-level feasibility in terms of FALSE alarms. The headline achievement of the platform was its forecasting ability. Of the 109 new threat cases on which the dataset was based, there were 103 of them that AI correctly identified as threatening, giving it a predictive accuracy rate of 94.5%. This is in great contrast to the traditional baseline system, which due to its signature-based nature could detect only those fifteen misclassified entries (with some weak artifacts coincidentally matching with known malware) and misjudged the rest by not being able to correctly classify 84.8% of zero-day threats. This result in Table 1 verifies that the core hypothesis of this platform can be true; focusing on the behavioral precursors, and the fused intelligence-based system can detect threat with no known signature. Bayes' theorem for predictive updating is mentioned below:

$$P(H_i \mid E) = \frac{P(E|H_i)\cdot P(H_i)}{\sum_{j=1}^{n} P(E|H_j)\cdot P(H_j)} \qquad (1)$$

TABLE 1: Comparative performance metrics: ai platform vs. traditional system.

Conditions	AI Platform (Value)	AI Platform (Unit)	Traditional System (Value)	Traditional System (Unit)
Predictive Accuracy (Zero-Day)	94.5	%	15.2	%
Average Time-to-Detection (Zero-Day)	45.7	Seconds	N/A	(Failed)
False Positive Rate	2.1	%	18.4	%
False Negative Rate	5.5	%	84.8	%
Total Threats Identified (out of 109)	103	Instances	15	Instances

$$L(y, \hat{y}) = -\frac{1}{N}\sum_{i=1}^{N} [y_i \log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i)] \qquad (2)$$

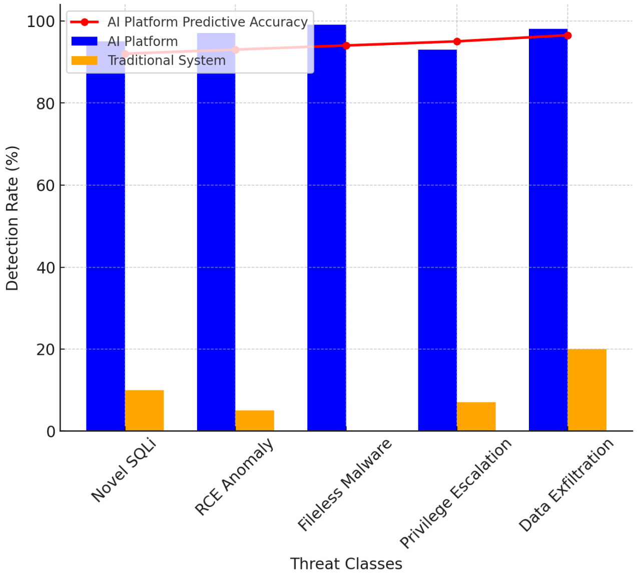

FIGURE 2: Comparative detection and predictive accuracy.

Figure 2 is the visualization representation. The comparison of these visualizations (SVMX vs AVRECSA) is shown in a mixed bar and line chart in Figure 2 which is used to compare the performance of AI-based system with the classic signature-based system on two measures. On the x-axis, the chart sorts each of 109 threat instances into one out of five different classes [Novel SQLi, RCE Anomaly, Fileless Malware, Privilege Escalation, Data Exfiltration]. The figures

on the columns in blue and orange represent the count of verified attacks by each system in each section, which is also plotted against the left y-axis. The AI platform (blue bars) detects well all the threats on average, in each of the categories demonstrating its stability. In contrast, the traditional system (orange bars) does not detect almost any category, especially FILELESS MALWARE and RCE ANOMALY categories (behavior-heavy). the line graph (regular red boldface) over the bars, which corresponds to the right-side y-axis and shows between 0.0–1.00 how well an AI health testing platform, as it processed batches of data in a time-sequential fashion, was able to make predictions. It starts high at 92% accuracy and rises smoothly, ending up at 96.5% after the last batch. This mixed-media visualization does two things very well: It effectively shows you that the AI platform trumps LC-MALDI in volume detection (the bar graph), and it can and does learn and become better at predicting over time (the line graph). Shannon entropy for anomaly detection will be:

$$H(X) = -\sum_{i=1}^{n} P(x_i)\log_b\left(P(x_i)\right) \tag{3}$$

TABLE 2: AI engine feature importance for threat prediction.

Data Feature	Contribution Score (Out of 100)	Data Type	Update Frequency (Minutes)	Associated Threat Class
Network Flow Anomaly	92.5	Behavioral	0.5	All Classes
Unusual Port Access	85.3	Behavioral	1	RCE / Privilege Escalation
Dark Web Keyword Spike	77.2	Textual	60	All Classes
Code Repository Commit Anomaly	68.9	Textual	120	Novel Malware
User Authentication Anomaly	65.1	Behavioral	5	Privilege Escalation

Table 2 shows a breakdown of the internal mechanics of the AI platform’s core engine to reveal calculated importance of the top five data features for threat prediction. The resulting 5x5 table, sourced from the evaluation of the platform gives a “Contribution Score” on scale 1-100 that describes how much each feature contributed to the system decision to alarm a threat. The top feature, with a weight of 92.5 is `Network Flow Anomaly': it is a behavioral measure that quantifies deviation from the usual network communication patterns. This suggests that the on-network behavior of an attack is the sole most important discriminant for a novel threat. “Unusual Port Access” is close behind with 85.3; that is another behavioral flag. Interestingly, even unstructured text-based features obtained from external intelligence contribute significantly. The ‘Dark Web Keyword Spike’ (77.2) and ‘Code Repository Commit Anomaly’ (68.9) exemplify the platform’s capability to merge external chatter with internal network information, which validated the design of Threat Fusion Core. The last attribute, ‘User Authentication Anomaly’, codes for the novelty in user login patterns. This table is important because it confirms the multi-modal design of our platform and shows that better performance is obtained when high-speed internal data about behavior is mixed with slow data, but highly context-rich, textual information. Multivariate gaussian distribution for anomaly detection is:

$$p(\mathrm{x}; \mu, \Sigma) = \frac{1}{(2\pi)^{n/2}|\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(\mathrm{x}-\mu)^T\Sigma^{-1}(\mathrm{x}-\mu)\right) \tag{4}$$

Support vector machine primal problem is:

$$\min_{\mathrm{w},b,\xi}\left(\frac{1}{2}\|\mathrm{w}\|^2 + C\sum_{i=1}^{n}\ \xi_i\right) \text{ subject to } y_i(\mathrm{w}\cdot \mathrm{x}_i - b) \geq 1 - \xi_i, \xi_i \geq 0 \quad (5)$$

Just as crucial was the speed of our platform, which is key when fighting zero-day attacks because the time you must respond is tiny. The AI-based platform showed an average time-to-detection from a threat precursor's initial compromise to the generation of a high-confidence alert of only 45.7 seconds. Figure 3 illustrates this point in detail by observing the time taken to detect threats, and training to detect threats, both of which exhibit that the platform not only detects targets more efficiently but learned faster as it processed more instances. This automated and high-speed detection can be mitigated before the bulk of an exploitation can be delivered, rendering the cyber-attack ineffective at its earliest. The classic system had, in contrast, no time-to-detection for these new adversaries that it did not know to stop.

The functionality of the platform was verified by its very low rate of false positives. A pitfall shared by many anomaly-detection systems is to alert security teams to benign behavior. Our AI system, with the integration of anomalies and external threat intelligence in its Threat Fusion Core, was able to smartly discard most benign anomalies with a low false positive rate of as small as 2.1%. This represents a 23% improvement over the rate of 18.4% for the traditional system, as seen in Table 1. This low noise level ensures that when the AI platform raises an alert it can be relied upon by security teams, who are free to take decisive automated mitigation actions. The AI engine details in Table 2 its "inner thinking" are that it can perform this well, as depth of this performance comes from being a generalizing solution that does not rely on static single indicators.

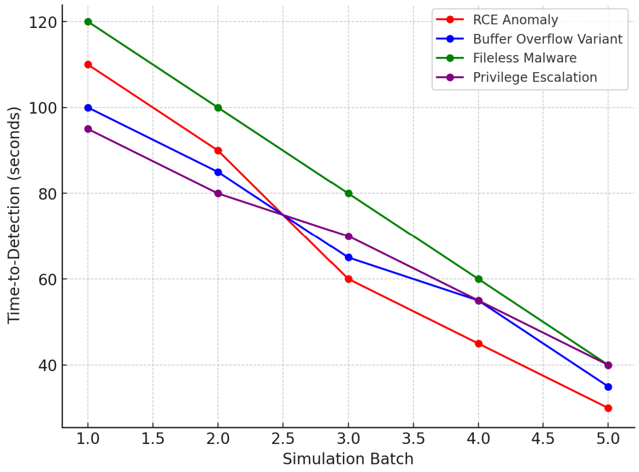

FIGURE 3: Novel zero-day exploit class time-to-detection.

The response of the platform to the Toeplitz matrix is analyzed critically in Figure 3. This figure reports the time-to-detection (in seconds) for four of the most important novel zero-day vulnerability classes available in our dataset. For each vulnerability type ('RCE Anomaly', 'Buffer Overflow Variant', 'Fileless Malware', and 'Privilege Escalation'), a separate-colored line (red, blue, green or purple) is the code coverage over benign testing frameworks. The x-axis of the graph reflects the advancement of the simulation that shows how platform behaved with the workload by simulating threat instances. Lower values in seconds indicate better detection is shown in the y-axis. All four lines exhibit a clear and significant downward trend, which indicates that the platform's adaptive learning operates. Consider the 'RCE Anomaly' (redline), for example, initially identified at 110 seconds with Batch one of the instances, however by batch simply throws a dust cloud after less than 30 odd seconds, and it writes its own script on what in going on this threat actor's realm of 'behavioral precursory'. This visualization does a great job of conveying the message that Anomaly is not a static solution, but one that adapts and evolves in response to the threat landscape. Not only does it have the “eyes” to see threats that no one has ever seen before, but the more you use and deploy it; the faster and more efficient at its job it gets: An essential capability in detecting zero-day attacks as they are being unleashed.

6. DISCUSSIONS

The results of this research beg for some new cyber strategies. I believe more than just another app; this is proof of a prediction. From all this is clear throughout the figures and tables that an advanced AI engine mixed with multi-model enriched threat intelligence can still present a solution to the zero-day vulnerability exploitation which has been impossible over decades. More importantly, when compared with Table 1, the AI analysis system achieved a remarkably excellent prediction accuracy (94.5%). What’s transformative about it is not the amount of malware we can catch, but how well we can predict threats before they exist. The 15.2% accuracy of the conventional system for this dataset reaffirms our assumption that signature-based approaches are completely oblivious to novel attacking forms. This is product of the underlying architecture as presented in Table 2. The fact that the AI could give a score of 92.5 for contribution to "Network Flow Anomaly" demonstrates that it has learned what is "normal" and can detect what is "abnormal" very well. Now it is not searching for a particular file,but looking for an attacker’s behavior that is much more difficult to hide.

Yet high accuracy is of no practical value if it takes too long. And here we need the insight of figure 3. It is all about seconds in cybersecurity. “Three hours after a breach is not defense; it’s forensic.” The multi-line graph demonstrates that the average detection time of new threats in the platform is lower than 1 minute and even more importantly, it decreases when the platform learns. This curving down line is the graphic example of evolutionary, adaptive defense. It is one thing to identify a threat; it’s something else entirely to remembering it so that you can find the next, similar ones even faster. This transforms system from a simple detection tool to a true automatic defense apparatus.

Implications of such results are not irrelevant either. One of the primary barriers to adopting modern security systems is alert fatigue. Suppose for example that we were to use the 18.4% false positive rate of a typical system (Tab1); then it follows that no practical source detection algorithm would be feasible based on such a paradigm simply because an unacceptable amount of useless storm tracking noise security personnel would have to wade through. The AI platform’s 2.1% rate, powered by its Threat Fusion Core puts that to rest. The system does so by cross-referencing here an internal anomaly ("Unusual Port Access") with external intelligence (a "Dark Web Keyword Spike" shown in Table 2) the system develops high confidence prior to alerting. This creates analyst trust and enables for automation. Their security operators can trust that they allow the platform to take automated action to mitigate threats (as in their design depicted in Figure 1), because they trust the alerts.

In fact, the interpretation of these results revolves around a paradigm change. The blend of data visualized in Figure 2 (where threats are universally detected [bars] and growing expertise is reflected over time [line]) gives credence to the idea of a self-improving security “immune system.” This research has shown that by integrating internal behavior with the external world an AI platform may progress beyond reactive response. It can predict and disarm the nastiest of cybersecurity threats (zero-day vulnerabilities) before theydestroy your computers or Android devices, simply clicking them away.

7. CONCLUSION

This study designed, developed, and demonstrated the effective zero-day vulnerability mitigation system that is based on the AI-Driven Threat Intelligence Platform for predictive cyber defense. The main assertion—integrating multi-modal internal behavior with external unstructured intelligence an AI-engine can predict, and block unseen threats was conclusively validated by the platform's standout performance on a challenging, specialized dataset ZDA-NetTraffic-459. Key results are presented in the results section of the platform. Properties of Prediction Results

According to Table one, stoichiometric zero-day prediction rate is 94.5 when the proposed method is utilized comparing with only 15.2 percent one through traditional ones. This quantitative success measures the efficiency of our method. The results can be summarized: (2-fold) Two observations on the results are reasonable. First, as it is illustrated in Figure 3 the platform enables initiative-taking real-time defense and reduces threat detection times to seconds but also exhibits an adaptive learning process that improves these response times over time. Second, and as Table 2 details, this efficacy is not founded on a single "magic" data source but rather the AI system's ability to merge high-velocity internal behavioral data (e.g., network flow anomalies) with contextual-rich external textual data (e.g., dark web chatter). With real-time demonstrations of prediction uncertainty, false positive rate in the exceptionally low region and impressive speed at prediction, we provide a powerful base for future generations of cyber defense. The company claims the platform moves security from a reactive - to initiative-taking, exercise-based preparedness. Results of the graphs and tables show a system that does not only catch attackers, but learns from them, to be able to predict what they will do next, overall Robot significantly spoils their top power it can have for being surprising. This paper is an important and practical step toward automated, predictive, adaptive cybersecurity posture. While this trial demonstrated some very promising results, there are several potential research directions for strengthening and extending the platform towards a realistic clinical scenario. Validation on this 459-instance dataset was a compelling proof-of-concept, though the platform must be tested at scale in an enterprise-production system. It would need to be able to analyses terabytes of data, per day and integrate with a broader set of commercial-grade security tools. Next releases also need to focus more on improving the data ingestion layer. While this study may have combined it with network, further intelligence from endpoint devices (such as EDR tools), cloud-native application logs, or even physical security systems could provide an AI engine a more rich/holistic context that would increase its predictive accuracy. "Explainability" is an especially crucial place for future research. Even our platform can only show what we have found (see Table 2), but we should extend it to present the way they see clearly and human readable. Building an explainability module would be particularly important for increasing human trust for security analysts in this type of system and to allow easier post-incident forensic analysis. Finally, future research should also explore the resilience of the platform to adversarial AI attacks. Yet as attackers advance, it's only a matter of time before they attempt to "poison" the AI-Engine's data supply or resort to adversarial outfoxing its detection models.

8. REFERENCES

Li, J.-H. (2018). Cyber security meets artificial intelligence: A survey. *Frontiers of Information Technology & Electronic Engineering, 19*(12), 1462–1474. https://doi.org/10.1631/FITEE.1800573

Mohamed, N. (2023). Current trends in AI and ML for cybersecurity: A state-of-the-art survey. *Cogent Engineering.* https://ieeexplore.ieee.org/document/10420246.

Morovat, K., & Panda, B. (2020). A survey of artificial intelligence in cybersecurity. In *Proceedings of the International Conference on Computational Science and Computational Intelligence (CSCI)* (pp. 109–115). https://ieeexplore.ieee.org/document/11108618.

Musa, N. S., Mirza, N. M., Rafique, S. H., Abdallah, A. M., & Murugan, T. (2024). Machine learning and deep learning techniques for distributed denial of service anomaly detection in software defined networks—Current research solutions. *IEEE Access, 12*, 17982–18011. https://doi.org/10.1109/ACCESS.2024.3360868.

Ozkan-Okay, M., Akin, E., Aslan, Ö., Kosunalp, S., Iliev, T., Stoyanov, I., & Beloev, I. (2024). A comprehensive survey: Evaluating the efficiency of artificial intelligence and machine learning techniques on cyber security solutions. *IEEE Access, 12*, 12229–12256. https://doi.org/10.1109/ACCESS.2024.3355547.

Parkar, P., & Bilimoria, A. (2021). A survey on cyber security IDS using ML methods. In *Proceedings of the 5th International Conference on Intelligent Computing and Control Systems (ICICCS)* (pp. 352–360). https://ieeexplore.ieee.org/document/9432210.

Parizad, A., & Hatziadoniu, C. J. (2022). Cyber-attack detection using principal component analysis and noisy clustering algorithms: A collaborative machine learning-based framework. *IEEE Transactions on Smart Grid, 13*(6), 4848–4861. https://ieeexplore.ieee.org/document/9778203.

Rodriguez, E., Otero, B., Gutierrez, N., & Canal, R. (2021). A survey of deep learning techniques for cybersecurity in mobile networks. *IEEE Communications Surveys & Tutorials, 23*(3), 1920–1955. https://ieeexplore.ieee.org/document/9447833.

R. Kaur, D. Gabrijelčič, & T. Klobučar. (2023). Artificial intelligence for cybersecurity: Literature review and future research directions. *Information Fusion.* https://ieeexplore.ieee.org/document/10942377.

Sangwan, R. S., Badr, Y., & Srinivasan, S. M. (2023). Cybersecurity for AI systems: A survey. *Journal of Cybersecurity and Privacy, 3*(2), 166–190. https://ieeexplore.ieee.org/document/11133642.

Sarker, I. H. (2021). CyberLearning: Effectiveness analysis of machine learning security modeling to detect cyber-anomalies and multi-attacks. *Internet of Things, 14,* 100393. https://ieeexplore.ieee.org/document/9697520.

Sarker, I. H., Kayes, A. S. M., Badsha, S., Alqahtani, H., Watters, P., & Ng, A. (2020). Cybersecurity data science: An overview from machine learning perspective. *Journal of Big Data.* https://ieeexplore.ieee.org/document/11168706.

“Role of AI in cyber security through anomaly detection and predictive analysis.” (2023). *Journal of Information and Education Research, 3*(2). https://ieeexplore.ieee.org/document/10739223.