# Early Detection of Classifier Health Deterioration Under Zero-Day Cyberattacks

Chhaya Katiyar[1], Priscila Silva[2]

Electrical & Computer Engineering, University of Texas at El Paso

[1]ckatiyar@utep.edu, [2]pdsilva@utep.edu

*Abstract*—**Traditional performance metrics such as the F1-score evaluate whether a cyberattack classifier is making correct decisions, but they provide limited insight into how confidently those decisions are being made as threat environments evolve. Consequently, deployed classifiers may experience increasing uncertainty and health deterioration long before substantial performance degradation becomes apparent. To address this challenge, this paper proposes a Prognostics and Health Management (PHM)-oriented framework for online monitoring of cyberattack classifier degradation under zero-day attack conditions. A binary Support Vector Machine (SVM) classifier is trained to distinguish benign from malicious traffic using the CIC-IDS2017 dataset, while previously unseen attack classes are progressively introduced during testing. The proposed approach converts the SVM decision margin of each incoming sample into a risk measure that is recursively accumulated through a SVM Health Index ($SHI$). Experimental results show that $SHI$ provides a stable and interpretable representation of classifier health, outperforming Cluster Mean Distance, Kolmogorov-Smirnov, and Page-Hinkley monitoring approaches. Moreover, $SHI$ identifies emerging degradation trends despite the classifier maintaining a Macro F1-score above 97%, demonstrating its potential as a leading indicator of degradation for proactive maintenance.**

*Index Terms*— **Cyberattack Classification, Prognostics and Health Management, Online Health Monitoring, Zero-Day Attacks**

## I. Introduction

Deep learning models are widely deployed as classifiers for anomaly detection [1] and zero-day cyberattack recognition [2]. Although considerable effort has been devoted to improving detection performance [3], the ability to assess classifier health after deployment has become increasingly important as previously unseen attack patterns emerge over time [4]. Under such conditions, a classifier may continue producing correct predictions while simultaneously experiencing increasing uncertainty. As a result, conventional performance metrics alone may provide an incomplete view of the operational condition of classifiers, delaying maintenance and reducing the long-term reliability of defense systems.

Support Vector Machines (SVMs) [5] have been extensively adopted for Network Intrusion Detection Systems (NIDS) [6] because of their strong generalization capability under limited training data [7] and suitability for real-time deployment [8]. To further improve detection performance and zero-day attack recognition capabilities [9,10], SVMs have been integrated with particle swarm optimization [11], ant colony optimization [12], hyper-heuristic frameworks [13], artificial neural networks [14], autoencoders [15], deep feature extraction [16], and adaptive kernel methods [17]. Despite these advances, the problem of monitoring the operational health of deployed SVM classifiers under evolving cyber threats remains largely unexplored. To assess classifier behavior when there are distributional differences between known and unknown data, researchers investigated Page-Hinkley [18], cluster mean distance [19], Kolmogorov–Smirnov [20], and repair score [21] methods. However, they mostly rely on offline analysis or window-level tests and therefore, uncertainty associated with emerging zero-day threats may remain undetected until degradation becomes evident in traditional metrics.

To address these limitations, this paper proposes a continuously health monitoring of deployed cyberattack classifiers operating under zero-day attacks scenarios to identify degradation trends before they become apparent in conventional performance metrics. Specifically, this work develops a SVM Health Index ($SHI$), which transforms the decision margin associated with each incoming sample into an uncertainty-based risk measure and recursively accumulates this information over time to enable monitoring of long-memory degradation while accounting for fluctuations from dynamic network traffic. The proposed $SHI$ is evaluated against established techniques such as Page-Hinkley ($PH$), Cluster Mean Distance ($CMD$), and Kolmogorov–Smirnov ($KS$) tests, to assess their ability to identify early degradation

in emerging zero-day threats scenarios. Experimental results on a public dataset demonstrate that the proposed $SHI$ better interpret classifier performance degradation even when the classifier maintains a high F1-score, demonstrating the potential of $SHI$ as an early-warning indicator for proactive maintenance decisions.

## II. Support Vector Machine Classifiers

A Support Vector Machine (SVM) [22] performs classification by identifying an optimal decision boundary, referred to as a hyperplane, to maximize the separation margin between different classes. In NIDS [5], SVMs analyze network traffic observations, such as packet payloads, and construct decision boundaries that separate benign traffic from malicious cyberattacks. To improve the separation of complex patterns, SVMs commonly employ kernel functions [17], which transform the original input data into higher-dimensional feature spaces where classes become easier to separate.

Classifiers performance is commonly evaluated using the F1-score [23], which accounts for false positives (FP) and false negatives (FN). In cybersecurity, the Macro F1-score [24] is also adopted because it provides a balanced assessment of classification performance across both benign and malicious traffic. Although attack detection remains the primary objective of intrusion detection systems, monitoring performance on both classes provides a more comprehensive view of classifier behavior under evolving operational conditions.

## III. Methodology

Inspired on Prognostics and Health Management (PHM) [25,26], which assess system health to support timely maintenance decisions, this section presents the proposed framework for online health monitoring of Support Vector Machine (SVM) cyberattack classifiers to identify deterioration during classifier operation. By continuously evaluating the uncertainty associated with incoming samples, the proposed approach provides early indications of evolving threats and changes in the operational environment before substantial degradation becomes evident in conventional performance metrics. Figure 1 illustrates the proposed framework.

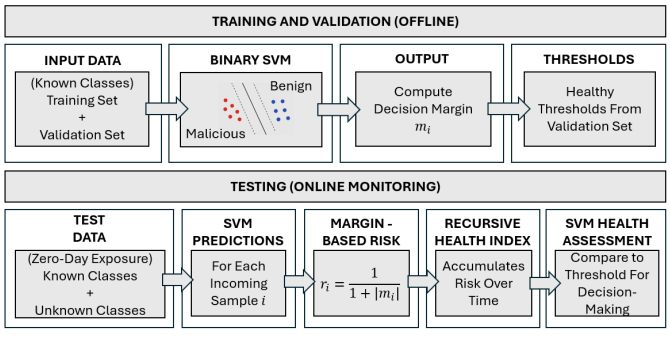


Fig. 1. Proposed framework

Initially, a balanced dataset composed of benign traffic and known cyberattack samples is divided into training and validation subsets. During training, a binary SVM classifier learns a decision boundary that maximizes the separation margin between the benign and malicious classes, while model hyperparameters are adjusted during validation until satisfactory classification performance is achieved in terms of the Macro F1-score. To test the binary SVM, the classifier is exposed to conditions involving both previously unseen samples from the known classes and zero-day cyberattacks that were not observed during training, where the proportion of zero-day attacks are gradually increased. The objective is to assess how classifier health evolves as exposure to previously unseen threats increases and to determine whether the proposed monitoring framework can identify growing uncertainty and deterioration trends.

### A. Recursive SVM Health Index

For each incoming sample, the SVM computes a decision margin that quantifies the distance between the sample and the separating hyperplane. Samples located far from the decision boundary are classified with high confidence, whereas samples near the hyperplane indicate increased uncertainty and operational risk. To quantify this uncertainty, the decision margin is transformed into an instantaneous uncertainty risk ($r_i$) as

$$r_i = \frac{1}{1 + |m_i|} \tag{1}$$

where ($m_i$) denotes the signed decision margin associated with sample ($i$). The absolute value ensures that uncertainty depends only on the distance to the decision boundary, regardless of whether the sample is classified as benign or malicious. This transformation maps the margin into a bounded interval ($0 < r_i \leq 1$), where risk values close to zero correspond to highly confident classifications located far from the hyperplane, while values approaching one indicate samples positioned near the decision boundary and therefore associated with greater classification uncertainty.

Although the decision-margin risk can be considered instantaneous classification uncertainty, individual high-risk samples may arise from transient traffic variations or atypical observations that do not necessarily indicate persistent classifier degradation. Therefore, rather than analyzing individual samples alone, the proposed metric recursively accumulates information on uncertainty over time through a SVM Health Index ($SHI$) defined as

$$SHI_i = \lambda SHI_{i-1} + (1 - \lambda) r_i \tag{2}$$

where $\lambda \in [0, 1]$ is a weighting parameter that determines how strongly historical operating conditions influence the current health assessment. The component $\lambda SHI_{i-1}$ preserves information accumulated from past observations, whereas $(1 - \lambda) r_i$ incorporates the uncertainty associated with the newly observed sample. As a result, the proposed SHI captures sustained changes in classifier

behavior while attenuating the effects of short-term variability and isolated anomalous observations. The recursive structure is motivated by Exponentially Weighted Moving Average (EWMA) monitoring methods [27] widely used in PHM applications. Unlike conventional EWMA formulations that smooth raw measurements directly, the proposed $SHI$ recursively aggregates uncertainty information derived from the SVM decision margin, producing a continuous health measure capable of tracking gradual degradation associated with evolving zero-day cyber threats. As exposure to previously unseen attacks increases, samples tend to migrate closer to the decision boundary, resulting in elevated uncertainty and a corresponding increase in $SHI$, thereby providing an early indication that retraining or alternative defensive actions may be required before substantial performance degradation becomes evident in conventional metrics such as the Macro F1-score.

### B. SVM Health Index Critical Threshold

A healthy operational threshold ($SHI_{crit}$) is defined as a high percentile of the healthy $SHI$ distribution

$$SHI_{crit} = Q_p\left(SHI_{\text{healthy}}\right) \tag{3}$$

where $Q_p(\cdot)$ represents the $p$-th percentile operator (e.g., $95^{th}$ or $99^{th}$ percentile), and $SHI_{\text{healthy}}$ corresponds to the set of $SHI$ values obtained from validation samples belonging exclusively to the known classes used during training, representing the nominal operating conditions where no zero-day attacks are present. This threshold represents the upper operational limit of expected classifier uncertainty during healthy behavior. By estimating $SHI_{crit}$ directly from healthy operational data, the proposed framework adapts the degradation threshold to the intrinsic uncertainty characteristics of the deployed classifier. Persistent exceedance of $SHI_{crit}$ indicates sustained degradation behavior.

### C. Existing Monitoring Methods

To evaluate the effectiveness of the proposed $SHI$, comparisons are performed against three established monitoring and distribution-shift techniques commonly used to assess changes in classifier behavior under evolving operational conditions:

*1) Page-Hinkley ($PH$) Test [18]:* monitors cumulative deviations between incoming margin-based risk observations and their running mean to detect persistent changes in classifier behavior over time. An alarm is generated whenever the $PH$ test exceeds a predefined threshold, indicating a sustained shift in the uncertainty characteristics of the incoming traffic.

*2) Cluster Mean Distance ($CMD$) [19]:* computes the average Euclidean distance between samples within a testing window and the nearest centroid derived from the benign and malicious training classes. Larger ($CMD$) values indicate that incoming traffic increasingly differs from the operational patterns observed during training, suggesting potential exposure to previously unseen attack behaviors.

*3) Kolmogorov–Smirnov ($KS$) Statistic [20]:* quantifies the discrepancy between the training distribution and the current testing-window distribution by measuring the maximum distance between their empirical cumulative distribution functions. Larger ($KS$) values indicate stronger distributional deviation from the benign and malicious traffic characteristics learned during training, thereby providing evidence of evolving operational conditions and potential zero-day threats.

## IV. Illustrations

To evaluate the proposed framework, CIC-IDS2017 dataset [28] is used, with 1,410,255 network traffic samples in 15 distinct classes, as summarized in Table I.

TABLE I
NIDS Benchmark CIC-IDS2017 Dataset

| Class | Name | Samples | Class | Name | Samples |
|---|---|---|---|---|---|
| 0 | Benign | 362108 | 8 | DoS Slowhttptest | 80542 |
| 1 | Infiltration | 115007 | 9 | DoS Hulk | 250000 |
| 2 | Bot | 2543 | 10 | DoS GoldenEye | 128122 |
| 3 | PortScan | 830 | 11 | Heartbleed | 13486 |
| 4 | DDoS | 241405 | 12 | Brute Force | 11754 |
| 5 | FTP-Patator | 31843 | 13 | XSS | 3341 |
| 6 | SSH-Patator | 48165 | 14 | Sql Injection | 12 |
| 7 | DoS slowloris | 121097 | | | |

To construct a balanced dataset, 4,000 randomly selected benign samples (Class 0) and 4,000 randomly selected malicious samples from the known attack classes (Classes 1–9) of the CIC-IDS2017 dataset were used for training and validation. The resulting binary dataset was partitioned such that 80% of the samples were allocated for training, while the remaining 20% were reserved for validation and threshold estimation. A binary SVM with a Radial Basis Function (RBF) kernel [17] was applied due to its effectiveness in modeling nonlinear decision boundaries commonly encountered in cybersecurity datasets. The hyperparameters were tuned using the validation dataset until stable classification performance exceeding 90% Macro F1-score was achieved. After training, the SVM decision margin was used as the primary source of uncertainty information, enabling margin-based risk estimation and subsequent online monitoring through the proposed $SHI$.

### A. Illustration I: Sensitivity Analysis of $SHI$

After training and validation, a testing dataset containing 2,000 samples was constructed using reserved benign and known malicious samples that were not used during training or validation, together with malicious samples from unseen attack classes (Classes 10–14) to emulate zero-day attack traffic. Testing samples were introduced sequentially in random order to simulate an online network traffic environment. To evaluate degradation behavior under increasingly challenging operating conditions, the testing sequence was divided into 20 windows of 100 samples each, and the proportion of zero-day attacks was gradually increased from 0% to 100% using a linear progression. Although such a progression may not occur in practice, this controlled scenario was intentionally applied to facilitate the analysis of classifier health as exposure to previously unseen threats

increases over time. To assess the robustness of the proposed framework, the testing procedures were repeated twenty times independently. In each repetition, a new testing dataset was generated using different randomly selected samples and a different ordering of known and unseen traffic. The results presented throughout this paper therefore correspond to the average behavior across the twenty runs.

During each testing run, the $SHI$ was computed for every incoming sample according to Equation (2) using $\lambda = (0.5, 0.75, 0.95)$ to evaluate the influence of recursive memory accumulation on classifier health monitoring. Since no prior health history exists for the first observation, $SHI_1$ was initialized using the margin-based uncertainty risk associated with the first incoming sample. Subsequent samples were then incorporated recursively according to Equation (2). Figure 2 shows the $SHI$ trajectories over the 2,000 testing samples for the different $\lambda$ values assumed, averaged across the twenty independent runs. Confidence intervals were computed from the standard deviation observed across the runs to quantify variability in the monitoring behavior. The figure also illustrates the corresponding healthy threshold ($SHI_{crit}$), computed using the $95^{th}$ percentile of the $SHI_{\text{healthy}}$ from the validation dataset for each $\lambda$ configuration. The $95^{th}$ percentile was selected because percentile-based thresholds are commonly adopted in anomaly detection and statistical process monitoring to represent the upper bound of nominal operational uncertainty while remaining robust to isolated fluctuations and outliers [29]. Values exceeding $SHI_{crit}$ indicate that the classifier is experiencing uncertainty levels beyond those observed during nominal operation, suggesting increased exposure to previously unseen attack patterns and potential operational degradation.

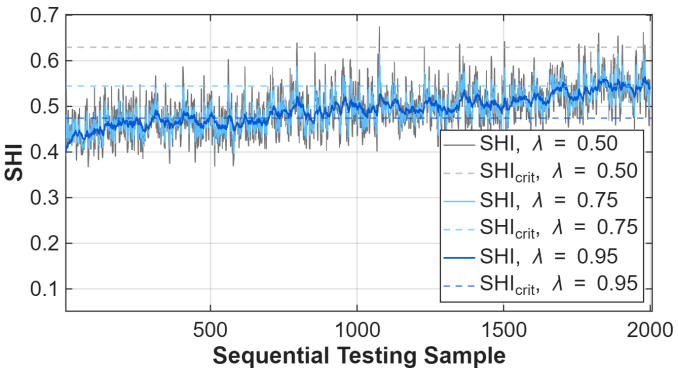


Fig. 2. Sensitivity Analysis of $SHI$ for different $\lambda$

Figure 2 illustrates that smaller memory parameters produce more responsive $SHI$ trajectories, whereas larger values emphasize long-term uncertainty accumulation. Specifically, $\lambda = 0.5$ exhibits greater sensitivity to short-term fluctuations in the incoming traffic, while $\lambda = 0.75$ provides a balance between responsiveness and smoothing. In contrast, $\lambda = 0.95$ generates the smoothest health trajectory by placing greater weight on previously observed uncertainty, thereby reducing the influence of isolated variations. As exposure to unseen attacks increases, $\lambda = 0.95$ reveals the most stable and interpretable degradation trend, making it particularly suitable for long-term health monitoring. Consequently, $\lambda = 0.95$ was selected for the remaining analyses presented in this paper.

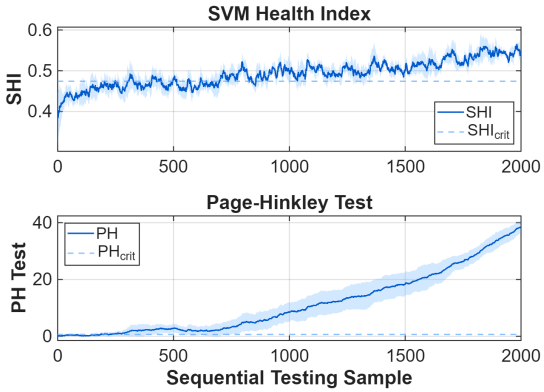


Fig. 3. Online monitoring comparison: $SHI$ vs $PH$

### B. Illustration II: Comparison of SHI and Existing Metrics

The effectiveness of the proposed $SHI$ was evaluated by comparing its monitoring behavior against the Page-Hinkley ($PH$) test, a commonly adopted method for detecting persistent changes in sequential data streams. For this comparison, the $PH$ test was updated for each incoming sample using the margin-based uncertainty risk defined in Equation (1). A corresponding critical threshold, $PH_{crit}$, was determined from the validation dataset and represents the $95^{th}$ percentile $PH$ value observed under healthy operating conditions, and it is also shown in Figure 3. As a reminder, the profiles of both methods shown in Figure 3 are the averages for twenty testing runs, and the shaded bands represent the 95% confidence intervals derived from the variability among the independent experiments.

Figure 3 shows that both $SHI$ and $PH$ increase as exposure to previously unseen attacks grows, indicating that both methods are capable of detecting changes in the operational environment. However, the proposed $SHI$ produces smoother trajectories and narrower confidence intervals across the twenty independent testing runs, demonstrating greater consistency and robustness to random variations in the incoming traffic. In contrast, the variability of the $PH$ test increases throughout the testing sequence due to its cumulative change-detection mechanism. An important observation is the behavior of the corresponding healthy thresholds. The $PH$ test exceeds $PH_{crit}$ almost immediately and remains above the threshold for the remainder of the experiment, indicating high sensitivity to even small deviations from nominal operating conditions. While this characteristic is advantageous for rapid change detection, it provides

limited insight into the severity or progression of the degradation. In contrast, the proposed $SHI$ remains below $SHI_{crit}$ during the initial portion of the testing sequence and crosses the threshold only after uncertainty has accumulated sufficiently to indicate sustained deviation from healthy behavior. Furthermore, once $SHI$ exceeds $SHI_{crit}$, it continues to increase gradually, providing a more interpretable representation of degradation progression. Consequently, while $PH$ functions primarily as a sensitive change detector, $SHI$ acts as a continuous health indicator capable of supporting operational assessment and maintenance decision-making under evolving zero-day attack conditions.

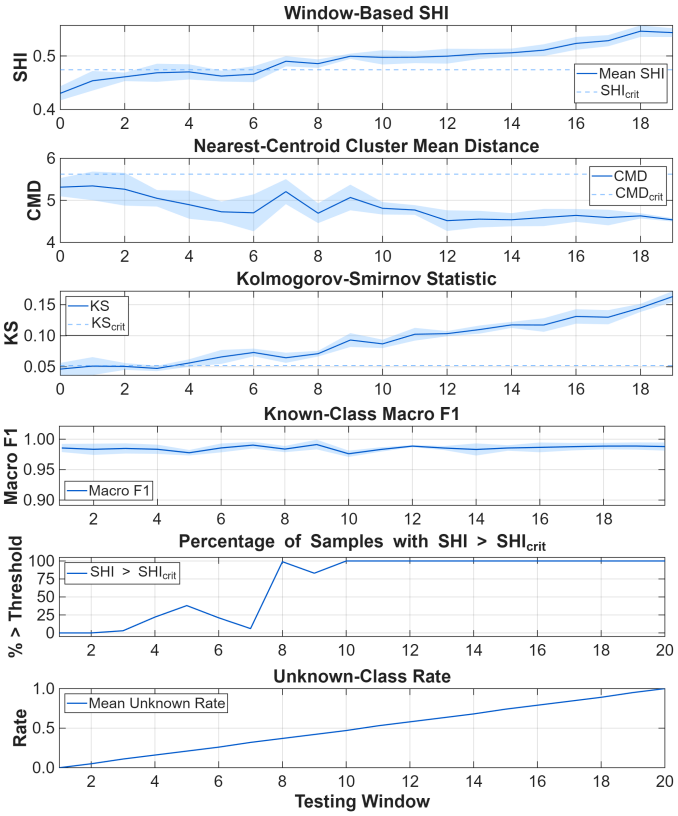


Fig. 4. Window monitoring comparison: $SHI$ vs $CMD$ vs $KS$

Because $CMD$ and $KS$ are inherently window-based metrics, their behavior cannot be evaluated at the individual sample level. Therefore, the testing sequence was partitioned into 20 non-overlapping windows containing 100 samples each. Within each window, the mean value of the proposed $SHI$ was computed and compared against the corresponding $CMD$ and $KS$ statistics. Reference thresholds $SHI_{crit}$, $CMD_{crit}$, and $KS_{crit}$ were derived from the validation, representing their $95^{th}$ percentile under healthy operating conditions. Figure 4 summarizes the evolution of these monitoring indicators as the proportion of unseen attacks increases throughout the testing sequence. For completeness, the figure also reports the corresponding Macro F1-score, the percentage of samples within each testing window for which $SHI > SHI_{crit}$, and the percentage of unseen attacks present in each window. All curves represent the average behavior obtained from twenty independent experiments, with the shaded bands indicating the associated 95% confidence intervals.

Figure 4 shows that $SHI$ increases steadily as the proportion of unseen attacks grows and eventually exceeds its healthy threshold, indicating progressive degradation in classifier health. A similar trend is observed for the $KS$ statistic, which exceeds its threshold earlier, as soon as the testing distribution begins to diverge significantly from the training data. Although $KS$ detects distribution shifts earlier, not every distributional change results in classifier degradation. Because $SHI$ is computed from the SVM decision margin, it directly quantifies the classifier's uncertainty and is therefore better aligned with the health of the classifier than $KS$, which only measures changes in the input data distribution. In contrast, $CMD$ remains relatively stable throughout the experiment and never exceeds its threshold, suggesting that distance-to-centroid measures may be less sensitive when previously unseen attacks remain close to the benign and malicious traffic distributions in the feature space. Moreover, its wider confidence intervals indicates greater variability and reduced consistency across the twenty independent testing runs relative to $SHI$ and $KS$.

An important observation is the behavior of the percentage of samples for which $SHI > SHI_{crit}$. During the initial testing windows, only a small fraction of samples exceed the healthy threshold. However, as the proportion of unseen attacks increases, this percentage rises rapidly and eventually reaches 100%, suggesting that classifier degradation is not driven by a small number of isolated anomalous samples, but rather by a systematic change affecting the overall traffic stream. Interestingly, the Macro F1-score remains nearly constant, with very narrow confidence intervals throughout the experiment, despite the increasing $SHI$, $KS$, and threshold-exceedance percentage. This behavior occurs because many unseen attacks can still be correctly classified as malicious, thereby preserving classification performance, while simultaneously exhibiting reduced separation from the decision boundary. As a result, classifier uncertainty increases even though the final predictions remain correct. These results suggest that the proposed $SHI$ provides a leading indicator of classifier degradation that may not be apparent from conventional performance metrics alone. In other words, the Macro F1-score measures the correctness of the classifier's decisions, whereas $SHI$ characterizes the confidence and health of those decisions under evolving threat conditions. As a result, $SHI$ provides an early indication of emerging degradation, enabling corrective actions before security vulnerabilities arise.

## V. Conclusion and Future Work

This paper presented a framework for online degradation monitoring of binary Support Vector Machine

(SVM) cyberattack classifiers operating under zero-day attack conditions. The proposed approach was inspired on techniques from Prognostics and Health Management, to transforms the SVM decision margin into an uncertainty risk and recursively accumulates this information through a SVM Health Index ($SHI$). Experimental results using the CIC-IDS2017 dataset demonstrated that $SHI$ provides a more interpretable measure of classifier health, allowing a more direct connection to classifier degradation when compared to common monitoring methods from the literature such as Page-Hinkley, cluster mean distance, Kolmogorov-Smirnov tests. Moreover, $SHI$ revealed increasing uncertainty and degradation trends even when the Macro F1-score remained unchanged, highlighting the limitations of relying solely on traditional performance metrics for monitoring deployed classifiers. Therefore, the proposed $SHI$ can serve as a leading indicator of degradation, enabling proactive retraining and maintenance decisions before reliability and security are compromised.

Future research will investigate automated retraining strategies triggered by $SHI$-based degradation alarms, extend the framework to deep learning and multi-class classifiers, and develop predictive models capable of forecasting future $SHI$ trajectories to support proactive maintenance and resilience management.

## REFERENCES


[1] A. A. Eshmawi, A. Al-Nowami, M. Mirza, N. Abu-Raya, R. Al-Thabit, S. Shiaeles, J.-G. Choi, and I. Ashraf, "Machine learning-based network monitoring for cybersecurity threat detection," *Journal of Network and Systems Management*, vol. 34, no. 1, p. 27, 2025.

[2] G. Baye, P. Silva, A. Broggi, N. D. Bastian, L. Fiondella, and G. Kul, "varmax: Towards Confidence-Based Zero-Day attack recognition," in *Machine Learning for Communications and Networking (MILCOM)*, (Washington, USA), p. 6, Oct. 2024.

[3] M. W. A. Ashraf, A. R. Singh, A. Pandian, R. S. Rathore, M. Bajaj, and I. Zaitsev, "A hybrid approach using support vector machine rule-based system: detecting cyber threats in internet of things," *Scientific Reports*, vol. 14, no. 1, p. 27058, 2024.

[4] P. M. Joe Prathap, W. V. Dani, L. S. Beevi, M. J. Kiran, and M. V. Reddy, "Investigation on zero-day attacks: Evolution, myths, complications, avoidance and case study analysis," in *2025 6th International Conference on Data Intelligence and Cognitive Informatics (ICDICI)*, pp. 1061–1066, 2025.

[5] M. Ajdani and H. Ghaffary, "Design network intrusion detection system using support vector machine," *International Journal of Communication Systems*, vol. 34, no. 3, p. e4689, 2021. e4689 IJCS-19-1162.R2.

[6] S. Kudithipudi, N. Narisetty, G. R. Kancherla, and B. Bobba, "Evaluating the efficacy of resampling techniques in addressing class imbalance for network intrusion detection systems using support vector machines," *Ingénierie des Systèmes d'Information*, vol. 28, no. 5, pp. 1229–1236, 2023.

[7] S. SakthiMurugan, A. Sanjay Kumaar, V. Vignesh, and P. Santhi, "Assessment of zero-day vulnerability using machine learning approach," *EAI Endorsed Transactions on Internet of Things*, vol. 10, p. 7, 11 2023.

[8] M. Khule, D. Motwani, and D. Chauhan, "Enhancing network intrusion detection with support vector machines: A comparative study of feature selection techniques," in *2024 5th International Conference on Smart Electronics and Communication (ICOSEC)*, pp. 1281–1286, 2024.

[9] A. K. Dubey, R. Kumar Dubey, A. Shukla, and P. Dubey, "Optimizing cybersecurity: Leveraging support vector machines for real-time threat detection," in *2024 First International Conference on Innovations in Communications, Electrical and Computer Engineering (ICICEC)*, pp. 1–6, 2024.

[10] A. Narayanan, Y. Hazaimeh, K. Muthukumarasamy, S. Ganesan, K. G. Kumar, and J. L. Prasanna, "Enhance cybersecurity monitoring with the use of ml tools: Monitoring enterprise network traffic for anomalies and real-time threat classification," in *2026 International Conference on Data Science, Agents and Artificial Intelligence (ICDSAAI)*, vol. 1, pp. 1–6, 2026.

[11] A.-C. Enache and V. V. Patriciu, "Intrusions detection based on support vector machine optimized with swarm intelligence," in *IEEE 9th IEEE International Symposium on Applied Computational Intelligence and Informatics (SACI)*, pp. 153–158, 2014.

[12] N. Mishra and S. Mishra, "Support vector machine used in network intrusion detection," *IOSR Journal of Engineering (IOSRJEN)*, pp. 25–27, 2018.

[13] N. R. Sabar, X. Yi, and A. Song, "A bi-objective hyper-heuristic support vector machines for big data cyber-security," *IEEE Access*, vol. 6, pp. 10421–10431, 2018.

[14] A. Kajal, "A hybrid approach for cyber security: Improved intrusion detection system using ann-svm," *Indian Journal of Computer Science and Engineering*, vol. 11, no. 4, pp. 412–425, 2020.

[15] H. Hindy, R. Atkinson, C. Tachtatzis, J.-N. Colin, E. Bayne, and X. Bellekens, "Utilising deep learning techniques for effective zero-day attack detection," *Electronics*, vol. 9, no. 10, 2020.

[16] R. El-Sayed, A. El-Ghamry, T. Gaber, and A. E. Hassanien, "Zero-day malware classification using deep features with support vector machines," in *2021 Tenth International Conference on Intelligent Computing and Information Systems (ICICIS)*, pp. 311–317, 2021.

[17] S. Allagi, T. Pawan, L. S. Rodriguez-Baca, P. Bammigati, C. F. C. P. De La Vega, N. Chumuang, and C. Munjandira, "An adaptive cyber threat detection via scaled svm kernels: A data-driven perspective," in *2025 International Conference on Computing Technologies (ICOCT)*, pp. 1–6, 2025.

[18] F. Makhmudov, G. Juraev, O. Yusupov, P. Nasriddinova, and D. Kilichev, "Drift-aware online ensemble learning for real-time cybersecurity in internet of medical things networks," *Machine Learning and Knowledge Extraction*, vol. 8, no. 3, 2026.

[19] C. Cariou, S. Le Moan, and K. Chehdi, "A novel mean-shift algorithm for data clustering," *IEEE Access*, vol. 10, pp. 14575–14585, 2022.

[20] G. Fasano and A. Franceschini, "A multidimensional version of the kolmogorov–smirnov test," *Monthly Notices of the Royal Astronomical Society*, vol. 225, pp. 155–170, 03 1987.

[21] P. Silva, G. Baye, N. Costagliola, N. D. Bastian, G. Kul, and L. Fiondella, "A repair-time trigger for cyberattack classifiers," in *MILCOM 2025 - 2025 IEEE Military Communications Conference (MILCOM)*, pp. 1–6, 2025.

[22] D. S. Kim and J. S. Park, "Network-based intrusion detection with support vector machines," in *Information Networking* (H.-K. Kahng, ed.), (Berlin, Heidelberg), pp. 747–756, Springer Berlin Heidelberg, 2003.

[23] M. Hossin and M. N. Sulaiman, "A review on evaluation metrics for data classification evaluations," *International journal of data mining & knowledge management process*, vol. 5, no. 2, p. 1, 2015.

[24] M. C. Hinojosa Lee, J. Braet, and J. Springael, "Performance metrics for multilabel emotion classification: Comparing micro, macro, and weighted f1-scores," *Applied Sciences*, vol. 14, no. 21, 2024.

[25] N.-H. Kim, D. An, and J.-H. Choi, *Prognostics and Health Management of Engineering Systems: An Introduction*. Springer, 2017.

[26] P. Silva, "Phm-based modeling for cyberattack classifier performance," *Annual Conference of the PHM Society*, vol. 16, no. 1, 2024.

[27] S. W. Roberts, "Control chart tests based on geometric moving averages," *Technometrics*, vol. 42, no. 1, pp. 97–101, 2000.

[28] A. Rosay, E. Cheval, F. Carlier, and P. Leroux, "Network intrusion detection: A comprehensive analysis of cic-ids2017," in *8th International Conference on Information Systems Security and Privacy*, pp. 25–36, SCITEPRESS-Science and Technology Publications, 2022.

[29] V. J. Hodge and J. Austin, "A survey of outlier detection methodologies," *Artificial Intelligence Review*, vol. 22, no. 2, pp. 85–126, 2004.