International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Cyber Threats Detection Based on Artificial Neural Networks Using Event Profiles

Veerabhadra Shanthi, Pacharlla Ravi Srinivasu, Etcherla Vandana, Yabaji Nitin Kumar, Rameswarapu Sreedhar Vignan
Department of Computer Science & Engineering, Institute of Technology, India

Abstract

The rapid proliferation of interconnected systems has elevated the sophistication of cyber threats far beyond what signature-based defenses can address. This paper introduces a novel framework that employs Artificial Neural Networks (ANNs) trained on structured Event Profiles — rich, multi-attribute representations that aggregate raw network events into semantically meaningful units encompassing temporal signatures, entity context, behavioral vectors, and risk metadata. Unlike raw event streams, Event Profiles encode higher-order relationships across correlated activities, providing the ANN with contextually enriched input that dramatically improves discrimination between benign and malicious behavior. Evaluated on the CICIDS-2017 and NSL-KDD benchmarks, the framework achieves a mean detection accuracy of 97.8% with a false-positive rate of 2.6%, outperforming all evaluated baselines including LSTM-based approaches. End-to-end alert latency is reduced to 12 ms through an event-profile micro-service pipeline that integrates natively with enterprise SIEM platforms.

Keywords: Artificial Neural Networks, Cyber Threat Detection, Event Profiles, Intrusion Detection System, Behavioral Analytics, Event Aggregation, Deep Learning, Network Security, SIEM, Anomaly Detection.

1. Introduction

The exponential growth of Internet-connected devices and cloud-native services has dramatically expanded the attack surface accessible to malicious actors. Traditional intrusion detection systems (IDS) that rely on hand-crafted rules or static signatures fail to counter polymorphic malware, zero-day exploits, and Advanced Persistent Threats (APTs) that deliberately evade pattern matching. There is therefore an urgent demand for adaptive, data-driven approaches capable of learning latent threat patterns directly from observable network behavior.

A key limitation of prior machine-learning approaches is their reliance on individual network events — single packets or flows — as the unit of analysis. Individual events lack the broader context needed to identify multi-step attack campaigns that unfold over seconds or minutes. This paper proposes Event Profiles: structured, composite representations that aggregate clusters of temporally and contextually related events into a single enriched object. Each Event Profile captures the temporal signature, entity context (source host, user, process), behavioral action sequence, and risk metadata (MITRE ATT&CK tags, CVE references) of an activity window, providing the ANN with dramatically richer discriminative input.

The proposed framework integrates an Event Profile Builder, an enrichment engine, a four-layer ANN inference module, and a SIEM-compatible alerting layer. We demonstrate that Event Profile-based inputs outperform raw event features on two standard benchmarks, achieving 97.8% mean accuracy and 12 ms alert latency — making the system suitable for real-time enterprise deployment.

2. Literature Survey

Intrusion detection using machine learning has evolved significantly over two decades. Early anomaly detection methods by Denning (1987) applied statistical baselines to audit logs. Subsequent work applied decision trees, SVMs, and naive Bayes to the KDD Cup 1999 dataset, achieving moderate accuracy on well-defined categories but struggling with concept drift and novel attack variants.

Deep learning substantially advanced the state of the art. LSTM networks capture temporal dependencies in sequential log data that feed-forward classifiers miss. Convolutional approaches applied to traffic matrices excel at encrypted

traffic classification. Graph Neural Networks model host communication graphs to expose lateral movement and beaconing invisible at the individual flow level.

Despite this progress, a fundamental gap persists: existing systems treat each network event atomically, losing the causal and contextual relationships between related events. Event correlation engines in SIEM platforms partially address this through rule-based aggregation, but cannot generalize to unseen attack patterns. The Event Profile paradigm bridges this gap by learning to aggregate events into semantically meaningful composite representations, enabling the ANN to reason about attack sequences rather than isolated indicators.

3. Proposed System

The system consists of six tightly coupled subsystems: an event ingestion layer, an Event Profile Builder, a profile enrichment engine, the ANN inference module, a threat classifier, and a SIEM alerting layer.

Figure 1: End-to-End System Architecture

↓	↓	↓	↓	↓	↓

3.1 Event Profile Builder

The Event Profile Builder ingests raw network events from the Kafka consumer group and groups them into profiles using a 60-second sliding window with 10-second slide. Events are co-assigned to the same profile if they share at least two of: source entity, destination entity, protocol family, or attack stage (as inferred by a lightweight pre-classifier). Each resulting profile is a structured object containing five attribute domains: temporal signature, entity context, behavior vector, risk metadata, and aggregation window statistics. On average, 12–40 raw events are fused into a single Event Profile.

Figure 2: Anatomy of an Event Profile

Profile ID	Time window	Source, dest, user, process	Protocol flags	ATT&CK tags

3.2 Profile Enrichment Engine

Profiles are enriched asynchronously with external threat intelligence: MITRE ATT&CK technique mappings are inferred from behavioral vectors using a lightweight trie-based matcher; CVE references are injected when vulnerable service versions are detected in entity context; geolocation and AS reputation scores augment source entity fields. Enrichment adds on average 8 additional feature dimensions, raising the total profile feature vector dimensionality to 86.

3.3 ANN Inference Engine

The ANN accepts 86-dimensional Event Profile feature vectors and passes them through four fully connected layers (86→256→128→64→8) with ReLU activations and batch normalization after layers 1 and 2. Dropout (p=0.3) is applied during training. The output softmax layer produces probability distributions over eight classes: seven attack categories and benign.

Figure 3: ANN Architecture

Profile Features	ReLU + BN	ReLU + BN	ReLU	ReLU	Softmax

7 Attack Classes + Benign					

4. Research Methodology

Figure 4: Event Profile Processing Pipeline

1. Raw Events	2. Profile Builder	3. Profile Enrich	4. ANN Inference	5. Profile Score	6. Threat Alert
→	→	→	→	→	→

4.1 Event Profile Dataset Construction

Training data was constructed by applying the Event Profile Builder to the CICIDS-2017 dataset, converting 2.8 million raw flows into 187,000 Event Profiles. Profiles were labeled by majority vote of constituent event labels. The NSL-KDD dataset was similarly processed to produce 64,000 profiles for cross-dataset validation. Synthetic APT profiles were generated using an adversarial model to address class imbalance for rare attack types, reducing the effective class ratio from 200:1 to 5:1.

4.2 Training Protocol

The ANN was trained for 120 epochs using the Adam optimizer ($lr=1\times10^{-3}$) with cosine annealing decay. Focal loss weighting penalized misclassification of minority attack profiles. A stratified 70/15/15 train/validation/test split preserved class proportions across all profile types. Early stopping with patience of 10 epochs on validation F1-score prevented overfitting. Training completed in approximately 55 minutes on a single NVIDIA A100 GPU.

4.3 Evaluation Metrics

System performance was assessed using per-class precision, recall, F1-score, and AUC-ROC, plus macro-averaged metrics. Profile construction overhead and end-to-end latency (raw event ingestion to alert emission) were measured under 500,000 raw events/second throughput to validate real-time viability.

5. Results and Discussions

Figure 5: Detection Accuracy by Attack Category

DoS/DDoS	94.1%	98.4%	+4.6%
Port Scan	92.5%	97.1%	+5.0%
Brute Force	90.8%	96.2%	+5.9%
SQL Inject	91.9%	97.3%	+5.9%
XSS	89.4%	95.8%	+7.2%
APT/Zero Day	82.0%	91.6%	+11.7%
Ransomware	87.6%	94.7%	+8.1%

5.1 Functional Testing

The system was evaluated against seven distinct attack scenarios replayed from pcap captures at wire speed. Table 1 summarizes test conditions. All attacks were detected within two ANN inference cycles. Crucially, Event Profile-based detection identified multi-step APT sequences that the raw-event baseline missed entirely in 3 of 5 replay instances, demonstrating the value of profile-level context.

Table 1: Functional Test Cases

TC-01	DoS/DDoS Flood	1,840	Yes	10 ms
TC-02	Port Scan	620	Yes	7 ms
TC-03	Brute-Force SSH	410	Yes	9 ms
TC-04	SQL Injection	280	Yes	8 ms
TC-05	XSS Payload	195	Yes	9 ms
TC-06	APT Lateral Mvt	105	Yes	12 ms
TC-07	Ransomware C2	260	Yes	11 ms
TC-08	Benign Traffic	38,000	No	—

5.2 Comparative Performance

Table 2 compares the ANN + Event Profiles system against five baselines on identical test data. The proposed system achieves the highest accuracy (97.8%) and lowest false-positive rate (2.6%) at the fastest latency (12 ms). Gains over the LSTM baseline are statistically significant ($p < 0.01$, Wilcoxon signed-rank test). The Profile Support column highlights that only the proposed system fully exploits the multi-attribute event profile structure.

Table 2: Comparative Analysis

Signature-Based IDS	81.3%	12.4%	~50 ms	No
Rule-Based Systems	78.6%	18.7%	~30 ms	No
SVM Classifier	89.2%	9.1%	~25 ms	Partial
Random Forest	91.5%	7.3%	~20 ms	Partial
LSTM-Based	93.8%	5.8%	~40 ms	Partial
ANN + Event Profiles (Ours)	**97.8%**	**2.6%**	**~12 ms**	**Full**

5.3 Discussion

The Event Profile representation is the primary driver of accuracy improvements. By encoding causal event ordering, temporal signatures, and entity context within each profile, the ANN captures behavioral fingerprints — such as the characteristic SYN-flood burst shape or APT low-and-slow reconnaissance cadence — that are invisible to aggregate per-flow statistics. Profile construction overhead (average 2.1 ms per profile) is offset by a 40× reduction in input volume (187,000 profiles vs. 2.8 million raw events), resulting in net latency improvement despite richer input.

6. Conclusion

This paper presented a framework for cyber threat detection using ANNs trained on Event Profiles — structured, multi-attribute representations that aggregate raw network events into contextually enriched composite objects. By elevating the unit of analysis from individual events to semantically meaningful profiles, the system achieves 97.8% mean accuracy with a 2.6% false-positive rate and 12 ms end-to-end alert latency across seven attack categories. The SIEM-compatible micro-service architecture enables drop-in enterprise deployment without workflow disruption. The results confirm that richer input representations — rather than simply larger or deeper models — are the most effective lever for improving detection performance in cybersecurity contexts. The Event Profile paradigm is model-

agnostic and can be applied to transformer encoders, GNNs, or ensemble methods, making it a broadly applicable contribution to the field.

7. Future Enhancements

Several extensions of the Event Profile framework merit future investigation. First, federated learning across organizations would enable collaborative profile model training without exposing raw network data. Second, dynamic profile window sizing — adapting the aggregation window based on detected attack stage — could further improve recall for slow-moving APT campaigns. Third, SHAP-based profile attribution would expose which profile attributes drove each classification decision, accelerating analyst triage. Fourth, extending profile construction to encrypted TLS traffic using JA3/JA3S fingerprints and certificate metadata would expand coverage to a growing threat surface. Finally, integrating real-time threat intelligence feeds into the profile enrichment pipeline would enable zero-day response without model retraining.

References

[1] Denning, D. E. (1987). An intrusion-detection model. IEEE Transactions on Software Engineering, 13(2), 222–232.

[2] Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. (2009). A detailed analysis of the KDD CUP 99 data set. IEEE CISDA.

[3] Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. ICISSP.

[4] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.

[5] Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). Focal loss for dense object detection. IEEE ICCV.

[6] Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. JAIR, 16, 321–357.

[7] Yin, C., Zhu, Y., Fei, J., & He, X. (2017). A deep learning approach for intrusion detection using recurrent neural networks. IEEE Access, 5, 21954–21961.

[8] Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed representations of words and phrases. NeurIPS.

[9] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.

[10] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. NeurIPS.