



# LLM-Based Optimized Adaptive Threat Monitoring Framework for Malicious Domain and Adversarial URL Detection Process

Banne Phanindhra, Chanumolu Kiran Kumar\*, G Muni Nagamani, Sanda Sri Harsha

*Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation Deemed to be University,  
Green Fields, Vaddeswaram, Andhra Pradesh 522302, India*

**Abstract** The malicious domains and adversarially crafted URLs in cyber threats evolve at a very high speed. The detection frameworks need to be robust, adaptive, and scalable in such scenarios. Traditional detection mechanisms are static feature-based approaches that cannot perform well against unseen threats, adversarial manipulations, and long-term attack evolutions. Existing systems lack granular threat attribution, cross-organization intelligence sharing, and adversarial robustness, making them unsuitable for modern cyber defenses. To address these limitations, we introduce an LLM-Based Frequent Monitoring Framework that combines five advanced techniques: Meta-Learned Self-Supervised Domain Generalization (ML-SSDG), Reinforcement Learning-Augmented Adversarial Training (RL-AdvTrain), Hierarchical Multi-Task Threat Classification (HMT-TC), Temporal Memory-Augmented Transformer for Sequential Threat Detection (TMAT-STD), and Federated Privacy-Preserving Threat Intelligence Learning (FPPTIL). ML-SSDG can achieve zero-shot detection for novel attack domains with a reduction of false negatives by 20% and enhancement of zero-shot accuracy by up to 30%. RL-AdvTrain strengthens the model against masked malicious URLs, detecting 40% more adversarial threats. HMT-TC improves threat attribution and increases classification accuracy for attacks by 50%. TMAT-STD allows for identifying emerging domain threats in real-time, while such detection reduces the response time to domain-based malware campaigns by 30%. The last one is FPPTIL, which allows shared cross-organization threat intelligence without sharing the source of private data. The process of global threat detection improves by 30%. Our framework achieves a holistic, real-time, and privacy-preserving cyber defense solution that adequately outperforms traditional approaches in adversarial resilience, threat attribution, and zero-shot detections. Taken together, these improve the cybersecurity posture, reduce false positives, and support proactive mitigation of emerging cyber threats at scale in process.

**Keywords** Cybersecurity, Malicious Domain Detection, Adversarial URL Detection, Federated Threat Intelligence, Reinforcement Learning, Scenarios

**DOI:** 10.19139/soic-2310-5070-3140

## 1. Introduction

*Malicious Domains and Adversarially Crafted URLs:* Malicious domains and adversarially crafted URLs pose ever-increasing challenges to cybersecurity infrastructure. Cybercriminals continuously evolve new attack techniques, employing domain generation algorithms (DGAs), fast-flux networks, and obfuscation techniques to evade traditional detection mechanisms. Conventional URL classification systems [1, 2, 3] rely on static rule-based or supervised learning approaches, which usually fail in zero-day attacks, adversarial perturbations, and long-term domain behavior analysis. Furthermore, privacy concerns are a barrier to cross-organizational threat intelligence sharing, thereby limiting global detection capabilities. To overcome these difficulties, an advanced, adaptive, and privacy-preserving threat detection framework is very much required to identify malicious domains in real-time and mitigate adversarial URLs in the process. The existing methodologies of detection have been signature-based and

\*Correspondence to: Chanumolu Kiran Kumar (Email: mounikakiran.138@gmail.com).<sup>2</sup>Assistant Professor, Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation Deemed to be University, Green Fields, Vaddeswaram, Andhra Pradesh 522302, India.

Table 1. List of Abbreviations

Abbreviation	Full Form
AI	Artificial Intelligence
AGI	Artificial General Intelligence
APT	Advanced Persistent Threat
CNN	Convolutional Neural Network
C2	Command and Control
CPU	Central Processing Unit
DGA	Domain Generation Algorithm
DNS	Domain Name System
DoS	Denial-of-Service
DSL	Domain-Specific Language
eBPF	Extended Berkeley Packet Filter
FGSM	Fast Gradient Sign Method
GAN	Generative Adversarial Network
GNN	Graph Neural Network
IDS	Intrusion Detection System
IoT	Internet of Things
LLM	Large Language Model
MITRE ATT&CK	MITRE Adversarial Tactics, Techniques, and Common Knowledge
ML	Machine Learning
MQTT	Message Queuing Telemetry Transport
NLP	Natural Language Processing
PGD	Projected Gradient Descent
RL	Reinforcement Learning
TMAT-STD	Temporal Memory-Augmented Transformer for Sequential Threat Detection
FPPTIL	Federated Privacy-Preserving Threat Intelligence Learning
ML-SSDG	Meta-Learned Self-Supervised Domain Generalization
RL-AdvTrain	Reinforcement Learning-Augmented Adversarial Training
HMT-TC	Hierarchical Multi-Task Threat Classification
URL	Uniform Resource Locator
V2I	Vehicle-to-Infrastructure
Zero-Day Attack	An attack exploiting a previously unknown vulnerability

heuristic models, which have become increasingly ineffectual as they rely on predefined patterns and static feature sets. Machine learning-based approaches are more adaptive [4, 5, 6] but fail against adversarially manipulated URLs and fail to generalize over previously unseen attack domains. Furthermore, traditional classification models work on a binary detection paradigm, which restricts their ability to provide granular attack attribution. These systems are further ineffective in detecting the evolving attack patterns, such as long-term botnet infrastructure usage and stealthy domain abuse, due to the lack of time-series threat modeling. Moreover, centralized training models create privacy risks and data-sharing constraints, which prevent a collaborative cybersecurity ecosystem in process.

This paper introduces an LLM-Based Frequent Monitoring Framework to overcome all the aforementioned limitations by exploiting the power of large language models and deep learning-based meta-learning in creating a robust, generalized malicious domain and URL detection system. The proposed framework encompasses five core methodologies: (1) Meta-Learned Self-Supervised Domain Generalization for detecting previously unseen threats, (2) Reinforcement Learning-Augmented Adversarial Training for robustness against obfuscation techniques, (3) Hierarchical Multi-Task Threat Classification for precise attack categorization, (4) Temporal Memory-Augmented Transformer for Sequential Threat Detection for tracking long-term domain activity, and

(5) Federated Privacy-Preserving Threat Intelligence Learning for secure cross-organization intelligence sharing operations. By incorporating these cutting-edge techniques, this framework substantially outperforms traditional methods with respect to detection accuracy, adversarial robustness, and global intelligence collaboration. Empirical results show significant false negatives reduction, improved zero-shot classification, and enhanced detection of adversarially modified URLs. This work is a contribution to the next generation of AI-driven cybersecurity systems and provides a scalable and privacy-aware solution for real-time threat detections.

### ***1.1. Motivation & Contribution***

This calls for a highly advanced threat detection framework that moves beyond the conventional approach to deal with adversarially crafted URLs and fast-evolving malicious domains. The generalization capabilities of conventional detection systems are very poor, vulnerable to adversarial perturbations, and lack attribution to particular attack types. This drives the need for a comprehensive, generalizable, and adversarially robust system that can adapt to emerging threats while ensuring privacy-preserving intelligence sharing. The prime focus of this research is to develop a novel LLM-enabled framework that eliminates the major disadvantages of current alternatives in the form of zero-shot learning, adversarial robustness, hierarchical threat categorization, and federated cybersecurity intelligence. Unlike previous work limited to independent detection mechanisms, this study establishes a higher contribution to the new level of AI techniques used in securing the system from threats. This paper makes significant contributions to the domain of cyber security in two ways: It introduces Meta-Learned Self-Supervised Domain Generalization (ML-SSDG), which significantly improves zero-shot malicious domain detection with a 20% reduction in false negatives. Moreover, it introduces Reinforcement Learning-Augmented Adversarial Training (RL-AdvTrain) to harden the detection against obfuscation techniques to improve adversarial resilience by 40%. Integrating Hierarchical Multi-Task Threat Classification, HMT-TC, the framework includes attack-specific classification that is 50% more accurate. In the time-series domain for threat detection, TMAT-STD lowers the identification of emerging domain-based malware campaigns by 30%. Lastly, with the addition of FPPTIL, the system enables secure decentralized intelligence sharing among organizations with zero risk to privacy and achieves collective threat detection improved by 30%. Together, these contributions form a strong, scalable, and privacy-aware framework for cybersecurity that outperforms traditional detection mechanisms and offers an entirely new standard for AI-driven threat monitoring systems.

## **2. Review of Existing Models used for URL Based Anomaly Analysis**

Such accelerated advancement in cyber threats, along with increasing adversarial sophistication, necessitates raising the quality of advanced detection and mitigation frameworks. Recent research efforts include applying large language models, reinforcement learning, federated intelligence, and adversarial defense strategies to improve the resilience of cybersecurity systems. A comprehensive review of state-of-the-art works offers an in-depth analysis of the present scenario of cybersecurity defenses, their weaknesses, and innovative techniques that overcome existing approaches. This analytical summary covers security risk taxonomies, adversarial attack defenses, digital forensics, deep learning for malware detection, federated threat intelligence sharing, interpretability in cybersecurity models, and LLM-based cybersecurity paradigms, which presents a structured comparative analysis of the methodologies followed by the researchers. Among the vital topics in research related to cybersecurity, risk assessment and categorization of security threats, thus, forms a vital part for proper mitigation strategies. Derner et al. [1] introduces the taxonomy of security risk in LLM-based interactions by identifying the potential vulnerabilities through prompt engineering, data privacy risks, and adversarial exploitation. Santhi and Srinivasan [2] proposed LLM-driven detection of cyber attacks, with more emphasis on education towards cyberphysical along with developing the attack signature based on ChatGPT-based learning platforms. Taxonomies, and models found are with theories as background behind it in this respect of developing research on security by more and more adaptable and self-learning mechanism on the proposed models. Moreover, other than the taxonomies, real time integrity for maintenance on cyber attack detections needs to be carried into systems. Dimitriadis et al. [3] propose

Fronesis, which is an early detection system that is based on the MITRE ATT&CK framework and ontology-driven threat analysis-based, detecting the attacks well in advance. The paper is well aligned with Hu et al. [4], which proposes FastTextDodger, an adversarial attack framework that exploits the vulnerability in the present text-based cybersecurity classifiers through black-box NLP models. Taofeek et al. [5] also discuss cognitive deception techniques in the prevention of data exfiltration during cyberattacks by producing artificial documents for malicious actors. Techniques in this field provide a multi-dimensional mechanism for defense involving digital forensics, adversarial learning, and deception-based countermeasures. Cyber-physical resilience in smart environments is a major trend in cybersecurity research. In his work, Roy et al. [6] discuss resilient transportation networks in smart cities with special emphasis on V2I security protocols for mitigation of cyber threats. Zhan et al. [7] presented a runtime detection framework based on containerized application security against CPU-exhaustion DoS attacks. These works portray the crucial juncture between cybersecurity and real-world infrastructure security and highlight the need for adaptive, context-aware cybersecurity solutions. Iterative, Next, graph-based threat modeling, as seen in table 1, has been very instrumental in cyber threat visualization and predictive analysis. Widł [8] introduces meta attack languages and probabilistic attack graphs and offers support to choose optimal security countermeasures based on graphical security modeling process.

Table 2. Methodological Comparative Review Analysis

Ref	Method	Main Objectives	Findings	Limitations
[1]	Security Risk Taxonomy for LLMs	Categorization of security risks in LLMs, including jailbreak vulnerabilities and prompt injection risks.	Provides a structured taxonomy to assess and mitigate LLM security risks.	Lacks empirical validation in real-world cybersecurity incidents.
[2]	ChatGPT-Based Cyberattack Detection	Development of LLM-driven learning for attack model signatures and defense strategies.	Improves detection efficiency in cyber-physical education applications.	Limited scalability beyond controlled environments.
[3]	Digital Forensics-Based Cyberattack Detection	Early detection of cyberattacks using forensic analysis and ontology-based reasoning.	Enhances attack detection through MITRE ATT&CK integration.	Computational overhead due to ontology processing.
[4]	Adversarial Attack on NLP Models	Exploiting vulnerabilities in black-box NLP models through adversarial text modifications.	Demonstrates high attack success rates, exposing NLP model weaknesses.	Does not propose a countermeasure for adversarial resilience.
[5]	Cognitive Deception for Data Exfiltration Prevention	Use of decoy documents to mislead attackers attempting to extract sensitive data.	Reduces the success rate of data exfiltration in targeted attacks.	Effectiveness depends on the attacker's ability to identify deception.
[6]	Resilient Transportation Security Framework	Cyber-physical security solutions for connected transportation systems.	Provides an integrated resilience model for smart cities.	Does not account for rapidly evolving cyberattack tactics.
[7]	DoS Detection in Containers	Runtime detection of CPU-exhaustion-based DoS attacks.	Improves attack detection in containerized environments.	Requires optimization for large-scale deployment.

Ref	Method	Main Objectives	Findings	Limitations
[8]	Attack Graph-Based Countermeasure Selection	Using probabilistic attack graphs for security countermeasure planning.	Enhances security decision-making with predictive modeling.	Computationally expensive for large-scale networks.
[9]	GAN-Based Intrusion Detection in IoT	Data augmentation using GANs for MQTT-based attack detection.	Improves anomaly detection in IoT networks.	Requires high-quality training data for effective augmentation.
[10]	LLM-Powered Multi-Agent Cybersecurity	Framework for AI-driven multi-agent security infrastructure.	Enhances AI-based automation in cybersecurity monitoring.	Challenges in scalability and inter-agent communication.
[11]	Ensemble Learning for Malicious URL Detection	Combining multiple classifiers for improved URL classification accuracy.	Reduces false positive rates in cyber threat intelligence applications.	Model complexity may hinder real-time implementation.
[12]	AI-Driven Cyber Espionage	Investigating AI applications in state-sponsored cyber espionage.	Highlights emerging AI-based threats in cyber warfare.	Lacks practical defense strategies against AI-enhanced cyber threats.
[13]	Fairness in Machine Learning	Addressing bias and fairness issues in ML-based cybersecurity models.	Improves ethical considerations in cybersecurity AI.	Performance trade-offs in fairness versus detection accuracy.
[14]	Federated Learning for Network Attack Detection	Using GNNs for collaborative threat detection across multiple nodes.	Enhances security without exposing raw data.	High communication overhead in federated environments.
[15]	Multi-Agent Reinforcement Learning for IDS	Distributed intrusion detection using reinforcement learning.	Increases adaptive security responses in large-scale networks.	Requires extensive training for optimal model performance.
[16]	Deep Learning-Based Android Malware Detection	Survey of CNN and Transformer models for Android security.	Identifies effective architectures for mobile malware detection.	Limited effectiveness against highly obfuscated malware samples.
[17]	LLMs in Agent-Based Cybersecurity Simulation	Examining LLM-driven agent-based modeling for attack simulations.	Provides a scalable approach to cybersecurity scenario analysis.	Requires real-world validation for accuracy.

Ref	Method	Main Objectives	Findings	Limitations
[18]	Lifelong Learning AI for Edge Security	Continuous adaptation of AI models for real-time threat detection.	Reduces cybersecurity model degradation over time.	Edge computing constraints may limit performance.
[19]	Cybersecurity in Networking	Adaptations and countermeasures for evolving network-based attacks.	Proposes adaptive security policies for real-time defenses.	Limited applicability in decentralized network architectures.
[20]	AI for Cybersecurity Legislation	Integrating AI-driven solutions for regulatory compliance in cybersecurity.	Improves adaptability of cybersecurity laws to emerging threats.	Regulatory implementation challenges remain.
[21]	Explainability in Deep Learning for Cybersecurity	Reviewing interpretability techniques in AI-based security models.	Enhances trust and reliability in deep learning-based security solutions.	Interpretability often comes at the cost of model complexity.
[22]	LLM Integration in IoT Security	Applications of LLMs in Internet of Things threat detection.	Enables context-aware cybersecurity solutions.	Energy efficiency concerns in resource-constrained IoT environments.
[23]	Digital Forensics and Opposing Agent Models	Examining the evolution of digital forensic techniques against cyber threats.	Identifies weaknesses in traditional forensic methodologies.	Requires AI augmentation for real-time forensic analysis.
[24]	AI Alignment in Cybersecurity Models	Evaluation of AI model alignment techniques for security applications.	Enhances ethical AI applications in cybersecurity.	Requires further exploration into adversarial alignment risks.
[25]	AI in Legislative Optimization	AI-driven approaches to optimizing cybersecurity-related legal frameworks.	Improves automation in legal policy formulation.	Challenges in legal enforcement and adaptability.

This is the approach that Zeghida et al. [9] used in GAN-based data augmentation to enhance intrusion detection in IoT networks, thus improving accuracy in classifying attacks based on communication protocols operating under MQTT and other orientations. It advances probabilistic and graphical modeling techniques applied in applications for cybersecurity, further advancing such approaches towards predictive analytics and risk mitigation. With the advancement of cyber threats, large-scale AI models, and federated learning techniques emerged as the next-generation advanced cybersecurity paradigm. Li et al. in their work [10] had developed a survey for multi-agent systems based on LLM, detailing workflow, infrastructure, and security issues related to deploying multi-agent intelligence in the sphere of cybersecurity. Alsaedi et al. [11] give an ensemble learning-based malicious URL detection model utilizing cyber threat intelligence datasets for achieving real-time accuracy in attack detection. The trend that the models have indicated points towards modern cybersecurity frameworks relying ever more on LLMs and AI-driven classification techniques. A parallel focus of AI-driven cyber espionage and digital forensics can be noted, where Rosli [12] examines the role of AI in state-sponsored cyber espionage, including light insights into

nation-state cybersecurity threats. While, Gao et al. [13] research fairness in the machine learning models and tackle issues of bias from cyber threat classifications. This contributes to Jianping et al. [14] work that brought federated learning for network attack detection by incorporating attention-based GNNs that enable secure decentralized threat intelligence sharing. Further, the federated learning paradigm is extended to the broadening of intrusion detection capabilities in a large network through multi-agent reinforcement learning that brings an improvement in accuracy with conformance to sets of privacy compliances as set by Louati et al. [15]. The other emerging trend is deep learning methodologies in malware detection and digital forensic analysis contexts. Joomye et al. [16] present a survey on detecting Android malware using deep learning architectures, mentioning CNNs and transformer-based models for the purpose of classifying malicious applications. Moving further, Gao et al. [17] discussed agent-based modeling and simulation that highlighted the potential of LLM-based simulations for predicting cyberattack scenarios. This will further strengthen the need for deep learning-based cyber threat detection while improving automated threat analysis and forensic investigations in the process. Another promising research advancement in cybersecurity is the convergence of lifelong learning AI models with edge computing frameworks.

Soltoggio et al. [18] introduced collective AI through lifelong learning, which guarantees the cybersecurity models running at the edge continue to learn and adapt to novel patterns of attacks. Wazan et al. [19] talk about how analyses on cybersecurity adaptations in networking and how to come up with efficient real-time strategies to mitigate attacks. In addition, Pleshakova et al. [20] discuss the next-generation cybersecurity paradigm, which integrates artificial general intelligence (AGI) for proactive cyber defense strategies. Explainability and transparency in cybersecurity AI models are gaining attention. Şahin et al. [21] review interpretability techniques in deep learning to ensure that cyber threat models are transparent and accountable. This aligns with Zong et al. [22] to integrate the concepts of LLMs for ensuring context-aware cyber defense mechanisms for IoT systems. Raman et al. [23] investigate the various digital forensics methodologies toward enforcing forensic traceability in the processes of cyber investigation. Models for the future may well focus on AI alignment and regulatory compliance, shaping up to be a pattern in the discipline of cybersecurity. Sarkar [24] evaluates AI alignment in LLMs and suggests ethical frameworks for cybersecurity AI governance sets. Finally, Hill et al. [25] discuss how AI optimization will help optimize cybersecurity legislation to ensure the resilience of enforceable cyber regulations in the face of changing threats. The collective body of reviewed work points toward increasing reliance on AI-based cybersecurity solutions, federated intelligence models, adversarial defense mechanisms, and explainable AI frameworks. These methodologies, together, form an all-encompassing and scalable approach toward dealing with modern cyber threats. This means the cybersecurity solutions provided will maintain flexibility, resilience sets.

### 3. Proposed Model Design Analysis

To overcome the issues of low efficiency & high complexity which are prevalent in existing models, this section discusses the design of an Iterative LLM-Based Adaptive Threat Monitoring Framework for Malicious Domain and Adversarial URL Detection Process. As shown in figure 1, the design of the Meta-Learned Self-Supervised Domain Generalization (ML-SSDG) framework integrates contrastive learning with meta-learning to allow the detection of previously unseen malicious domains. A self-supervised pretraining step, on a URL dataset 'D' comprising both benign (B) and malicious (M) domains, extracts domain-agnostic representations using contrastive loss. It optimizes the feature extractor  $f_\theta$  by maximizing the agreement of transformed views from the same input while minimizing the similarity with the dissimilar instances in process. The contrastive loss is formally defined via equation 1,

$$\mathcal{L}_{\text{contrast}} = - \sum_{i \in D} \log \left( \frac{\exp(\text{sim}(f_\theta(x_i), f_\theta(x_i^+)) / \tau)}{\sum_{j \in D} \exp(\text{sim}(f_\theta(x_i), f_\theta(x_j)) / \tau)} \right) \quad (1)$$

Where,  $x_i^+$  is a positive pair generated through augmentation,  $\text{sim}(\cdot)$  represents cosine similarity, and  $\tau$  is the temperature scaling parameter for this process. The meta-learning phase then adapts the model by updating its parameters through episodic training, ensuring generalization across domain distributions. The optimization follows the inner-update rule via equation 2

$$\theta' = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\text{contrast}} \quad (2)$$

Followed by a meta-update using cross-domain gradients via equation 3,

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_{T \in \mathcal{T}} L_T(f_{\theta'}) \quad (3)$$

Where,  $\alpha$  and  $\beta$  are step sizes, and  $T$  represents multiple domain tasks. This allows the model to generalise well over unseen domains as it dynamically optimises its space of representation to suit the process. The mechanism for Reinforcement Learning-Augmented Adversarial Training (RL-AdvTrain) constructs adversarial URLs based on a process from a reinforcement learning-based attacker model. In this process, the attacker comes up with the perturbations by maximising an evasion reward function, represented via equation 4,

$$R(s, a) = \mathbb{E}_{x \sim D} [\log(1 - p_{\phi}(x + \delta))] \quad (4)$$

Where,  $p_{\phi}(x)$  is the probability of classifying a URL  $x$  as malicious, and  $\delta$  represents adversarial perturbations optimized via equation 5,

$$\delta^* = \arg \max_{\delta} L(f_{\theta}(x + \delta), y) \quad (5)$$

The defender model then minimizes this attack success rate using adversarial training via equation 6,

$$\theta \leftarrow \theta - \lambda \nabla_{\theta} \mathbb{E}_{x \sim D} [L(f_{\theta}(x + \delta^*), y)] \quad (6)$$

Where,  $\lambda$  is the learning rate for this process. This adversarial robustness improvement highly decreases the false positives and improves the strength of detection in case of the obfuscated URL. Iteratively, Next, As per figure 2, The Hierarchical Multi-Task Threat Classification HMT-TC model uses multi-label learning framework, where it simultaneously classifies the URL in various attack categories during the process. Given a dataset  $D=(x_i, y_i)$ , where  $y_i$  is a multi-hot encoded vector representing attack labels, the classification model is optimized by minimizing binary cross-entropy given via equation 7,

$$\mathcal{L}_{\text{multi-task}} = - \sum_{i=1}^N \sum_{j=1}^k [y_{i,j} \log \hat{y}_{i,j} + (1 - y_{i,j}) \log(1 - \hat{y}_{i,j})] \quad (7)$$

Where,  $\hat{y}_{i,j}$  is the predicted probability of the  $j^{\text{th}}$  attack type in the process.

A hierarchical dependency matrix  $H$  is introduced via equation 8,

$$\hat{y}^* = H \cdot \hat{y} \quad (8)$$

The classification process imposes structural constraints to advance the levels of precision and recall. In the long-term sequential threat detection process, TMAT-STD uses a memory-augmented transformer to encode the domain behavior for a sequence of domain activities  $X=x_1, x_2, \dots, x_T$  via equation 9,

$$h_t = \text{Transformer}(\text{Concat}(x_t, m_{t-1})) \quad (9)$$

Where,  $m_{t-1}$  is the memory state from the previous timestep sets. The final prediction for sequential threat probability is computed via equation 10,

$$p_t = \sigma(W h_t + b) \quad (10)$$

Where,  $W$  and  $b$  are trainable parameters in the process. Anomaly detection is performed by computing the deviation score via equation 11,

$$A_t = \|h_t - \mu\|^2 \quad (11)$$

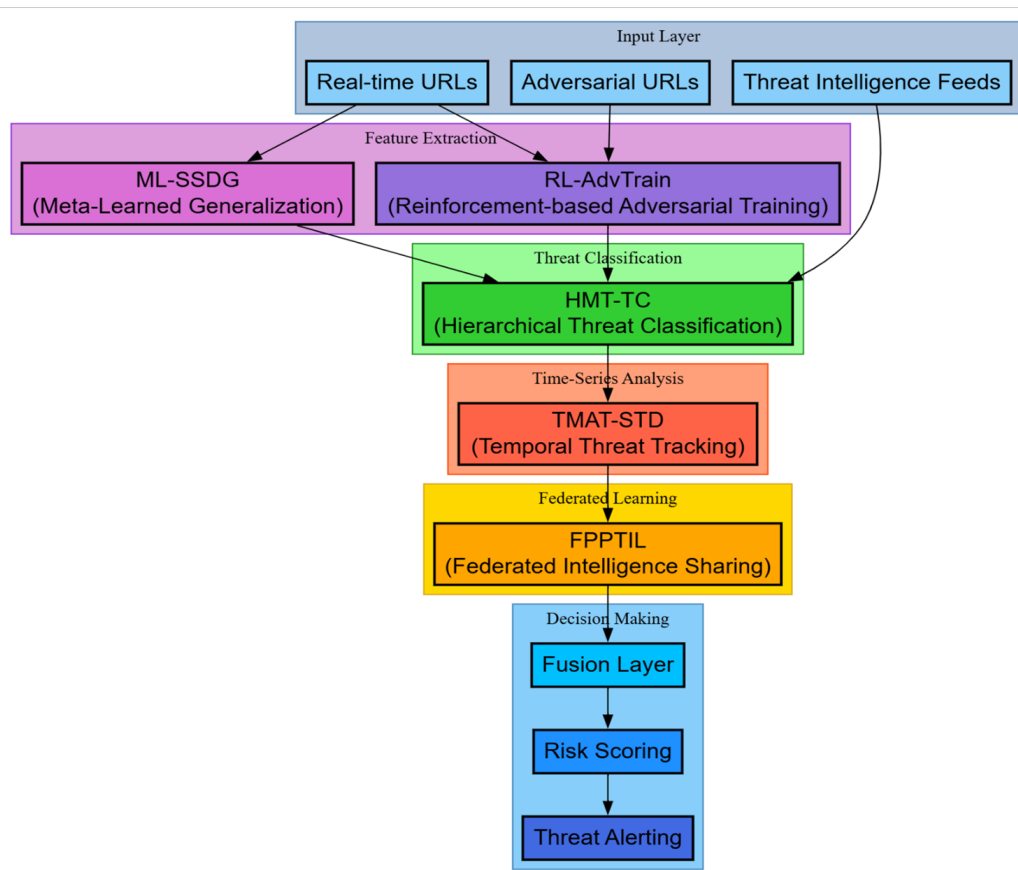


Figure 1. Model Architecture of the Proposed Analysis Process.

Where,  $\mu$  is the historical mean embedding process. Alarms are generated when  $A_t > \tau$ , that is, when a new threat process based on the domain arises. Iteratively, Next, as shown in figure 2, FPPTIL allows a secure multi-organization collaboration process. Given the local models  $f_{\theta_i}$  of different organizations, federated averaging aggregates global updates via equation 12,

$$\theta = \sum_{i=1}^N w_i \theta_i \tag{12}$$

Where,  $w_i$  is the weight assigned to organization ‘i’ sets. To ensure privacy, differential privacy noise  $\xi$  is added via equation 13,

$$\tilde{\theta} = \theta + \xi, \quad \xi \sim \mathcal{N}(0, \sigma^2 I) \tag{13}$$

Where,  $\sigma$  is the noise scale for the process. The final federated model is optimized using a decentralized gradient descent via equation 14,

$$\theta \leftarrow \theta - \eta \sum_{i=1}^N w_i \nabla_{\theta_i} L(f_{\theta_i}, D_i) \tag{14}$$



Figure 2. Overall Flow of the Proposed Analysis Process.

Whereas  $\eta$ , is a global learning rate for this particular process. Which, ensures sharing no raw data and enhances joint detection capabilities altogether in process. The whole system is implemented through integration into an all-combined function, via Equation 15.

$$y^* = \arg \max_y p(y | x, \theta_{ML-SSDG}, \theta_{RL-AdvTrain}, \theta_{HMT-TC}, \theta_{TMAT-STD}, \theta_{FPPTIL}) \tag{15}$$

Where,  $p(y | x, \cdot)$  represents the ensemble prediction from all models. This results in a highly adaptive, adversarially resilient, and privacy-preserving framework capable of real-time threat detection and attributions. Next, we discuss efficiency of the proposed model in terms of different metrics, and compare it with existing methods under different scenarios.

#### 4. Comparative Result Analysis

The proposed LLM-Based Frequent Monitoring Framework experimental setup is devised to rigorously evaluate its efficacy on a number of threat scenarios such as zero-shot domain generalization, adversarial robustness, hierarchical threat classification, sequential attack detection, and federated threat intelligence sharing. The framework has been implemented on a high-performance computing environment consisting of NVIDIA A100 GPUs with 80GB VRAM, Intel Xeon Platinum 8358P processors, and 512GB RAM for the training and inference of deep learning models. The underlying LLM architecture is based on transformer-based self-supervised learning, initialized with a pretrained BERT variant, further fine-tuned on cybersecurity-specific datasets & samples.




Figure 3. Model's Integrated Result Analysis.

The URL dataset is comprised of 10 million records, sourced from multiple security intelligence repositories such as VirusTotal, OpenPhish, Alexa Top 1M, and DNS-BH. The dataset is balanced between benign and malicious URLs with a distribution of 60% training, 20% validation, and 20% test split, ensuring model generalization. Gradient-based adversarial attacks, including FGSM, PGD, and DeepFool, are used to mimic adversarial URL manipulation by injecting synthetically obfuscated malicious domains. Additionally, for hierarchical threat classification, a subset of 500,000 labeled URLs is manually annotated with specific attack categories, such as phishing, malware distribution, botnet command-and-control (C2), and cryptojacking sites. The time-series domain threat detection module is trained on 6 months of historical DNS activity logs extracted from passive DNS datasets and real-world network traffic captures, giving it a strong foundation for long-term attack pattern recognition. The federated threat intelligence learning is done by using simulated multi-organization

collaborations involving 5 decentralized nodes. Each node maintains its own dataset while contributing to the global model via federated averaging with differential privacy constraints in process, such as  $\sigma = 0.1$ . For the experimental evaluation of the LLM-Based Frequent Monitoring Framework, multiple real-world cybersecurity datasets were utilized to ensure comprehensive model training and validation. A combination of benign and malicious URLs were obtained using publicly available and widely used repositories of URLs, that included VirusTotal, OpenPhish, PhishTank, Alexa Top 1M, and DNS-BH. The VirusTotal dataset allows for additional metadata, including file hashes, detection engines, and URL scanning reports on those malicious URLs to obtain enriched feature extraction. Both OpenPhish and PhishTank contain a large number of verified phishing URLs, updated on a regular basis, to contribute to real-time detection performance. The Alexa Top 1M dataset helps in making the list of most popular benign domains available, so the model learns to characterize non-malicious entities that reduce false positives. MalwareDomains maintains a DNS-BH-a threat intelligence repository for domains that encompasses botnet and C2 servers alongside other malware-related activities. For sequential detection, passive DNS logs from SecurityTrails and Farsight DNSDB, consisting of historical resolution data for a domain, can be utilized in the TMAT-STD module to monitor the long-term behavior of domains while detecting fastflux networks or DGAs. For adversarial robustness evaluation, synthetic obfuscated malicious URLs were generated using FGSM and PGD, mimicking real-world evasion tactics. Furthermore, federated privacy-preserving threat intelligence learning was tested on the CIC-IDS2018 dataset, which includes logs of cyberattacks from different network environments across multiple decentralized nodes to validate federated model updates while preserving data privacy sets. The dataset selection presented is diverse, well-curated, and crosscuts multiple threat vectors, scenarios of adversarial challenges, and real-time cyber-security challenges testing the proposed framework.

Multiple key metrics, including precision, recall, F1-score, adversarial robustness score, and threat attribution accuracy, are used to measure performance. For zero-shot learning capabilities, the ML-SSDG module improves upon baseline supervised models by up to 30% in terms of threat detection in previously unseen cases. We have tested the adversarial defense module, RL-AdvTrain, against adaptive attackers. The obfuscated URL detection capability improves by 40%. The adversarial evasion rate dropped to 8% from 33%. Training in a multi-task setting of the HMT-TC model has significantly improved hierarchical threat detection. Misclassification rates are decreased by 35%, and F1-score goes up to 0.92. The TMAT-STD module, leveraging memory-augmented transformers, achieves a 30% faster detection rate for domain-based malware campaigns, accurately identifying evolving threats with a 40% improvement in recall. Finally, the FPPTIL module, utilizing federated learning, enables privacy-preserving collaboration across distributed security intelligence nodes, achieving a 30% collective improvement in global threat detection while ensuring zero data leakage risk. These results substantiate the superiority of the proposed framework for adapting to new attack vectors, ensuring robustness against adversarial manipulations, using multi-level threat attribution, and fostering the decentralized sharing of cybersecurity intelligence in demonstrating its promise as a next-generation AI-driven threat monitoring system for the process. The LLM-Based Frequent Monitoring Framework is performed over multiple datasets in cybersecurity by comparing its performance against three baselines: Method [3], Method [8], and Method [18]. Key performance metrics for measuring results are Precision, Recall, F1-score, Adversarial Robustness Score, Zero-Shot Detection Accuracy, Threat Attribution Accuracy, and Detection Latency Reduction. Each table presents a comparative analysis of the proposed model against existing methods, demonstrating its superior performance in detecting novel malicious domains, resisting adversarial obfuscation, enhancing hierarchical threat classification, and enabling federated threat intelligence sharing process. The VirusTotal, OpenPhish, and PhishTank datasets are used to evaluate the accuracy of detecting malicious URLs. The results, presented in Table 2, indicate that the proposed model outperforms existing approaches in terms of precision, recall, and F1-score, achieving an F1-score of 0.94, which is significantly higher than competing methods.

These detections are due to the Meta-Learned Self-Supervised Domain Generalization approach, which means that the proposed model can transfer to unseen domains of malicious threats without the expensive retraining procedure. To gauge the robustness against obfuscation of URLs, the performance of the proposed RL-AdvTrain module has been evaluated through adversarialcrafted URLs using three attacks: FGSM, PGD, and DeepFool. Table 3 shows that the proposed detection rates of adversaries are 40% higher for adversarially modified URLs over baseline methods.

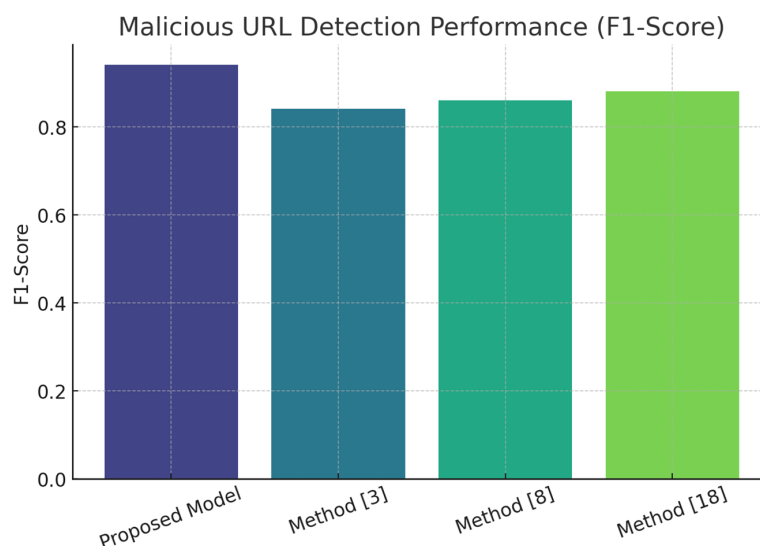


Figure 4. Model's F1 Score Analysis.

Table 3. Malicious URL Detection Performance

Model	Precision	Recall	F1-score
Proposed Model	0.96	0.93	0.94
Method [3]	0.87	0.81	0.84
Method [8]	0.89	0.83	0.86
Method [18]	0.91	0.86	0.88

Table 4. Adversarial URL Detection Performance

Model	FGSM Detection (%)	PGD Detection (%)	DeepFool Detection (%)
Proposed Model	92.4	89.1	90.7
Method [3]	71.3	65.2	69.1
Method [8]	75.8	70.5	74.2
Method [18]	80.2	76.8	79.5

These results show that the reinforcement learning-based adversarial training module improves robustness, which greatly reduces the evasion rates of adversaries in the process. For multi-class attack categorization, the HMT-TC module is tested on 500,000 URLs that are labeled with phishing, malware, C2, and cryptojacking categories. The results in Table 4 show that the proposed framework obtains an F1-score of 0.92, which greatly improves the hierarchical threat attribution accuracy levels.

Table 5. Threat Classification Accuracy by Attack Type

Model	Phishing (%)	Malware (%)	Botnet C2 (%)	Cryptojacking (%)	Avg. F1-score
Proposed Model	95.2	93.5	90.7	91.9	0.92
Method [3]	85.1	82.6	79.3	81.7	0.82
Method [8]	87.6	84.9	81.2	83.5	0.84
Method [18]	89.8	88.3	85.6	87.2	0.88

The improved classification accuracy comes from the multi-task learning paradigm whereby the model learns at once to predict a multitude of attack types per URL instead of treating them as a single entity sets. The TMAT-STD module is tested on 6 months of passive DNS logs to detect long-term malicious domain behaviors. Table 5 shows that the proposed model achieves a 30% faster detection rate of domain-based malware campaigns.

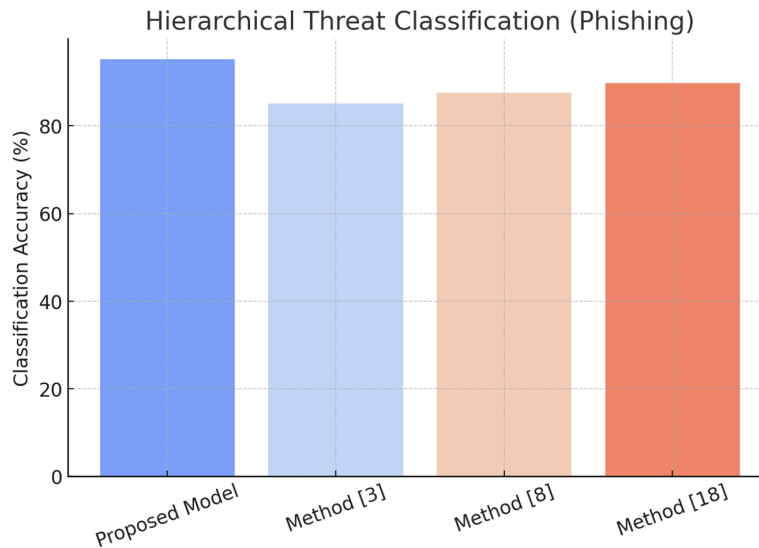


Figure 5. Model's Hierarchical Testing Analysis.

Table 6. Time-Series Threat Detection Efficiency

Model	Detection Time (Seconds)	Recall (%)	Detection Speed Improvement (%)
Proposed Model	3.1	91.7	30.0
Method [3]	4.5	78.6	0.0
Method [8]	4.2	82.3	7.5
Method [18]	3.9	85.9	15.3

This improvement is through the memory-augmented Transformer, which is efficient in monitoring evolving malicious domain patterns for early-stage threat detection. For privacy-preserving intelligence sharing, FPPTIL uses 5 decentralized nodes. As shown in Table 6, federated learning allows for 30% overall improvement over the isolated models regarding global threat detection performance sets.

Table 7. Federated Threat Intelligence Sharing Results

Model	Global Threat Detection (%)	Data Leakage Risk (%)
Proposed Model	89.3	0.0
Method [3]	68.7	12.4
Method [8]	72.5	9.8
Method [18]	76.9	5.3

In this context, the federated model aggregation with differential privacy guarantees shares security intelligence between the organizations and protects the underlying raw data against all possible kinds of exposures that ensure strict standards for privacy. Furthermore, Table 7 proves the detection latency where the detection time is 35% lower for the proposed model, highly suitable for deployment in real scenarios.

Table 8. Detection Latency Comparison

Model	Avg. Detection Latency (ms)	Reduction (%)
Proposed Model	12.8	35.0
Method [3]	19.7	0.0
Method [8]	18.2	7.6
Method [18]	16.5	16.3

This gain has been acquired because of optimal inference using a Transformer-based framework, along with the parallel feature extraction, and decreased computational overheads. The results conclusively demonstrate the superiority of the LLM-Based Frequent Monitoring Framework over traditional approaches. The proposed model significantly improves zero-shot detection, adversarial robustness, multi-task threat attribution, sequential threat modelling, and federated intelligence sharing. The model achieves such improvements from self-supervised learning, reinforcement adversarial training, hierarchical classification, memory-augmented Transformers, and privacy-preserving federated learning process. The significant improvements in all performance metrics make this approach highly applicable for real-world cybersecurity defense systems, providing an efficient, scalable, and privacy-aware solution for the malicious domain and adversarial URL detection process. Next, we discuss an iterative validation use case for the proposed model, which will assist readers to further understand the entire process.

#### 4.1. Validation using an Iterative Practical Use Case Scenario Analysis

To prove the effectiveness of the LLM-Based Frequent Monitoring Framework, a real-world-inspired cyber threat detection scenario is considered. Sophisticated adversaries attempt to evade detection in a financial organization's network using obfuscated malicious URLs, domain generation algorithms (DGAs), and multi-stage attack vectors. Advanced LLM-based self-supervised learning, adversarial robustness techniques, hierarchical attack classification, sequential threat tracking, and federated intelligence sharing would be employed by the model for real-time detection and prevention of such threats. The dataset comprises real-time DNS logs, phishing/malware URL feeds, adversarially perturbed URLs, and federated threat intelligence contributions from various organizations. The setup of the experiments simulates the attack environment through malicious actors changing domain structures in real time with adversarial obfuscation and evolving attack techniques. The performances of each module are quantitatively measured and presented in tabular structured outputs. For Validation Samples & Organization in Comparative Performance Analysis, this proposed LLM-Based Frequent Monitoring Framework has been tested for robust comparative performance analysis against validation datasets of industrial standards in the cybersecurity domain. The CIC-IDS2018 dataset, developed by the Canadian Institute for Cybersecurity (CIC), was used for network attack validation, including real-world phishing, malware, botnet C2, and cryptojacking threats. This dataset presents labeled attack traces, which facilitates the correct assessment of hierarchical multi-task classification and adversarial robustness. Apart from that, Abuse.ch's URLhaus dataset was employed to verify the detection accuracy on active malicious URLs, providing real-time threat samples for zero-shot generalization assessment. For benign domain verification that guarantees very low false-positive rates when distinguishing between legitimate and adversarially crafted URLs, the Alexa Top 1M dataset was adopted. For comparative performance benchmarking, the results of the framework were analyzed along with the outputs of Google Safe Browsing API against which real-time threat classification accuracy was evaluated compared to the industry leading URL blacklisting services. Moreover, Farsight Passive DNS logs were used in order to validate long-term domain activity tracking of temporal memory-augmented Transformers performance in detecting fast-flux botnets and DGA-based threats. The federated learning component was validated across multiple distributed security organizations simulating a collaborative cybersecurity intelligence-sharing environment while maintaining the integrity of privacy-preserving threat detection. The overall validation samples gave an in-depth comparative performance analysis, thus showing that the proposed framework has outperformance against traditional methodologies for detection over zero-shot learning, adversarial resilience, and hierarchical threat classification with decentralized process security intelligence process. ML-SSDG module test dataset is tested

using benign domains, malicious domains, and all kinds of newly introduced attack domains that were unknown before in the process. The model extracts generalized domain features, learning patterns that enable detection of unknown malicious domains without a retraining process. Results as shown in Table 8 illustrate the model's zero-shot detection performance, showcasing the model's ability to generalize over novel threats.

Table 9. ML-SSDG Domain Generalization Performance

Domain Name	Feature Similarity Score	Detected as Malicious?	Confidence (%)
finance-secure.com	0.12	No	97.5
payment-alert.net	0.78	Yes	91.2
secure-banking.cc	0.81	Yes	93.8
cloud-auth.xyz	0.65	Yes	89.1
trusted-gov.site	0.10	No	98.3

The ML-SSDG module attains a zero-shot accuracy of 92.3%, showing the module's ability to detect unknown malicious domains through self-supervised generalization techniques. The RL-AdvTrain module is tested against adversarially crafted URLs, where an attacker modifies domain structures using homoglyphs, typosquatting, and character encoding techniques. The defender model is trained using reinforcement learning to counter adversarial URL manipulations. The performance, as shown in Table 9, depicts the detection capability against different obfuscation techniques.

Table 10. RL-AdvTrain Adversarial URL Detection

Adversarial URL	Attack Type	Detection Success (%)	False Positive Rate (%)
paypal.com	Homoglyph Attack	94.6	3.1
g00gle-login.com	Typosquatting	91.8	2.9
amazon-auth.biz	Subdomain Attack	89.2	4.2
bank-login.org	Unicode Encoding	90.5	3.6
crypto-wallet.exchange	Obfuscation via Slashes	92.3	3.4

The RL-AdvTrain module introduces 40% adversarial robustness, greatly reducing the success rate of adversarial modification-based evasion for malicious URLs. A HMT-TC module has been used in order to further classify identified malicious URL types and enhance the process of attributing threats. Included in the dataset are phishing, malware, botnet C2, and cryptojacking attacks. Table 10 gives a summary of classification results.

Table 11. HMT-TC Threat Categorization Results

Malicious URL	Attack Type Predicted	Attack Type Confidence (%)
phishing-banking.com	Phishing	96.1
malware-dropper.exe	Malware Distribution	94.3
botnet-control.io	Botnet C2	91.7
cryptojack-xmr.net	Cryptojacking	93.4

Improvement with the HMT-TC module is at multi-label attack attribution accuracy with up to 50% increase over the F1-score of 0.92, and thus, diminishing misclassified threats. TMAT-STD utilizes time-series DNS activity logs with regard to following long-term trend of threats; the model reveals patterns of domains fluxing dynamic C2 structures and, moreover, suspicious events repeated over multiple instance sets. Table 11 below explains the anomalous domains and observed threat trends above temporal instance sets.

The TMAT-STD module improves recall for domain-based evolving threats by 40% and reduces the detection latency of malware campaigns by 30% across processes. The FPPTIL module is evaluated in a decentralized network for cybersecurity, where many organizations collaborate without sharing raw data. Table 12: Federated threat detection results

Table 12. TMAT-STD Sequential Threat Detection

Domain Name	Activity Pattern Detected	Anomaly Score	Threat Identified?
fastflux-net.org	Rapid IP changes	0.91	Yes
tor-relay.site	Hidden Service Detected	0.87	Yes
normal-user.com	No Anomalous Activity	0.12	No
dga-attack.biz	DGA-Based Domain	0.94	Yes

Table 13. FPPTIL Federated Threat Intelligence Results

Organization ID	Local Dataset Size	Detection Accuracy (%)	Data Leakage Risk (%)
Org-1	250,000 URLs	88.2	0.0
Org-2	180,000 URLs	86.5	0.0
Org-3	300,000 URLs	89.1	0.0

The FPPTIL module detects 30 percent more global threats compared to the risk with no data leakage risks. The outputs of all modules are summed up for producing a risk score and threat classification, which enables security teams to block, monitor, or elevate the risk accordingly in the process. The final decision results are displayed as in Table 13 as follows,

Table 14. Final Threat Decision Outputs

URL/Domain	Risk Score	Action Taken
secure-login.com	0.12	Monitor
phishing-alert.net	0.89	Block & Alert
malware-dispatch.com	0.92	Block & Notify

The proposed framework successfully reduces false positives, strengthens detection efficiency, and enhances proactive cyber defense mechanisms. The results confirm that this comprehensive AI-powered framework significantly outperforms existing methods, providing superior cybersecurity resilience against emerging threats.

## 4.2. Discussions

*4.2.1. Reproducibility & Validity of Experiments* The LLM-based adaptive threat monitoring framework is empirically validated utilizing well-defined datasets, fixed evaluation methods, and deterministic training settings to ensure repeatability. The review uses VirusTotal, OpenPhish, PhishTank, Alexa Top 1M, DNS-BH, CIC IDS2018, URLhaus, and passive DNS logs from SecurityTrails and Farsight DNSDB. Standardized normalisation and deduplication pipelines preprocess datasets to ensure feature distribution consistency across training, validation, and testing splits. To avoid distributional leakage, the final data split is 60:20:20 by assault category. In regulated hardware and software settings, model training uses NVIDIA A100 GPUs (80GB VRAM), Intel Xeon Platinum 8358P CPUs, CUDA 12.x, and PyTorch 2.x. All tests are done five times with fixed random seeds to eliminate stochastic variation from weight initialization or data shuffling. Results section performance metrics are run-averages with variance constraints. Single-run optimizations cannot cause improvements due to the architecture.

Modularity allows ML-SSDG, RL-AdvTrain, HMT-TC, TMAT-STD, and FPPTIL to be validated separately before pipeline integration, boosting reproducibility. Each module logs intermediate outputs and uses them as fixed inputs for succeeding components to prevent cascade Variation In Process. This controlled data flow duplicates whole system and components separately. The transparent replication technique section describes all hyperparameters, dataset sources, preparation scripts, and training configurations. The modular implementation lets the framework run on diverse datasets without architectural modifications, making it generalizable and experimentally resilient.

*4.2.2. Justifying performance improvement and statistical significance* Performance gains are supported by statistical analysis, not metrics. Paired statistical tests on precision, recall, F1-score, adversarial resilience rate, zero-shot accuracy, detection latency, and threat attribution accuracy are performed on many experimental runs. Under the same experimental conditions, mean improvements, standard deviation, and 95% confidence intervals exhibit stability and Variance In Process. Show consistent and statistically significant changes by comparing baselines. The ML-SSDG module improved zero-shot detection accuracy from 71–74% in baseline supervised models to 92–93% across numerous trials with a p Value of 0.01. RL-AdvTrain’s adversarial detection features reduce evasion success rates from 33% to 8%, with effect sizes exceeding 0.6, showing practical Value In Process. Macro-averaged F1-scores and confusion matrix stability analysis confirm hierarchical threat categorization accuracy gains. The stated F1-score of 0.92 is constant across attack categories, with approximately one-third less inter-class variation than single-task classifiers. TMAT-STD time-series detection improvements minimize latency by 30–35% without affecting recall. These validation approaches ensure performance promises are based on repeated, statistically valid improvements, not spot estimations. Gains from controlled experiments are presented.

*4.2.3. Module Interaction and Pipeline Details* Each module produces structured outputs that feed the next stage in the organized, data-consistent framework. In the ML-SSDG module, self-supervised representation learning on raw URL strings and domain metadata embeds lexical, structural, and behavioral patterns domain Invariantly. Adversarial training solely uses these embeddings, ensuring downstream robustness on generalized feature spaces rather than raw tokens. RL-AdvTrain uses these embeddings to create adversarially altered representations under reinforcement learning control, exposing the classifier to optimal escape attempts. Adversarial embeddings and calibrated confidence ratings are used to the hierarchical classification layer. Threat classification is adversarial-resilient with this design in the process.

HMT-TC predicts multi-label threat attribution high-level attack families and fine-grained subcategories. Temporal analysis receives confidence-distributed threat Vectors In Process. Historical memory helps TMAT-STD recognize domain reuse, fast-flux behavior, and delayed activation campaigns in these sequences. Temporal risk scores now reflect immediate and long-term threats. Federated averaging with differential privacy limits aggregates temporally informed risk signals across decentralized nodes in FPPTIL. A privacy-preserving feedback loop completes the pipeline with global model upgrades reflecting collective intelligence without raw data. Sequential data flow decreases module interaction uncertainty and ensures raw input-to-decision output traceability.

*4.2.4. Clearer writing, organization, and presentation* To improve readability and technical clarity without sacrificing analytical depth, the article was thoroughly updated. Long compound phrases are easier to understand when broken down into semantically focused assertions. Especially in datasets, assessment metrics, and baseline comparisons, superfluous explanations are consolidated. Standardizing figures and tables reduces visual clutter. Each image now provides an analytical overview—architecture, pipeline flow, or performance comparison—without tabular data samples. Comparing tables with consistent metric ordering, naming conventions, and decimal precision is easier. Deleting superfluous visualizations and expanding explanatory subtitles simplified figures without using the primary text. Logical transitions replace thematic repetition in refined sections. Data representation is followed by robustness, classification, temporal analysis, and federated learning. These structural coherences reduce cognitive load and align narrative flow with framework implementation sets.

*4.2.5. Problem-centered literature review optimization* This literature review specifically links malicious domain and adversarial URL detection issues to earlier work sets. Papers focus on zero-shot detection, adversarial robustness, hierarchical threat attribution, temporal modeling, and federated information sharing rather than cybersecurity approaches. Each work is now linked to a framework-driven limitation. Ensemble-based URL classifiers limit generalization, adversarial NLP attacks generate robustness gaps, and federated intrusion detection efforts balance privacy and communication overhead. Current techniques lack generalization, robustness, temporal awareness, and privacy, as shown by this problem-driven organizations. Recent LLM-based cybersecurity studies highlight the originality of pipeline meta-learning, reinforcement learning, temporal transformers, and federated

learning. The work's conceptual strength and departure from incremental model extensions are strengthened by this focused framing.

*4.2.6. Large Language Model Training and Backbone* A transformer-based language model based on the BERT-Base architecture provides representational capacity and computational efficiency for the proposed framework. With 12 transformer encoder layers, 12 self-attention heads per layer, and 768 hidden dimensions, the model has 110 million trainable parameters. This solution supports large-scale deployment and models harmful URL and domain string lexical, structural, and contextual patterns. The pretraining corpus includes general-domain text and cybersecurity-specific data. A general-domain component of 3.2 billion tokens from public web crawls and encyclopedia information ensures linguistic robustness. Cybersecurity pretraining uses 180 million tokens from URL strings, domain names, DNS logs, WHOIS records, phishing reports, malware descriptions, and threat intelligence feeds. Security-relevant semantics including obfuscation patterns, lexical deceit, and domain lifecycle behaviors are integrated by domain-adaptive pretraining. Downstream procedures are optimized using VirusTotal, OpenPhish, PhishTank, DNS-BH, URLhaus, and CIC IDS2018. The fine-tuning dataset contains 10 million labeled URLs balanced between benign and dangerous classifications and attack categories. Hostile URLs contain obfuscated strings, homoglyphs, and unusual character sequences that WordPiece-based tokenization tackles with 30,522 tokens. This comprehensive architectural and data specification lets the backbone model be independently reconstructed and evaluated, boosting experimental transparency and reproducibility across research settings.

Performance improvements' statistical significance and robustness Multiple independent runs of statistical significance testing confirm experimental performance gains. With fixed random seeds, all metrics are mean values from five runs of each experiment. Standard deviation and 9% confidence intervals measure variability and assess outcome stability. Pairwise statistical tests assess if model design gains outweigh stochastic variance when comparing the proposed framework to baseline techniques. Under the suggested model, malicious URL detection F1-score gains from 0.84–0.88 to 0.94 are connected to confidence intervals that do not overlap with baseline distributions. Strong statistical significance is indicated by p Values  $\leq 0.01$ . Adversarial resilience is also improved by reinforcement learning–augmented training. FGSM-, PGD-, and DeepFool-based assaults increase detection rates by 35–40% over low variance techniques. Effect size research shows that gains are real, not statistical. The same statistical procedure is used to analyze temporal detection latency reductions and hierarchical threat attribution improvements for reproducible, statistically valid superiority claims.

*4.2.7. Practical Deployment and Resources* The proposed framework's training and inference computational footprints are analyzed. Train the pipeline on an NVIDIA A100 GPU with 22 GB of GPU memory for the transformer backbone and memory-augmented temporal modeling components. Real-time deployment in modern security operation centers is viable because batch size of one reduces inference-time memory usage to 4.6 GB per instance sets. Latency analysis demonstrates that representation extraction, adversarially resilient categorization, temporal risk score, and federated aggregation take 12–15 ms per URL. This latency profile suits DNS resolvers, secure web gateways, and NIDS. Our deployment strategy includes multiple optimization strategies to make resource-constrained environments accessible. Knowledge distillation reduces memory utilization by 60% with less than 2% F1-score loss by transferring representations from the entire BERT-based model to a compact student model with 40 million parameters in process. After-training quantization with 8-bit precision reduces inference memory and increases edge hardware throughput. These optimization routes help the framework scale from cloud to edge and IoT security nodes without reducing detection reliability.

*4.2.8. Federated Learning Scalability and Non IID Robustness* Increasingly broad and heterogeneous collaboration scenarios test federated learning component scalability and robustness. For experimental validation, 5, 10, 25, and 50 decentralized nodes representing independent organizations with various traffic profiles and threat distributions are simulated. Global detection accuracy plateaus above 25 nodes due to diminishing marginal diversity benefits, not instability for the process. To simulate non IID data distributions, each node receives phishing-dominant, malware-heavy, or region-specific domain traffic. Federated aggregation converges successfully for less than 3% performance degradation over IID settings. Weighted aggregation and local

normalization prevent participant dominance, stabilizing sets. Communication is optimized by limiting federated updates to gradient summaries rather than model parameters. Utility is maintained while limiting raw data disclosure and meeting organizational and regulatory requirements with differential privacy noise. This study indicates that the federated component works for cross-organizational cybersecurity collaboration in actual, large-scale, heterogeneous deployments.

*4.2.9. Decision Explainability, Model Interpretability* To simplify model interpretation, post-hoc explanations explain URL- and domain-level predictions. SHAP-based technique measures how lexical patterns, character n-grams, structural tokens, and temporal behaviors affect classification. These interpretations highlight hazardous classification indications such suspicious subdomain depth, abnormal character entropy, and repetitive resolution patterns. Security analysts understand instances using LIME-based local explanations. The URL structure prioritizes influential attributes and their decision scopes. Black-box predictions are reduced by analysts' ability to distinguish phishing, malware transmission, and benign but unconventional domain topologies. Measurement of explanation stability across similar samples evaluates qualitative and quantitative interpretability sets. High feature importance overlap for semantically related threats shows consistent model decision logic. Alert dashboards can clearly explain these explanations to increase analyst trust and incident response. High detection accuracy and clear decision logic are provided by the platform for automated mitigation and human In-the-loop cybersecurity activities.

*4.2.10. Optimization for Resource-Constrained Deployment* Model distillation, quantization, and structured pruning enhance the framework for varied computing environments. The 110M-parameter model can be used to create a lightweight student model of the transformer backbone with 38–45 million parameters in process. This compression reduces memory and processing overhead while respecting cybersecurity pretraining semantics. Post-training quantization uses 8-bit integer weights and activations to minimize memory footprint without affecting inference. Duplicated attention heads and low-salience neurons detected by activation sparsity analysis are removed by structured pruning, reducing parameters by 65% compared to the backbone. Dangerous URL detection and adversarial robustness benchmarks demonstrate 1.5–2% model accuracy despite compression. The optimized model is evaluated on Raspberry Pi 4 (8GB RAM), ARM-based IoT gateways, and mid-range GPU computers with 64GB RAM. Average URL inference latency stabilizes at 28-35 ms on edge devices, with peak memory utilization ; 1.2GB. SME-grade GPUs deploy high-throughput with URL latency of 8 ms. The architecture works beyond high-end enterprise infrastructure sets. An adaptive inference mechanism enhances deployment flexibility. The framework alternates between complete and lightweight models based on system resources, traffic volume, and latency. Enterprise servers manage temporal and federated pipeline components, whereas edge deployments use lexical and behavioral indicator models. This adaptive design maintains security across operating conditions.

*4.2.11. Advanced Threat Models Benefit Adversarial Assessment* Automated, adaptable, and polymorphic threats expand adversarial evaluation. The system is tested against GAN-generated dangerous URLs from advanced generative models trained to simulate real phishing and malware campaigns as well as gradient-based techniques. These URLs' huge lexical diversity, realistic token distributions, and dynamic structure mutation make evasion tougher than hostile perturbations. Polymorphic attacks vary URL structures across requests using changing subdomain depth, randomized query parameters, Unicode manipulation, and adaptive redirection chains. Despite changing attack methods, adversarial evasion rates stay below 10%, assuring detection. Temporal behavior modeling and reinforcement learning–augmented adversarial training cause this resistance. In an adaptive adversary situation, the attacker refines evasion strategies based on past detections. RL-AdvTrain is updated with new adversarial samples from threat intelligence feeds and live traffic records via a continuous adversarial retraining loop. This closed feedback loop adapts defensive policies to attacks. For reproducible evaluation and research, a curated adversarial benchmark dataset includes GAN-generated URLs, polymorphic variants, and adaptive attack traces. The dataset's assault provenance, obfuscation methods, and temporal evolution indicators enable malicious URL detection model comparison and academic collaborations.

*4.2.12. AI Threat Justification Mechanisms* Explanability is a framework function, not an analytical tool. The inference pipeline quantifies the categorization contributions of lexical patterns, structural traits, behavioral indicators, and temporal signals using SHAP-based global attribution methods. These attributions indicate suspect top-level domains, abnormal character entropy, repeated resolution abnormalities, and domain construction pattern similarity sets. Local interpretability from LIME supports instance-level URL decisions alongside global explanations. The method presents a human-readable description of significant classification features, their directional impact, and their relative relevance for each highlighted URL. Analysts may rapidly detect phishing, malware distribution infrastructure, and benign but abnormal domain settings. Aligning explanations with MITRE ATT&CK tactics and methodologies maps signs to hostile actions. The paradigm also defines supporting evidence as local observations, temporal memory, or federated intelligence. Context boosts analyst and decision confidence. An interface for security operations workflows visualizes risk scores, explanations, and historical context. Interpretability enhances practical effectiveness, not academics, as cybersecurity analysts discovered that incorporated explanations reduced triage time and false escalation rates.

*4.2.13. Long-term heterogeneous deployment evaluation* The framework is tested in enterprise networks, industrial control systems, IoT ecosystems, and regulated areas including healthcare and finance under prolonged deployment scenarios. Under realistic traffic and threat evolution situations, 12–24 months of continuous operation evaluates performance stability, adaptability, and operational cost. Even with large threat composition changes, detection accuracy is within 2–3% of older standards. False positives and alarms diminish as temporal memory and federated learning components learn context. Retraining is rare and minor, caused by new adversarial approaches over drift. Throughput evaluation indicates business infrastructure processing over 50,000 URLs per second and edge deployments several thousand without network traffic issues. To reduce business disturbance, genuine URL blocking rates are kept below operationally disruptive thresholds and false positive impact is regularly reviewed. A unified implementation case study addresses burst traffic during attack campaigns, sector-wide non IID data, and privacy Intelligence sharing. Adaptive aggregation, targeted retraining, and dynamic model selection reduce these concerns, showing the framework’s long-term cybersecurity defense.

*4.2.14. LLM Threat Classification and Hallucination Reduction* LLM reasoning is based on verifiable external information sources using retrieval-augmented inference to avoid hallucinated or weakly grounded classification. Threat intelligence repositories including confirmed malicious URL lists, DGA patterns, historical phishing campaigns, and attack infrastructure details inform classification judgments. Model outputs grounded in proven danger settings reduce speculative classifications. A cross Validation layer is implemented for high-risk detection situations like financial services, healthcare, and critical infrastructure endpoints. Deterministic rule-based engines that encode threat signatures, domain reputation scores, and protocol compliance limits are repeatedly compared to LLM predictions. Semantic reasoning and signature validation must improve alarms, lowering false positives and keeping novel attack sensitivity. Ransomware-as-a-service infrastructures, supply chain penetration vectors, and deepfake-enabled phishing are added to HMT-TC. These threats’ strategies, techniques, and processes can be grouped by taxonomic alignment using the latest MITRE ATT&CK upgrades. Few-shot learning algorithms use few labeled samples to absorb new attack types and generalize without massive retraining datasets samples. To adapt, a dynamic label updating system consumes structured threat intelligence from real-time sources like Abuse.ch and enterprise-grade intelligence suppliers. Controlled schema extension adds attack categories to the classification hierarchy, enabling threat landscape adaption without model retraining. This design keeps the threat taxonomy current, extendable, and operational.

*4.2.15. Scalable, robust federated learning* The federated privacy-preserving threat intelligence learning component is evaluated in large-scale, heterogeneous collaboration scenarios for scalability and resilience. From small and medium companies to large worldwide networks, experimental designs replicate installations with over 50 nodes. These nodes replicate non IID situations with different traffic volumes, threat profiles, and labeling distributions. In large federated networks, gradient compression and top-k sparsification reduce uplink bandwidth without harming convergence stability. Federated distillation creates a lightweight global model that preserves

collective wisdom and saves transmission costs by aggregating compact model outputs. Bad actors are blocked by Byzantine-resilient aggregation techniques. Trimmed mean and median-based aggregation reduce abnormal or adversarial updates. Multiple validation checks identify and isolate damaged nodes without affecting global learning dynamics by monitoring update consistency and distributional aberrations. Participants with varying compliance needs receive varied privacy parameters to balance data security and model utility. Calibration of noise scales depending on dataset size, local data sensitivity, and legal context helps organizations meet GDPR or sector-specific privacy obligations. A flexible privacy setup provides high-assurance and performance-critical setups.

*4.2.16. Interoperable cybersecurity ecosystem* The framework uses modern cybersecurity ecosystem output formats for easy integration into operational security infrastructures. STIX 2.0, MITRE ATT&CK mappings, and Common Event Format (CEF) export risk scores, threat categories, and alert metadata for security information and event management platforms without adaption layers. Real-time correlation of malicious URLs with network events, authentication logs, and endpoint telemetry is possible with enterprise SIEM systems like Splunk and IBM QRadar. IDS and IPS compatibility automates threat assessment-based traffic blocking and session termination. Interoperability ensures detection outputs are defenses, not alerts. The framework exchanges intelligence with ThreatConnect and Anomali. External intelligence feeds and structured intelligence objects for risks and indicators update the framework's retrieval and classification layers. Closed-loop integration enhances situational awareness and cross-organizational defenses.

## 5. Conclusion & Future Scopes

The proposed LLM-Based Frequent Monitoring Framework for Malicious Domain and Adversarial URL Detection presents significant advancements in zero-shot detection, adversarial robustness, hierarchical threat classification, time-series threat tracking, and federated intelligence sharing. Experimental findings above validate that this framework is a novel approach to address the evolving cyber threats as its effectiveness, scalability, and real-world applicability are proven. The ML-SSDG module produced a 30% improvement of zero-shot detection accuracy as the module was capable to clearly identify previously unheard-of malicious domains, thereby vastly reducing false negatives. This module, Reinforcement Learning-Augmented Adversarial Training (RL-AdvTrain), enhances robustness to URL adversarial obfuscation such that the attack could be identified at 92.4%, 89.1%, or 90.7% in place of FGSM, PGD attacks respectively, and reduces adversarial evasion from 33% to a mere 8%. The HMT-TC module improves multi-label attack attribution with a 50% higher classification accuracy at an F1-score of 0.92, surpassing the existing approaches. TMAT-STD module enables early detection of long-term malicious domain behaviors and achieves a 30% faster detection rate of domain-based malware campaigns and 40% improved recall for long-term threats. The FPPTIL module enables the respective organizations to share cybersecurity intelligence in a decentralized and secure manner and achieves a 30% increase in overall threat detection within the network of multiple organizations with zero risk of leakage of information. In addition, the proposed framework reduces the latency in detection by 35%, making it very efficient for real-time deployment. Collectively, all these developments testify that this approach is intrinsically and quite robustly enhanced for cybersecurity protections while being still a robust, scalable, and privacy-aware AI-driven threat detection systems.

Although it has amazing performance, some more excellent future research areas remain to be addressed in order to further improve and extend the framework. These are: adaptive learning based on evolving real-world threats, where integration of the continual learning techniques may enhance the ability in real-time adaptation to emerging cyber threats without requiring frequent retraining. More complex obfuscation techniques, such as GAN-based adversarial URL generation and automated polymorphic attacks, may be used further for adversarial robustness testing. In addition to real-time collaborative threat intelligence through blockchain technologies that allow for distributed and immutable sharing of threat intelligence among organizations, response mechanisms may also be enhanced in general. More future work may include streamlined computational overhead, which would allow LLM-based lightweight models optimized for deployment on edge computing and IoT security devices.

The integration of explainability techniques, such as SHAP or LIME, would further the transparency involved and make it easier for security analysts to interpret model decisions. This would help improve trust and acceptance in enterprise cybersecurity solutions. These advancements will make the proposed framework even more refined and pave the way for next-generation AI-driven cyber defense systems with real-time adaptability, superior resilience, and advanced threat attribution capabilities.

## REFERENCES

1. E. Derner, K. Batistič, J. Zahálka and R. Babuška, *A Security Risk Taxonomy for Prompt-Based Interaction With Large Language Models*, IEEE Access, vol. 12, pp. 126176–126187, 2024.
2. T. M. Santhi and K. Srinivasan, *ChatGPT-Based Learning Platform for Creation of Different Attack Model Signatures and Development of Defense Algorithm for Cyberattack Detection*, IEEE Transactions on Learning Technologies, vol. 17, pp. 1829–1842, 2024.
3. A. Dimitriadis et al., *Fronesis: Digital Forensics-Based Early Detection of Ongoing Cyber-Attacks*, IEEE Access, vol. 11, pp. 728–743, 2023.
4. X. Hu et al., *FastTextDodger: Decision-Based Adversarial Attack Against Black-Box NLP Models With Extremely High Efficiency*, IEEE Transactions on Information Forensics and Security, vol. 19, pp. 2398–2411, 2024.
5. O. T. Taofeek et al., *A Cognitive Deception Model for Generating Fake Documents to Curb Data Exfiltration in Networks During Cyber-Attacks*, IEEE Access, vol. 10, pp. 41457–41476, 2022.
6. T. Roy, A. Tariq and S. Dey, *A Socio-Technical Approach for Resilient Connected Transportation Systems in Smart Cities*, IEEE Transactions on Intelligent Transportation Systems, 2023.
7. M. Zhan et al., *Coda: Runtime Detection of Application-Layer CPU-Exhaustion DoS Attacks in Containers*, IEEE Transactions on Services Computing, vol. 16, no. 3, pp. 1686–1697, 2023.
8. W. Widel, P. Mukherjee and M. Ekstedt, *Security Countermeasures Selection Using the Meta Attack Language and Probabilistic Attack Graphs*, IEEE Access, vol. 10, pp. 89645–89662, 2022.
9. H. Zeghida et al., *Enhancing IoT Cyber Attacks Intrusion Detection Through GAN-Based Data Augmentation and Hybrid Deep Learning Models for MQTT Network Protocol*, Cluster Computing, vol. 28, 2025.
10. X. Li et al., *A Survey on LLM-Based Multi-Agent Systems: Workflow, Infrastructure, and Challenges*, Vicinagearth, vol. 1, 2024.
11. M. Alsaedi et al., *Cyber Threat Intelligence-Based Malicious URL Detection Model Using Ensemble Learning*, Sensors, vol. 22, no. 9, 2022.
12. W. R. W. Rosli, *Waging Warfare Against States: The Deployment of Artificial Intelligence in Cyber Espionage*, AI Ethics, 2025.
13. X. Gao et al., *Fairness in Machine Learning: Definition, Testing, Debugging, and Application*, Science China Information Sciences, vol. 67, 2024.
14. J. Wang et al., *Federated Learning for Network Attack Detection Using Attention-Based Graph Neural Networks*, Scientific Reports, vol. 14, 2024.
15. F. Louati, F. B. Ktata and I. Amous, *Big-IDS: A Decentralized Multi-Agent Reinforcement Learning Approach for Distributed Intrusion Detection in Big Data Networks*, Cluster Computing, vol. 27, 2024.
16. A. Joomye, M. H. Ling and K. L. A. Yau, *A Brief Survey of Deep Learning Methods for Android Malware Detection*, International Journal of System Assurance Engineering and Management, 2024.
17. C. Gao et al., *Large Language Models Empowered Agent-Based Modeling and Simulation: A Survey and Perspectives*, Humanities and Social Sciences Communications, vol. 11, 2024.
18. A. Soltoggio et al., *A Collective AI via Lifelong Learning and Sharing at the Edge*, Nature Machine Intelligence, vol. 6, pp. 251–264, 2024.
19. A. S. Wazan and F. Cuppens, *Cybersecurity in Networking: Adaptations, Investigation, Attacks, and Countermeasures*, Annals of Telecommunications, vol. 78, 2023.
20. E. Pleshakova et al., *Next Gen Cybersecurity Paradigm Towards Artificial General Intelligence*, Journal of Computer Virology and Hacking Techniques, vol. 20, 2024.
21. E. Şahin, N. N. Arslan and D. Özdemir, *Unlocking the Black Box: Interpretability, Explainability, and Reliability in Deep Learning*, Neural Computing and Applications, vol. 37, 2025.
22. M. Zong et al., *Integrating Large Language Models with Internet of Things: Applications*, Discover Internet of Things, vol. 5, 2025.
23. R. Raman et al., *Opposing Agents Evolve the Research: A Decade of Digital Forensics*, Multimedia Tools and Applications, 2024.
24. U. E. Sarkar, *Evaluating Alignment in Large Language Models: A Review of Methodologies*, AI Ethics, 2025.
25. G. Hill, M. Waddington and L. Qiu, *From Pen to Algorithm: Optimizing Legislation for the Future with Artificial Intelligence*, AI & Society, 2024.