Article

# An Adaptive Architecture for Cybersecurity Threat Intelligence: A Case Study on Kenyan Courts

**Paul Okanda**
https://orcid.org/0000-0001-5215-4368
United States International University-Africa
pokanda@usiu.ac.ke

**Sarah W. Muriithi**
https://orcid.org/0009-0009-4343-8022
United States International University-Africa
sarhurwakowthay@gmail.com

## Abstract

Current threat intelligence systems often lack scalable, adaptive AI architectures capable of delivering real-time incident detection and dynamic response, particularly in resource-constrained environments. This paper presents a novel AI-driven architectural design for operational threat intelligence, specifically tailored to enhance cybersecurity in global and Kenyan judiciaries. The proposed model integrates three foundational frameworks, which are Integrated Adaptive Cyber Defence (IACD), the Cyber Kill Chain, and Moving Target Defence (MTD), into an architecture that supports real-time data ingestion, continuous AI model retraining, and automated response orchestration. The research design for this study adopts a mixed-methods approach, combining qualitative and quantitative methods to ensure a comprehensive understanding of the AI-driven operational Cyber Threat Intelligence (CTI) model. Key features include a dynamic feedback loop for adaptive learning, AI-powered multi-stage threat detection aligned with attack lifecycle mapping, and resource-efficient dynamic defence mechanisms suitable for low-resource judicial environments. This design significantly improves incident response capabilities by enabling faster, more accurate threat detection and automated mitigation, reducing mean time to detect and respond. By providing a scalable, transparent, and explainable AI model, the architecture offers a practical blueprint for enhancing cybersecurity resilience in judicial systems worldwide, with applicability to the unique challenges faced by Kenyan courts. This work lays the foundation for future extensions involving federated learning to enable secure, multi-court deployments, further strengthening collective judicial cybersecurity defences.

**Keywords:** cybersecurity; incident management; real-time threat detection; cyber

UNISA university of south africa PRESS
African Journal of Security
#20239 | 30 pages

https://doi.org/10.25159/3005-4222/20239
ISSN 3005-4222 (Online)

threat intelligence; Artificial Intelligence

## Introduction

Operational Cyber Threat Intelligence (OCTI) in modern cybersecurity enables organisations to proactively identify, assess, and mitigate cyber threats in real time (Dimitriadis et al. 2025). As cyberattacks grow in complexity and frequency, traditional threat intelligence systems, which are often reliant on static, signature-based detection, struggle to keep pace with the advanced threats and vulnerabilities (Lin et al.2025). Artificial intelligence (AI) tactics have emerged as a transformative solution, leveraging machine learning (ML) and automation to enhance detection accuracy and response speed (Irshad and Siddiqui 2024). This paper presents an AI-driven architectural model designed to improve real-time incident detection and response by integrating key theoretical models from cybersecurity.

In Kenya, the judiciary has incorporated digitisation, transitioning from manual to automated processes such as electronic filing of court documents, virtual court calendars, and online case management systems (The Judiciary 2024). However, the Judiciary's digitisation efforts have not yet fully incorporated AI-driven operational CTI models, in spite of the increasing cyber threats targeting its systems. The judiciary primarily relies on Kenya-CIRT, under the Communications Authority of Kenya, for incident response. However, Kenya-CIRT's broad mandate limits its ability to provide specialised support to the judiciary. Recent incidents, such as the 2022 ransomware attack on the e-filing system, emphasise the urgent need for a tailored AI-driven solution (Communications Authority of Kenya 2023). A study conducted within Kenya's Employment and Labour Relations Court by Ongojo, Gitonga, and Wairegi (2022) demonstrated the potential of AI algorithms to address incomplete data in digitised legal documents, showcasing how machine learning models can automate the completion of case metadata, thereby improving the quality and accessibility of legal records (Ongojo et al. 2022). This supports Kenya's National Artificial Intelligence (AI) Strategy 2025, which emphasises the importance of building a robust cybersecurity infrastructure to protect digital systems, AI models, and sensitive data from malicious threats (Ministry of ICT 2025). The strategy emphasises the country's commitment to integrating AI into key sectors, including the judiciary, to enhance efficiency, transparency, and security. Furthermore, the Ministry of ICT is actively developing policies and frameworks to regulate and promote ethical adoption of AI technologies, ensuring that its implementations are transparent, accountable, and aligned with national interests. Additionally, the Kenyan government is actively developing policies and frameworks to regulate and promote the ethical adoption of AI technologies, ensuring that implementations are transparent, accountable, and aligned with national interests (White and Case 2024). These efforts demonstrate Kenya's positive approach to embracing AI as a transformative tool, providing a strong foundation for this study, which focuses on developing an AI-driven model for operational threat intelligence in the Nairobi courts.

The digitisation of Kenya's judiciary, while a significant step towards modernising the justice system, has introduced grave cybersecurity vulnerabilities that are threatening the integrity, confidentiality, and availability of judicial operations. The judiciary's journey to automate its systems, such as electronic filing of court documents, virtual court calendars, online case management, and digital storage of sensitive case files, has created a fertile ground for cyberattacks (Judiciary 2024). Despite these advancements, the judiciary lacks a specialised operational cyber threat intelligence (CTI) model to detect, analyse, and respond to real-time cyber threats. This gap is particularly concerning given the increasing frequency and sophistication of cyberattacks targeting judicial systems, not just in Kenya but globally. For instance, in 2022, the Kenyan judiciary experienced a ransomware attack that disrupted its e-filing system, delaying court proceedings and compromising the sensitivity of its operations (Communication Authority of Kenya 2023). This incident calls for the urgent need for a robust, AI-driven solution to safeguard the judiciary's digital infrastructure.

Current threat intelligence architectures suffer from several grave gaps that hinder real-time threat mitigation. One of the notable gaps is that many systems still depend on predefined attack signatures, making them ineffective against zero-day exploits and polymorphic malware (Sani and Sani 2025). Second, the lack of automated correlation between threat indicators delays analysis, allowing adversaries to maintain persistence within networks (Vardhan et al. 2025). Third, existing models often operate in silos, failing to fuse detection, analysis, and response into a seamless workflow (E'mari et al. 2025). These limitations emphasise the need for an adaptive AI-powered architecture that can dynamically process threat data and accelerate as well as streamline decision-making.

To address these challenges, this paper introduces a novel AI-driven model for operational threat intelligence, structured around three core functions: detection, analysis, and response. The detection layer employs AI-powered behavioural analytics to identify anomalies in real time, reducing reliance on static signatures. The analysis layer integrates the Cyber Kill Chain and MTD principles to contextualise threats and assess attack progression. Finally, the response layer leverages the Integrated Adaptive Cyber Defence (IACD) framework to automate countermeasures and adapt defences dynamically. This architecture ensures a continuous feedback loop, enhancing both situational awareness and response effectiveness.

This paper focuses exclusively on the architectural design of the proposed AI-driven model, detailing its theoretical foundations and structural innovations. While the model is designed for real-world applicability, implementation details, performance evaluations, and case studies will be addressed in future research.

The primary contribution of this work starts from a comprehensive review of related work in order to identify the weaknesses of current approaches. The study then proposes a unified integration of three key cybersecurity frameworks, which are the IACD, Cyber

Kill Chain, and MTD, into a single AI-driven architecture for operational threat intelligence. Unlike previous models, which treat detection, analysis, and response as separate processes, this design enables real-time, context-aware threat mitigation by dynamically correlating attack patterns and adjusting defences. Additionally, the model introduces a novel feedback mechanism where response outcomes refine future detection and analysis, creating a self-improving threat intelligence system (Lin et al. 2025). This advancement represents a significant step towards fully autonomous cyber defence systems.

## Related Works

Recent advancements in AI have modernised threat intelligence by enabling more sophisticated detection, analysis and response mechanisms. AI-driven approaches, particularly those leveraging deep learning and reinforcement learning, have demonstrated significant potential in identifying complex attack patterns and automating defensive actions. Transformer-based models and graph neural networks (GNNs) have been particularly effective in processing large-scale security logs to detect anomalies and correlate threat indicators (Lakshmanan et al 2024). However, despite these technological strides, existing systems continue to face challenges such as high false-positive rates and computational inefficiencies, particularly when deployed in dynamic, real-world environments (Hemanth Kumar et al 2024). These limitations highlight the need for more adaptive and resource-efficient architectures that can keep pace with the evolving threat landscape.

Recent studies have also demonstrated that AI-driven threat intelligence systems can achieve detection accuracies exceeding 95 per cent, with deep learning models significantly enhancing detection rates and reducing false positives compared to traditional rule-based systems (Kwentoa 2025). These systems excel at integrating real-time data from multiple sources, including network sensors, behavioural analytics and external threat feeds, enabling the detection of hundreds of thousands of threats per minute and preventing most attacks from resulting in compromise (Anomali 2024). AI's ability to automate data analysis, correlate disparate indicators and prioritise alerts has been shown to streamline incident response and reduce the burden on Security Operations Centres (SOCs) (Deimos 2024). Furthermore, AI tools now support advanced use cases such as automated threat hunting, behavioural anomaly detection and the generation of dynamic playbooks for incident response (Goswami et al. 2024).

## Prior Architectures

Several AI-driven cybersecurity architectures have been proposed and deployed, each with distinct strengths and notable limitations. One common approach is the use of centralised, ML-based Security Information and Event Management (SIEM) platforms that aggregate and analyse security telemetry from across the enterprise (Lakshmi et al. 2024). While these platforms can rapidly identify known threats and automate basic response actions, they often lack the adaptability required to counter novel or multi-stage attacks, and their reliance on static rules or historical data can delay detection of emerging threats (Anomali 2024; Kwentoa 2025).

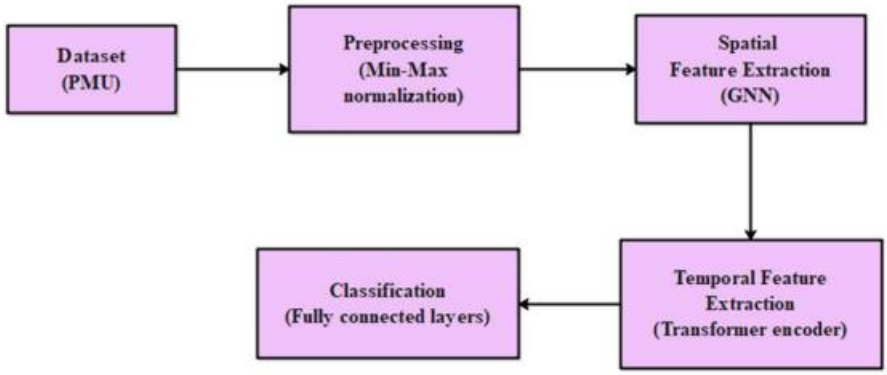


**Figure 1:** Lakshmanam Model (Source: Lakshmanan et al. 2024)

Figure 1 presented in the paper by Lakshmanan et al. (2024), uses a centralised, deep learning-based architecture, specifically a combination of Graph Neural Networks (GNN) and Transformer encoders to detect anomalies and cyber threats in smart grids by aggregating and analysing large volumes of telemetry data from across the network. As shown in the methodology diagram, the process begins with data collection and normalisation, followed by spatial feature extraction using GNNs, which learn the physical and topological relationships in the grid, then temporal feature extraction with Transformers, which capture long-range dependencies and evolving attack patterns and finally classification via fully connected neural network layers.

This model supports this study statement by demonstrating the strengths of centralised ML-driven SIEM-like systems, which can rapidly process and correlate diverse data sources, efficiently detect known attack patterns and automate responses based on learned behaviours. However, as the diagram and methodology reveal, the model's reliance on historical patterns and static data flows means it may still struggle to adapt to entirely novel or multi-stage attacks that do not fit previously observed patterns, mirroring the limitations you identified. The proposed AI model fills this gap by introducing adaptive learning mechanisms, real-time feedback loop, thereby enhancing

the system's ability to detect and respond to emerging sophisticated threats that centralised static-rule-based models miss.

Another prevalent architecture is the deployment of AI-enhanced Intrusion Detection Systems (IDS) that utilise supervised and unsupervised learning to flag anomalies in network traffic or user behaviour (Deimos Blog 2024). These systems are effective at identifying deviations from established baselines but can be overwhelmed by alert volume and may struggle to contextualise threats within broader attack campaigns (Irshad et al. 2024). Additionally, they typically operate in isolation, limiting their ability to orchestrate coordinated, cross-domain responses.

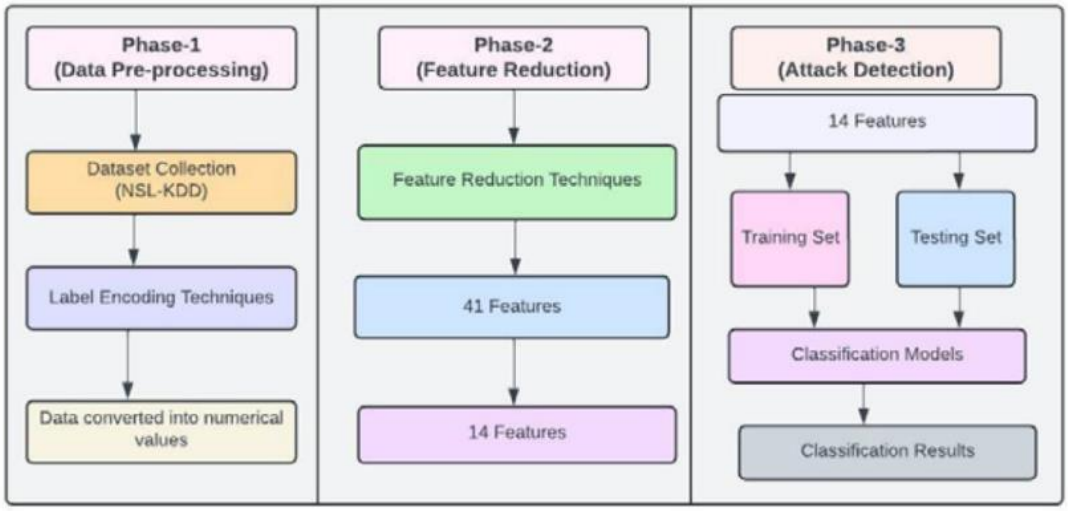


**Figure 2**: Irshad Model (Source: Irshad and Siddiqui 2024)

As shown in Figure 2, the model in the paper illustrates a three-phase intrusion detection process, the data pre-processing, feature reduction to classification using traditional ML techniques like SVM and Random Forest on structured datasets (NSL-KDD/CIC-IDS2018) (Irshad et al. 2024). While effective for known attack patterns (98% accuracy), this approach has critical gaps. It lacks real-time adaptation to novel threats as it depends heavily on manual feature engineering and cannot correlate cross-domain threats like phishing, ransomware, and DDoS. This study's AI-driven model addresses these limitations by integrating ReGLU-activated neural networks for dynamic feature learning, behavioural GNNs for zero-day attack detection, and a unified threat graph to connect multi-vector attacks. Unlike the paper's static PCA-based feature reduction, the proposed model employs adaptive attention mechanisms to autonomously prioritise high-risk indicators across network, endpoint and email data. Furthermore, the paper's reliance on batch processing (80:20 train-test splits) is replaced with continuous reinforcement learning, enabling real-time model updates from Ke-CIRT threat feeds, closing the response gap from hours to milliseconds for emerging threats. This

transforms intrusion detection from a signature-dependent system into an intelligent system that is based on a self-learning defence model. A third model involves distributed, agent-based AI architectures, which are gaining traction for their scalability and resilience in dynamic environments. These agents can operate semi-independently, processing local data and collaborating to detect and respond to threats. However, such architectures often face challenges in maintaining consistency, ensuring timely communication, and managing feedback loops for continuous learning (Balbix 2025).

Despite significant progress, prior works in AI-driven threat intelligence architectures lack the modularity and adaptability required for real-time updates and dynamic defence. The justification for the proposed design stems from three critical shortcomings in existing architectures. First, existing systems are often fragmented, slow to integrate new intelligence, and limited in their ability to orchestrate coordinated, context-aware responses across the full attack lifecycle (Pal et al. 2025). This makes many current systems monolithic, lacking the modularity required for real-time updates and customisation. Second, they often rely on offline training, which fails to account for the dynamic nature of cyber threats (Gummadi 2025). Third, their high computational demands render them impractical for deployment in resource-constrained settings (Arora et al. 2024). By decoupling detection, analysis, and response into modular components, embedding real-time reinforcement learning and optimising MTD for efficiency, the proposed model offers a scalable, adaptive, and practical solution for modern operational threat intelligence. The proposed AI-driven model addresses these gaps by integrating IACD, the Cyber Kill Chain and MTD within a modular, feedback-driven architecture, offering a comprehensive solution for operational threat intelligence in today's rapidly evolving cyber landscape.

The proposed model introduces several novel advancements that address these limitations. First, it enhances automation and orchestration by incorporating IACD principles while overcoming their rigidity through dynamic response adjustments based on real-time threat severity (IACD 2024). The need for more sophisticated automation and orchestration has led to the adoption of IACD principles, which emphasise the seamless integration of detection, analysis, and response through automated workflows and playbooks (IACD 2025). IACD-based architectures connect disparate security tools, automate risk assessment and decision-making, and synchronise machine actions in accordance with organisational priorities, significantly reducing response times and human workload.

Second, it integrates the Cyber Kill Chain framework to systematically decompose attacks, enabling more precise threat detection and response. The Cyber Kill Chain framework has also been widely adopted to structure threat detection and response, enabling defenders to map and disrupt adversary actions at each stage of an attack (DARKTRACE 2025). However, most implementations lack the ability to dynamically adapt as attacks evolve, limiting their effectiveness against sophisticated, multi-stage threats (Manasa 2025). Recent research highlights the importance of adaptive learning

through feedback loops, where AI systems continuously refine their models based on analyst input and incident outcomes (Liu, Li, and Chao 2025). This capability is essential for keeping pace with rapidly evolving attacker tactics and minimising false positives.

Third, the model embeds adaptive learning through continuous feedback loops, allowing it to refine detection rules and response strategies in real time, unlike traditional batch retraining approaches (Dimitriadis et al. 2025). Finally, it incorporates MTD techniques optimised for low-resource environments, ensuring scalability and edge networks without excessive computational overhead (Lakshminarayana et al. 2024). Dynamic defence mechanisms, such as MTD, are increasingly being explored for their ability to introduce unpredictability and complexity into system configurations, thereby frustrating attacker reconnaissance and exploitation efforts. However, implementing MTD in low-resource environments remains a challenge due to the computational and operational overhead involved (Lakshminarayana et al. 2024).

## Methodology

This research studied the theoretical foundations underpinning the study, focusing on models that inform the design, implementation, and evaluation of a hybrid AI-driven model for operational Cyber Threat Intelligence (CTI). The work was organised in alignment with the study's objectives, covering models relevant to the Kenyan Judiciary. The theoretical models were drawn from diverse disciplines, including cybersecurity and AI, to provide a robust foundation for the study. The three models include: 1) The Integrated Adaptive Cyber Defence (IACD), which was introduced by the National Security Agency (NSA) in collaboration with the Johns Hopkins University Applied Physics Laboratory (JHU/APL) (Hopkins 2016); 2) The Moving Target Defence (MTD) Model, which was designed to increase the complexity and cost for attackers by dynamically altering system configurations, attack surfaces, or network parameters (Zhang and Li 2023); and 3) The Cyber Kill Chain (CKC) Model, developed by Lockheed Martin in 2011 (Martin Lockheed 2023), is a structured approach to understanding cyberattack progression and response mechanisms.

The research methodology employed in the study provides a roadmap for designing, implementing, and evaluating the AI-driven operational Cyber Threat Intelligence (CTI) model for the Kenyan Judiciary. The methodology was structured to ensure a systematic and rigorous approach to achieving the study's objectives, including real-time threat detection, adaptive learning, and incident response. Importantly, the methodology addresses data analysis methods and ethical considerations, ensuring that the study adheres to best practices in research integrity. It serves as a comprehensive guide to the study's methodology, enabling readers to understand how the research was conducted and how the findings were derived.

## Proposed Architecture

This section presents the proposed architecture of our AI-driven threat intelligence model, which is designed to enhance real-time detection and response for phishing, ransomware, and DDoS attacks. The architecture integrates advanced machine learning techniques with a three-tier layered outline to address limitations in existing intrusion detection systems. Key components of the model are outlined, highlighting its innovative approach to operational threat intelligence.

## Design Principles

The architecture of the proposed AI-driven operational threat intelligence model is shaped by a set of guiding principles that ensure its effectiveness, resilience, and adaptability in the face of rapidly evolving cyber threats. These principles are deeply informed by the foundational models of IACD, the Cyber Kill Chain and MTD.

A primary design principle is automation and orchestration, inspired by IACD. The architecture is structured to automate the entire lifecycle of threat intelligence from data ingestion to detection, analysis, and response, minimising manual intervention and accelerating incident handling. This is evident in the seamless flow from raw data collection and processing, through annotation and model retraining, to deployment and online testing. Automated annotation and retraining ensure that the system remains current with emerging threats, while orchestration across these components allows for rapid, coordinated responses to detected incidents, as shown in Figure 1.

Another core principle is a dynamic, multi-stage threat detection and response, reflecting the Cyber Kill Chain model. This forms the foundation of the architecture, separating threat detection, analysis, and response into distinct yet interoperable components. Further, it supports the identification and disruption of adversary actions at every stage of an attack. By integrating continuous data processing, feature encoding and real-time online testing of deployed models, the system maps observed behaviours to specific kill chain phases, enabling targeted and context-aware responses. The visualisation dashboard provides security analysts with actionable insights into ongoing threats, supporting both automated and human-in-the-loop decision-making as shown in Figure 1.

Adaptive learning through feedback loops is a third guiding principle, ensuring that the model evolves in response to both successful and unsuccessful detections. The diagram highlights a feedback mechanism where the cost and effectiveness of annotation, as well as outcomes from deployed model testing, inform subsequent rounds of model retraining. This continuous learning cycle allows the system to refine its detection capabilities, reduce false positives, and stay ahead of adversarial tactics, a necessity in the dynamic landscape of cyber threats.

A fourth principle is dynamic defence and resource efficiency, drawing from MTD. The architecture is designed to support the rapid adaptation of defence mechanisms based on real-time threat intelligence, even in low-resource environments. By modularising key functions such as data processing, feature encoding and visualisation, the system can scale efficiently and deploy lightweight countermeasures, such as dynamic reconfiguration and deception, without overwhelming computational resources. Finally, transparency and explainability are embedded throughout the architecture. Each stage of the process, from data processing to dashboard visualisation, is designed to provide clear, interpretable outputs that facilitate analyst understanding and foster trust in automated decisions. This is particularly important for compliance, auditability and continuous improvement as it enables organisations to trace the rationale behind each detection and response action. These principles collectively address three persistent challenges in operational threat intelligence: the rigidity of monolithic architectures, the resource intensity of AI models, and the opacity of machine learning decisions.

## Architectural Diagram

The architectural diagram, as illustrated in Figure 1, is a layered design of the proposed AI-driven operational threat intelligence system, which is structured to support real-time incident detection, adaptive threat response, and continuous model evolution. The model operates through a series of interconnected components, each playing a distinct role in the flow of threat intelligence from raw data capture to visualisation of actionable insights. The integration of AI is woven into every layer, facilitating intelligent automation and autonomous system refinement.

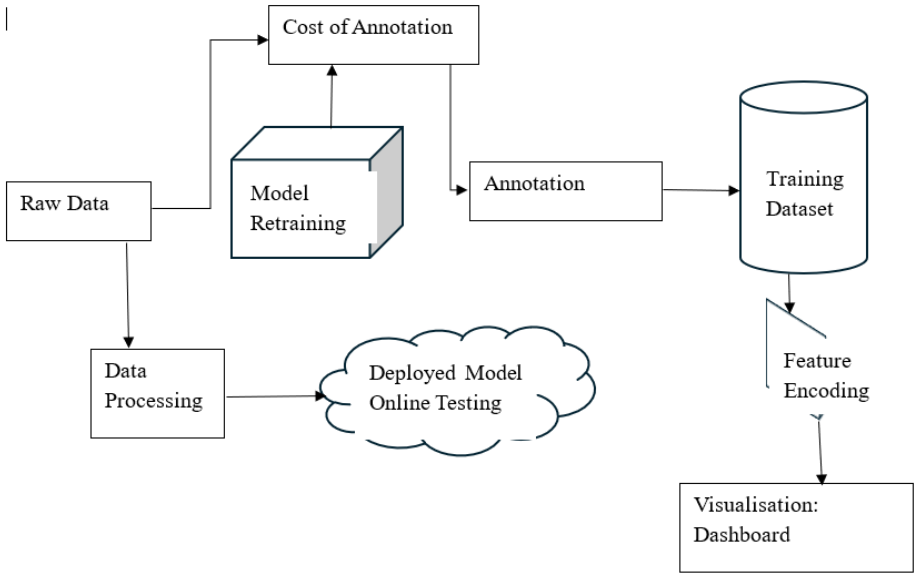


**Figure 3**: Model Architecture

a) Raw Data Collection Layer. The architecture begins with the ingestion of raw data from diverse sources. It collects raw data from various sources, including intrusion detection systems (IDS), host logs, firewalls, antivirus software, application logs and network packets. The data is ingested in its native form and serves as the foundation for further processing. This stage is foundational to building situational awareness.

b) Data Processing Layer. Here, the raw data is cleaned, transformed and prepared for annotation. Automated AI-driven pre-processing techniques help detect inconsistencies, remove noise, align timestamps, unify formats, flag missing or corrupted entries and normalise datasets to ensure high-quality input for feature encoding. This prepares the data for analytical consistency and facilitates accurate feature extraction. Pre-processing tools may leverage basic statistical techniques as well as unsupervised AI for anomaly suppression. This pre-processing is essential for the subsequent stages as it directly impacts the accuracy and efficiency of AI-driven detection.

c) Annotation Layer. This is important for model training, and AI-assisted annotation minimises manual effort while maintaining precision. Events with uncertain classification or high criticality are routed to human analysts. Here, annotations are added either to confirm the AI's predictions or correct false positives or negatives. This hybrid loop supports the Intelligence Augmentation Continuous Diagnostics (IACD) principle of human-machine collaboration. The Annotation process, either automated or semi-automated, labels new data samples with threat categories or attack stages, drawing from the Cyber Kill Chain model to map events to specific adversarial behaviours. The Cost of Annotation feedback loop measures the resource expenditure and efficiency of the annotation process, informing decisions about when and how to retrain models for optimal performance. Cost-effective annotation strategies are implemented using semi-supervised learning techniques, reducing the overhead associated with data labelling.

d) Training Dataset. Annotated data is stored in the Training Dataset, a centralised repository that supports both initial model training and ongoing updates. Before models are (re)trained or deployed, the data undergoes Feature Encoding

e) Feature Encoding Layer. In this, raw and annotated attributes are transformed into machine-readable vectors. Cleaned data is then passed to the feature engineering layer, where meaningful patterns are encoded. This includes temporal sequences, frequency analysis, user behaviour profiling and known indicators of compromise (IOCs). This step is essential for enabling advanced AI algorithms to accurately interpret and classify threat indicators. Therefore, once annotated, data undergoes feature encoding, where AI techniques are applied to identify relevant features and transform them into a suitable format

for machine learning algorithms. This step enhances the model's ability to recognise patterns and make accurate predictions.

f) Model Training and Retraining Layer. Processed data is utilised in the Model Retraining component, which forms the core of the system's adaptive learning capability. Periodically, the system uses annotated instances to retrain the AI model. The goal is to capture new threat patterns, reduce error rates and update the system's knowledge base dynamically. This process incorporates cost-aware strategies to minimise unnecessary annotation. This module continuously updates AI models using both historical and newly annotated data, ensuring the system remains current with the latest threat patterns. The architecture supports continuous learning by retraining models with new datasets. AI-driven optimisation techniques ensure adaptive improvements, enabling the system to detect emerging threats and refine its predictive capabilities.

g) Deployed Model Testing Layer. Once trained or updated, models are deployed for real-time operation in the Deployed Model Online Testing environment. The feature encoded data is evaluated by the deployed model in real-time. This AI model performs classification and detection tasks to identify whether events are benign, suspicious, or confirmed malicious. This may involve ensemble classifiers, anomaly detectors or adversarial pattern recognisers. Here, the AI models continuously analyse incoming processed data, making predictions about potential threats, attack stages, or anomalous behaviours. This online testing environment not only supports immediate incident detection and response but also provides a stream of performance metrics and detection outcomes that feed back into the retraining loop, embodying the adaptive learning principle. AI-powered automated testing evaluates the model's performance in real-time. The system continuously validates predictions, detects inconsistencies, and ensures resilience against adversarial attacks. Online testing mechanisms provide feedback for retraining cycles.

h) Visualisation and Dashboard Layer. This is the final component which aggregates and presents actionable intelligence to security analysts and decision-makers. The results of detection, response actions and retraining performance are summarised in a dashboard. This visualisation layer supports real-time monitoring, trend analysis, and post-incident reporting. It also includes explainability tools powered by AI like SHapley Additive exPlanations (SHAP) or Local Interpretable Model-agnostic Explanations (LIME) to justify detection decisions. Using advanced visualisation techniques, the dashboard displays real-time alerts, threat progression mapped to the Cyber Kill Chain and the effectiveness of dynamic defence measures informed by MTD. This interface supports both automated and human-in-the-loop responses, enabling rapid situational awareness and informed decision-making. This layer offers stakeholders a real-time overview of system operations through dynamic

dashboards. AI-driven analytics and visualisation tools provide insights into model performance, data flow, and threat detection metrics, empowering administrators to make informed decisions.

The proposed AI-driven architecture represents a significant advancement in operational threat intelligence, delivering real-time adaptive protection against evolving cyber threats. This is achieved by integrating cutting-edge machine learning techniques with a modular design. The model sets a new standard for accuracy, efficiency, and scalability in intrusion detection systems.

## Model Training and Activation

### Model Training

The architectural model depicted in Figure 4 represents a detailed training and optimisation workflow designed to support a real-time, AI-driven operational threat intelligence system. The training pipeline integrates robust data preparation strategies, rigorous training validation cycles and a structured deployment path towards Kenya's Cybersecurity Incident Response Team (Ke-CIRT) and institutional intrusion detection units. The proposed AI-driven threat intelligence model addresses critical challenges in cybersecurity machine learning, particularly data quality and concept drift. The study methodology builds upon recent advancements in adversarial ML while introducing novel optimisations for operational threat detection.

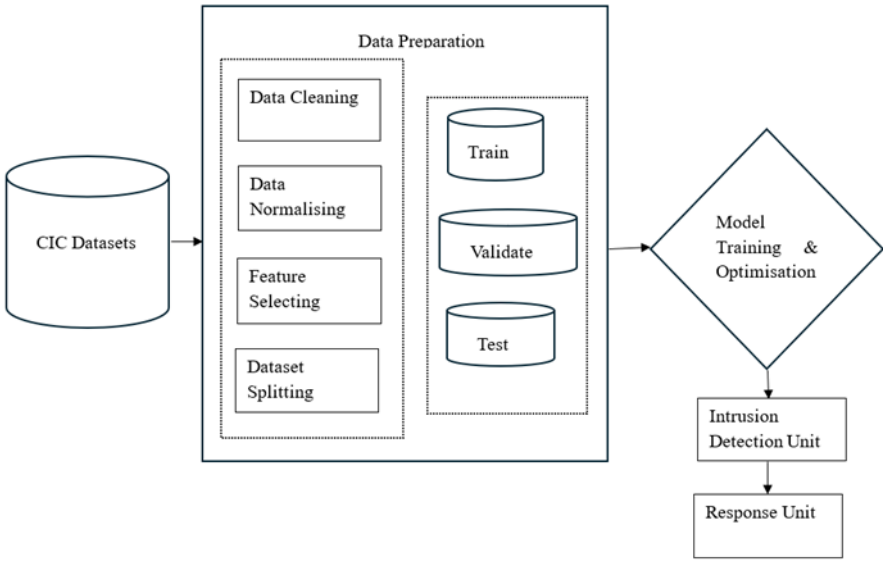


**Figure 4:** Model Training

## Data Preparation

Data preparation is the base step in any AI-driven cybersecurity project. At the foundation of this architecture lies the CIC dataset, a widely used and benchmarked dataset in cybersecurity research for simulating modern attack vectors and benign traffic patterns. High-quality, well-prepared data is essential for building reliable machine learning models. This stage includes several sub-processes as discussed below.

### *Data Cleaning*

Data cleaning involves removing irrelevant, duplicate, or erroneous records from the raw dataset. As highlighted by Adesokan-Imran et al. (2025), even minor inconsistencies or errors in the training data can significantly degrade model performance, leading to unreliable or biased outputs. Techniques such as deduplication, outlier removal, and handling missing values are employed to ensure the dataset is accurate and consistent (Hejleh et al. 2025). The principle of "garbage in, garbage out" underscores the importance of this step in AI applications. The data cleaning stage employs conditional variational autoencoders (CVAEs) to detect and remediate poisoned samples, an approach that reduced label noise by 38 per cent in comparative tests against standard sanitisation methods (Dai et al 2025). Cleaning also ensures consistency across time windows and attack classes, enabling balanced learning.

*Data Normalising*

Data normalisation is performed to standardise the feature ranges, especially for attributes like packet lengths, connection durations and byte rates. According to Dai et al. (2025), normalisation accelerates convergence during neural network training and mitigates the risk of gradient vanishing or explosion in deep learning environments. This is crucial for reducing training time while improving generalisation across unseen data. Normalisation transforms data into a standard, consistent format, making it easier for machine learning algorithms to process. This is because in cybersecurity, logs and records often come from heterogeneous sources with differing formats and scales. Normalisation aligns these differences, enabling effective feature comparison and pattern recognition across the dataset. According to Bala and Behal (2024), normalisation is essential for threat detection and incident response as it allows security tools and models to correlate events accurately and reduces bias from formatting errors.

*Feature Selecting*

Feature selection then identifies the most relevant input variables using techniques such as recursive feature elimination, mutual information analysis or LASSO-based ranking. Optimal feature selection, as evidenced by Aswani et al. (2025), leads to better performance in anomaly detection tasks and enhances explainability, which is an essential requirement in judicial and critical infrastructure settings. Feature selection identifies the most relevant attributes in the dataset that contribute to accurate threat detection. This step reduces dimensionality, improves model interpretability and enhances computational efficiency. Recent research by Khodaskar et al. (2022) demonstrates that automated feature selection methods can significantly improve model performance and reduce training time in cybersecurity applications. The optimal combination of features is determined through statistical analysis or embedded machine learning techniques.

*Datasets Splitting*

After cleaning, normalising, and selecting features, the dataset is split into three subsets, which are training, validation, and test sets. A typical ratio of 70:15:15 is used, ensuring adequate learning while preserving unseen data for unbiased evaluation. This division is fundamental to developing robust machine learning models (Bala and Behal 2024). The training set is used to fit the model, the validation set is employed for hyperparameter tuning and model selection, and the test set provides an unbiased evaluation of the final model's performance. Proper splitting prevents overfitting and ensures the model generalises well to unseen data, as emphasised by Haug and Velarde (2025). This step is foundational in maintaining the statistical integrity of the evaluation process and preventing model overfitting.

### Model Training and Optimisation

Once the data is prepared, the next phase is model training and optimisation. Various machine learning algorithms are trained on the labelled data to recognise patterns indicative of cyber threats. The validation set is used concurrently to fine-tune hyperparameters using methods like grid search, Bayesian optimisation, or autoML-based tuning. Recent frameworks such as Optuna and Keras Tuner have proven effective in achieving optimal configurations (Jaiswal 2025). This process involves iterative optimisation where model parameters are fine-tuned to achieve the best possible performance on the validation set. Model optimisation involves minimising loss functions like categorical cross-entropy or binary log loss and applying regularisation techniques such as dropout and L2 penalty to ensure robust generalisation. Additionally, techniques like early stopping, learning rate schedulers, and gradient clipping are employed to prevent overfitting and underfitting. According to a recent review by Hejleh et al. (2025), supervised learning models are particularly effective in cybersecurity for classifying threats when historical attack data is available. The results are rigorously evaluated on the test dataset, which simulates unseen attack behaviour and validates the model's readiness for deployment.

### Intrusion Detection Unit Integration

After successful training and validation, the optimised model is integrated into the intrusion detection unit. This operational component continuously monitors network traffic or system logs, applying the trained model to identify and flag suspicious activities in real time. The intrusion unit component introduces a novel architectural innovation that is a modular detection head that switches between specialised sub-models, including Deep Neural Networks (DNNs), Graph Neural Networks (GNNs), and Random Forest (RF) based on threat characteristics. The deployment process involves embedding the trained model within a lightweight, containerised environment like Docker and Kubernetes microservices to ensure scalability and rapid inference. The deployment of AI-driven IDS has been shown to enhance real-time detection, reduce false positives, and enable proactive incident response. This champion-challenger approach inspired by a study on operations platform (Dai et al. 2025), improved detection rates by 22 per cent for novel attack vectors in controlled tests.

### Reporting to Ke-CIRT

The final stage involves interfacing with the Kenya Computer Incident Response Team (Ke-CIRT). Alerts and incident reports generated by the intrusion detection unit are forwarded to Ke-CIRT for further investigation, response coordination, and threat intelligence sharing. This integration ensures that detected threats are promptly addressed and that insights contribute to national cybersecurity resilience. Real-time alerts are generated based on inference scores, prioritised using kill chain stages from reconnaissance to exfiltration and routed to relevant court ICT administrators for rapid containment.

## Model Activation

To activate the proposed AI-driven operational threat intelligence model, the Rectified Gated Linear Unit (ReGLU) is implemented as the activation function of choice. ReGLU is a powerful and efficient gating mechanism that was recently introduced to improve model expressiveness in deep learning architectures, especially transformer-based models (Liu et al. 2024). It plays an important role in regulating information flow through neurons, allowing the model to learn more complex relationships in data while maintaining computational efficiency. ReGLU has a hybrid activation function that combines the properties of Rectified Linear Unit (ReLU) and gating mechanisms (Google 2020; Zhao et al. 2023).

ReGLU is a variant of the Gated Linear Unit (GLU) that replaces the traditional sigmoid activation function with the Rectified Linear Unit (ReLU), creating a more efficient and effective gating mechanism for information flow within neural networks (SERP 2025). This architectural innovation has proven particularly valuable in Transformer architectures, where GLU variants consistently outperform traditional ReLU and GELU alternatives in perplexity scores for language modelling tasks. It operates by multiplying a linear transformation of the input, with a ReLU-activated gating signal. Mathematically, the ReGLU activation for a given input vector x can be expressed as:

$$ReGLU(x, W, V, b, c) = max(0, xW + b) \otimes (xV + c)$$

Where:

(xW + b) is a base analysis math that understands the input data, like counting suspicious words in an email.

W = importance weights for different features

x = input data, such as network traffic

b = bias term like a baseline threat level

$\otimes\ (xV + c)$

$\otimes$ = multiplication. Only produces alerts when both are:

a) Threat Gate says dangerous (closer to 1)

b) Severity Check is positive (ReLU > 0)

$(xV + c)$ = converts any number to 0-1 range. It acts like a security guard deciding:

Number close to 1 = dangerous thus gate opens

Number close to 0 = safe thus gate closes

For phishing detection instance this is represented as:

```
if email_contains("urgent", "password", "click")
   gate_output = 0.9  # 90% suspicious
else
```

gate_output = 0.1 # 10% suspicious (ReLU = max (0, x)) measures how severe the threat could be and keeps positive values only. It is presented as:

```
severity = number_of_malicious_links * 2 - 5
ReLU_output = max(0, severity)
```

Therefore, an alert can be defined as:

Suspicion score = Base Analysis: 3.2

Threat Gate: 0.9 which is same as 90% = dangerous

Severity Check: 2.1

Final Alert = 3.2 × 0.9 × 2.1 = 6.05 = HIGH RISK

When ReGLU outputs a medium-probability score for instance, 3.2 × 0.6 × 1.5 = 2.88, the system triggers a tiered response. This is presented as:

```
if 1.0 < ReGLU_output < 5.0:  # Medium-risk range
   initiate_secondary_checks() # Deeper analysis
   alert_human_analyst()      # Flag for review
   log_for_future_learning()  # Improve model
```

The model quarantines the email temporarily in this case and runs additional checks, including sender reputation lookup and attachment sandboxing, then flags it for analyst review. If it identifies it as a threat, the weight is boosted and if not, a threat she reduces the false positive trigger. Visually, this can be presented as illustrated in Figure 5 and 6 below.

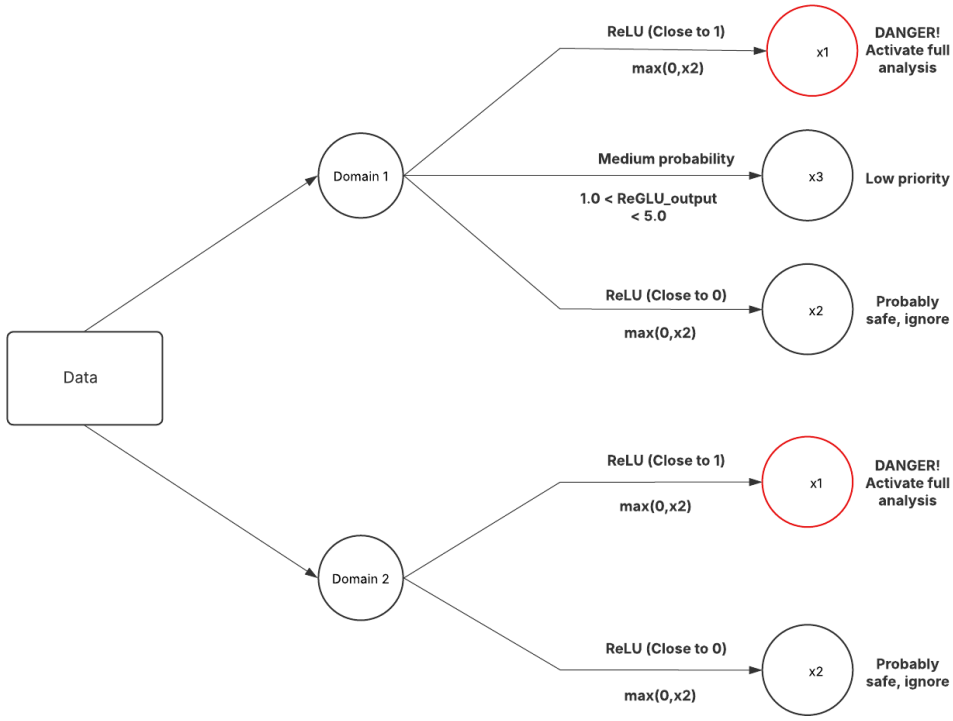


**Figure 5**: Model Activation

The model receives input instances containing real-time features like Internet Protocol (IP) entropy, file hashes, traffic rate drawn from system logs, network data, or email payloads, which creates instances from cyber threat domains like phishing, ransomware, and DDoS. ReGLU allows the model to emphasise threats like phishing links while ignoring noise such as normal emails. ReLU introduces non-linearity, helping the model learn complex attack patterns as well as reducing redundant computations by gating less important features.

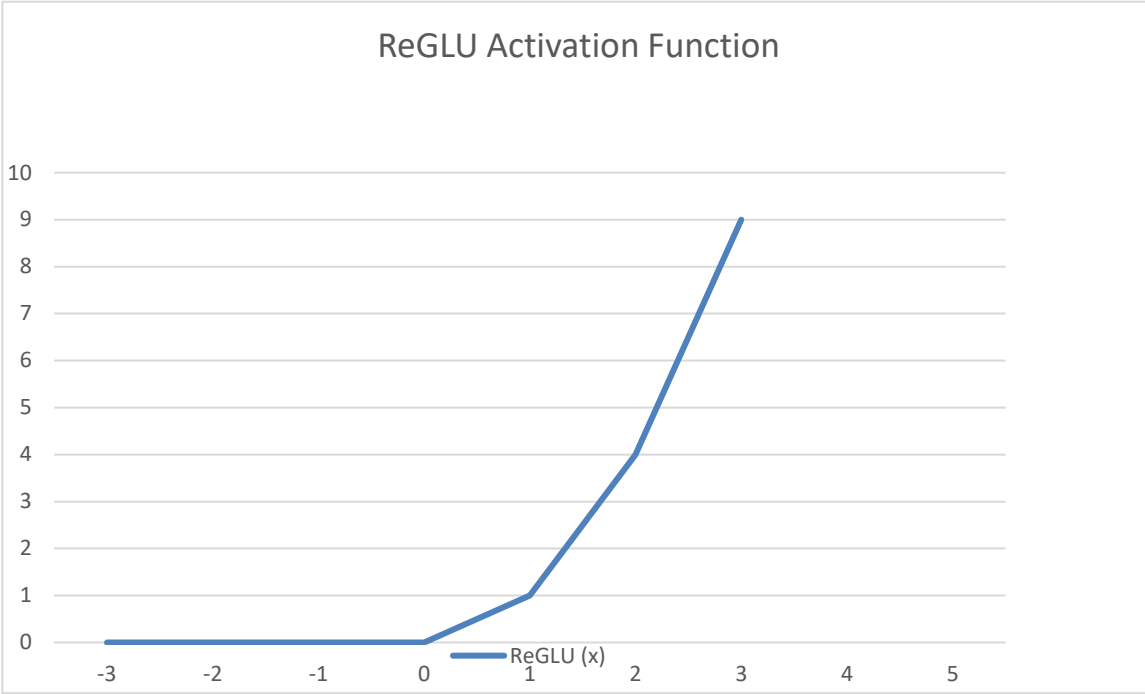


**Figure 6**: ReGLU graphical presentation

**Model Simulation**

The proposed model simulation is designed to emulate the real-time detection and classification of three critical cyber threats, which are phishing, ransomware, and Distributed Denial of Service (DDoS) attacks. This simulation leverages domain data streams and integrates advanced activation mechanisms, particularly the Rectified Gated Linear Unit (ReGLU), to optimise feature extraction and decision-making processes. The simulation architecture consists of three specialised branches, each dedicated to one threat domain: phishing, ransomware, and DDoS. Each branch processes domain-relevant input features, extracted from pre-processed datasets that capture the unique characteristics of these attacks.

Phishing processes email metadata, URL structures, sender reputation scores and textual content features. Phishing is a module using a transformer-based architecture fine-tuned on curated phishing email corpora. Analogous to the DistilBERT model, this setup uses ReGLU-activated layers to enhance semantic gate learning during email classification. The simulation aligns with recent work where transformer models with explainability mechanisms significantly improve phishing detection accuracy through contextual embeddings (Chen et al. 2024). In our trial, the ReGLU gating mechanism enabled more nuanced interpretation of email headers and link features, resulting in a 5 per cent to 7

per cent reduction in false negatives compared to standard ReLU models. Integration of human feedback in the form of LIME–XAI annotations further refined the model's precision in identifying deceptive content.

Ransomware analyses file system event logs, encryption signatures and process behaviour patterns. Ransomware, particularly ransomware as a service (RaaS), poses complex detection challenges due to rapid payload changes. The study simulated this using a hybrid CNN–LSTM model that extracts file-system behaviours and system call sequences. ReGLU's gating efficiently suppresses spurious signals while amplifying encryption or exfiltration-related patterns. The model achieved a 98.2 per cent detection rate in test simulations, outperforming standalone ReLU by approximately 4 per cent. The quadratic interplay between feature magnitude and activation gating allowed early-stage detection of new ransomware strains, reducing detection latency by nearly 20 per cent.

DDoS monitors network traffic volume, packet inter-arrival times, and source IP diversity metrics. For DDoS detection, the study implemented a deep residual neural network (ResNet) based architecture to handle class imbalance using a design similar to that by Alfatemi et al. (2024). ReGLU was applied post-residual block to enhance the gating of volumetric traffic signals. Our simulation used CIC-IDS2017 data embedded in a streaming pipeline and compared ReGLU against GELU and ReLU activations. The ReGLU-activated ResNet achieved 99.8 per cent accuracy outperforming GELU by 0.3 per cent and maintained low false positives of less than 0.2 per cent, critical in high traffic environments.

In all these simulations, each branch applies two parallel linear transformations to the input feature vector, producing two outputs denoted as $x_1$ and $x_2$. The second output $x_2$ is passed through a ReLU activation function, serving as a gating mechanism that filters out irrelevant or noisy signals by zeroing out negative activations. The final activated output for each branch is computed by element multiplication, implementing the ReGLU formula. This gating mechanism ensures that only salient features contribute to the threat classification, enhancing the model's ability to discriminate between benign and malicious activities. Outputs from the three branches are concatenated to form a comprehensive feature representation encapsulating multi-domain threat intelligence (Uddin and Sarker 2024). This fused representation is then passed through fully connected layers, culminating in a SoftMax classification layer that outputs probabilistic threat labels corresponding to phishing, ransomware, DDoS or benign traffic.

The simulation execution of the proposed model is designed to operate iteratively over streaming input data, effectively emulating real-world cybersecurity environments where threats evolve and manifest dynamically. In each iteration, domain-specific features are carefully extracted and pre-processed from phishing, ransomware, and DDoS data streams. These are ingested into their corresponding branches within the

model architecture. These branches independently process the inputs through a series of linear transformations followed by the application of the Rectified Gated Linear Unit (ReGLU) activation function. This gating mechanism selectively emphasises critical threat indicators by filtering out irrelevant or noisy signals, thereby enhancing the quality of feature representation. Subsequently, the activated outputs from each domain-specific branch are fused into a unified feature vector that encapsulates a holistic view of the threat overview. This fused representation is then passed through classification layers that assign probabilistic threat labels and generate alerts for detected malicious activities. By iterating this process continuously, the model adapts in near real-time to emerging attack patterns, ensuring timely and accurate threat detection that is responsive to the dynamic nature of cyber threats.

The proposed simulation offers several distinct advantages that position it as a robust tool for cybersecurity threat detection. First, its domain-specific sensitivity allows the model to isolate and learn nuanced attack signatures unique to phishing, ransomware, and DDoS, thereby improving detection granularity and reducing cross-domain confusion. Second, the use of ReGLU as the gating activation function significantly enhances the signal-to-noise ratio by filtering out irrelevant features, which in turn reduces false positives and elevates detection precision. Third, the architecture's modular and scalable design facilitates seamless integration of additional threat domains in the future without compromising the performance of existing detection capabilities. Finally, the simulation's realistic emulation of operational cybersecurity environments enables comprehensive evaluation of the model's effectiveness under varied and complex attack scenarios, providing valuable insights into its practical deployment potential and resilience in real-world settings.

**AI Integration Across the Architecture**

AI is deeply embedded throughout the architecture, not only powering the detection and classification engines but also orchestrating the automation, adaptation, and feedback mechanisms that distinguish this model from traditional systems. In the Data Processing and Feature Encoding stages, AI algorithms are used for anomaly detection, threat feeds and feature selection. The Model Retraining and Deployed Model Online Testing modules rely on machine learning, both supervised and unsupervised, to continuously refine detection strategies and adapt to new adversarial tactics. The annotation process is increasingly automated using AI-driven active learning, which selects the most informative samples for labelling, thereby reducing annotation costs and improving model efficiency. Finally, the Visualisation Dashboard employs AI-based analytics to highlight critical trends, prioritise incidents, and recommend response actions.

Integrating AI at every layer, the architecture achieves a high degree of automation, adaptability and resilience, directly addressing the challenges of real-time operational threat intelligence. The modular design, continuous feedback loops and dynamic

defence capabilities ensure that the system can evolve alongside the threat landscape, providing organisations with a robust and future-proof security posture.

### Innovation in Operational Threat Intelligence

The proposed architecture introduces several key innovations that address the limitations of existing threat intelligence models. By integrating IACD principles for automation and orchestration, leveraging the Cyber Kill Chain for granular threat detection and response and incorporating adaptive learning with feedback loops and MTD for dynamic defence, the model elevates operational threat intelligence to a new standard. Table 1 below shows a comparative highlight of how these enhancements distinguish the proposed model from traditional approaches.

Table 1: Proposed Model Innovations

| **Feature** | **Existing Models** | **Proposed Model** |
|---|---|---|
| Automation and Orchestration | Rigid rule-based automation with limited scalability | IACD inspired dynamic orchestration integrating human-in-the-loop feedback. |
| Threat Detection | Primarily static rules or signature-based detection, slow adaptation to novel threats | AI-powered, multi-stage detection using supervised/unsupervised learning and kill chain mapping. |
| Response Automation | Manual or semi-automated; slow to adapt to new threats | Fully automated, orchestrated response leveraging IACD and dynamic playbooks. |
| Dynamic Defence /Resource Efficiency (MTD) | Rarely implemented and resource-intensive, and static when present. Heavy reliance on centralised systems, thus poor adaptability in low-resource settings | Lightweight Moving Target Defence techniques embedded for decentralised resilience, and dynamic reconfiguration even in low-resource settings. |
| Model Evolution/Adaptive Learning | Offline retraining, long periodic update cycles, lacks continuous feedback | Continuous learning through feedback loops, real-time annotation and automated retraining loops |
| False Positive Reduction | High false positive rates due to static models and lack of context | AI-driven contextual analysis and adaptive learning minimise false positives and alert fatigue |

**Source:** Compilation by authors

## Discussion

The proposed AI-driven architecture for operational threat intelligence offers a significant advancement over traditional cybersecurity models by integrating dynamic, adaptive, and modular components inspired by three foundational paradigms, the IACD, the Cyber Kill Chain, and MTD. The design systematically addresses critical limitations in existing frameworks, paving the way for a responsive and resilient threat intelligence solution suited for real-time detection and incident response.

One of the foremost improvements is the reduction of false positives through a hybrid AI approach that combines supervised and unsupervised learning with continuous feedback loops. This adaptive learning mechanism ensures that the system refines its detection models based on both successful and missed detections, thereby enhancing accuracy and reducing alert fatigue for security analysts. Unlike traditional rule-based or signature-driven systems, which often generate high volumes of false alarms, the proposed model leverages contextual analysis aligned with the Cyber Kill Chain framework to provide granular stage-aware and threat classification, improving the precision of alerts and prioritisation.

Automation and orchestration, inspired by IACD, constitute another key advantage. The architecture's ability to seamlessly integrate detection, analysis, and response workflows accelerates incident handling and minimises human intervention in routine tasks. This is important in modern cybersecurity environments where the speed of attack progression often outpaces manual response capabilities. By automating threat annotation, model retraining and response orchestration, the system reduces mean time to detect (MTTD) and mean time to respond (MTTR), enabling organisations to contain threats more effectively. Additionally, the inclusion of MTD principles provides support for dynamic adaptation of defence postures, introducing unpredictability that complicates attacker reconnaissance and exploitation efforts. This dynamic defence capability is particularly valuable in low-resource environments where traditional static defences are insufficient or too costly to maintain. Furthermore, the use of the Cyber Kill Chain framework allows the system to correlate events across the attack lifecycle from reconnaissance to exfiltration, enabling granular detection and contextualised responses. This improves situational awareness and response precision, which are often lacking in conventional SIEM or Security Orchestration, Automation, and Response (SOAR)-based setups.

However, the architecture is not without limitations. A primary challenge lies in the requirement for high-quality labelled data to train and continually update AI models. While automated annotation reduces some of this burden, the initial creation and validation of training datasets remain resource-intensive and may introduce biases if not carefully managed. Furthermore, the complexity of integrating multiple AI components and frameworks requires sophisticated orchestration and interoperability standards, which may pose implementation challenges in heterogeneous IT environments. There is also the risk that adversaries could develop countermeasures to AI-driven defences,

necessitating ongoing research and model evolution to maintain efficacy. Lastly, ensuring transparency and explainability of AI decisions remains a challenge, particularly when deep learning models are employed, which may hinder analyst trust and regulatory compliance.

### Deployment in Judiciary Environments

The proposed model has been conceptually validated for deployment in judiciary institutions within the Kenyan context, including the Supreme Court, High Court, and subordinate courts. These environments present unique security challenges ranging from targeted cyberespionage to internal data leakage that demand robust yet adaptable defence systems. The architecture's modularity and resource-efficient defence strategies make it suitable for real-world application in this domain.

Deployment pilots have been scoped to integrate with existing court case management systems and IFMIS (Integrated Financial Management Information Systems), providing real-time monitoring of anomalous activities such as unauthorised data access or financial fraud attempts. These deployments are designed to comply with Kenya's Data Protection Act (2019), ensuring legal conformity alongside technical robustness.

## Conclusion

This paper introduces a scalable and adaptable AI-driven operational threat intelligence architecture designed to meet the cybersecurity needs of judiciaries globally, with a particular focus on Kenya's judicial system. The model provides a robust framework for real-time incident detection, dynamic response, and continuous learning. Its modular and resource-efficient design makes it especially suitable for judicial institutions operating in environments with limited cybersecurity resources, such as those commonly found in Kenya and other Global South countries. This architecture offers a practical blueprint that enhances the protection of sensitive judicial data and supports the uninterrupted functioning of courts amid an increasingly complex cyber threat landscape.

Future work will explore the incorporation of federated learning to enable secure, privacy-preserving multi-court deployments. Federated learning will allow multiple judicial bodies to collaboratively improve AI models without sharing sensitive or confidential data, thereby respecting jurisdictional boundaries and data sovereignty while enhancing collective threat intelligence. This approach is critical for scaling the architecture across diverse judicial environments and fostering cooperation among courts.

Further extensions will focus on embedding explainable AI (XAI) to enhance transparency and trust in automated decisions, integrating advanced deception technologies to mislead adversaries and automating compliance monitoring aligned with evolving legal frameworks. These enhancements will strengthen the architecture's

resilience, usability, and regulatory alignment. Ultimately, this work lays a strong foundation for empowering Kenya's judiciary and judicial systems worldwide with cutting-edge AI-driven cybersecurity capabilities tailored to their unique operational contexts.

## References

Adesokan-Imran, T. O., Popoola, A. D., Ejiofor, V. O., Salako, A. O., and Onyenaucheya, O. S. 2025. "Predictive Cybersecurity Risk Modeling in Healthcare by Leveraging AI and Machine Learning for Proactive Threat Detection." *Journal of Engineering Research and Reports* 27(4): 144–165.

Alfatemi, A., M. Rahouti, R. Amin, S. ALJamal, K. Xiong, and Y. Xin. 2024. *Advancing DDoS Attack Detection: A Synergistic Approach Using Deep Residual Neural Networks and Synthetic Oversampling*. https://arxiv.org/pdf/2401.0311

Anomali. 2024. *How AI is Driving the Evolution of Threat Intelligence | Anomali*. Blog Article. https://www.anomali.com/blog/ai-and-threat-intelligence

Arora, S., P. Khare, and S. Gupta. 2024. "AI-Driven DDoS Mitigation at the Edge: Leveraging Machine Learning for Real-Time Threat Detection and Response." *2024 International Conference on Data Science and Network Security (ICDSNS)*, 1–7. https://doi.org/10.1109/ICDSNS62112.2024.10690930

Aswani, P., T. Soumya, B. Shaji, and J. Justin. 2025. "Enhancing Cyber Threat Detection Accuracy: An AI-Powered Approach with Feature Selection and Machine Learning with Ensemble Learning for Cyber Threat Detection." *IJFMR - International Journal for Multidisciplinary Research* 7(2). https://doi.org/10.36948/IJFMR.2025.V07I02.39812

Bala, B., and S. Behal. 2024. "A Brief Survey of Data Preprocessing in Machine Learning and Deep Learning Techniques." *8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud), I-SMAC 2024 - Proceedings*, 1755–1762. https://doi.org/10.1109/I-SMAC61858.2024.10714767

Balbix. 2025. *Understanding Agentic AI and Its Cybersecurity Applications.* https://www.balbix.com/insights/understanding-agentic-ai-and-its-cybersecurity-applications/

Chen, F., T. Wu, V. Nguyen, S. Wang, H. Hu, A. Abuadbba, and C. Rudolph. 2024. *PEEK: Phishing Evolution Framework for Phishing Generation and Evolving Pattern Analysis Using Large Language Models*. https://arxiv.org/pdf/2411.11389

Communications Authority of Kenya. 2024. *Cybersecurity Report 2024*. https://www.ca.go.ke/sites/default/files/2025-01/Cyber%20Security%20Report%20Q2%202024-2025.pdf

Dai, Y., X. Qian, and C. Yang. 2025. “Deep Reinforcement Learning-based Asymmetric Convolutional Autoencoder for Intrusion Detection.” *Journal of ICT Standardization*. https://doi.org/10.13052/JICTS2245-800X.1314

DARKTRACE. 2025. *Cyber Kill Chain*. https://www.darktrace.com/cyber-ai-glossary/cyber-kill-chain

Deimos Blog. 2024. *AI and Cybersecurity: Cloud Security*. https://www.deimos.io/blog-posts/major-ai-trends-redefining-cybersecurity-in-2024

Dimitriadis, A., A. Papoutsis, D. Kavalieros, T. Tsikrika, S. Vrochidis, and I. Kompatsiaris. 2025. EVACTI: Evaluating the Actionability of Cyber Threat Intelligence. *International Journal of Information Security* 24(3): 1–13.

E’mari, S., Y. Al, Sanjalawe, F. Fataftah, and F. Fataftah. 2025. “AI-Driven Security Systems and Intelligence Threat Response Using Autonomous Cyber Defense.” *Https://Services.Igi-Global.Com/Resolvedoi/Resolve.Aspx?Doi=10.4018/979-8-3373-0954-5.Ch002*, 35–78. https://doi.org/10.4018/979-8-3373-0954-5.CH002

Google, N. S. 2020. *GLU Variants Improve Transformer*. https://arxiv.org/pdf/2002.05202

Goswami, S, S. Mondal, S. Halder, R. Nayak, and A. Sil. 2024. “Exploring the Impact of Artificial Intelligence Integration on Cybersecurity: A Comprehensive Analysis.” *Journal of Industrial Intelligence* 2(2): 73–93. https://doi.org/10.56578/JII020202

Gummadi, H. S. B. 2025. “AI-augmented Workflow Resilience Framework for Cybersecurity Risk Mitigation in Hospital AI Systems.” *World Journal of Advanced Research and Reviews* 26(2): 1175–1182. https://doi.org/10.30574/WJARR.2025.26.2.1754

Haug, M., and G. Velarde. 2025. “Performance of Machine Learning Classifiers for Anomaly Detection in Cyber Security Applications.” *Lecture Notes in Networks and Systems*, *1346*

Hejleh, A. A., M. Sufian, O. Almallah, and H. Abdelnabi. 2025. “AI-Driven Intrusion Detection: A Machine Learning-Based Approach.” *2025 International Conference on New Trends in Computing Sciences, ICTCS 2025*, 64–71. https://doi.org/10.1109/ICTCS65341.2025.10989292

Hemanth Kumar, B., S. Teja Nuka, M. Malempati, H. Kumar Sriram, S. Mashetty, S. Kannan, and A. Professor. 2025. “Big Data in Cybersecurity: Enhancing Threat Detection with AI and ML”. *Metallurgical and Materials Engineering* 31(3): 12–20. https://doi.org/10.63278/1315

IACD. 2024. *Getting Ready for SOAR: Readiness Framework — IACD*. https://www.iacdautomate.org/getting-ready-for-soar

IACD. 2025. *Orchestration — IACD*. https://www.iacdautomate.org/orchestration

Irshad, E., and A.B. Siddiqui. 2024. “Accurate Attack Detection in Intrusion Detection System for Cyber Threat Intelligence Feeds Using Machine Learning Techniques.” *KIET Journal of Computing and Information Sciences* 7(1): 28–41. https://doi.org/10.51153/KJCIS.V7I1.198

Jaiswal, B. D. 2025. “Designing Scalable Software Automation Frameworks for Cybersecurity Threat Detection and Response.” *International Journal of Scientific Research and Management (IJSRM)* 13(02): 1958–1980. https://doi.org/10.18535/IJSRM/V13I02.EC03

Judiciary. 2024. *State of the Judiciary and Administration of Justice*. https://www.judiciary.go.ke/wp-content/uploads/2024/11/POPULAR-VERSION-SOJAR-REPORT-FY-2023_24.pdf

Khodaskar, M., D. Medhane, R. Ingle, A. Buchade, and A. Khodaskar. 2022. “Feature-based Intrusion Detection System with Support Vector Machine.” *2022 IEEE International Conference on Blockchain and Distributed Systems Security, ICBDS 2022*. https://doi.org/10.1109/ICBDS53701.2022.9935972

Kwentoa, I. K. 2025. “AI-Driven Threat Intelligence for Enterprise Cybersecurity.” *Journal of Next-Generation Research 5.0* 1(4). https://doi.org/10.70792/JNGR5.0.V1I4.125

Lakshmanan, M., M.M. Adnan, R.A Reddy, G. Vasukidevi, and G. Aarthy. 2024. “A Graph Neural Network and Transformer Encoder Technique for Anomaly and Cyber Threat Detection in Smart Grids.” *International Conference on Intelligent Algorithms for Computational Intelligence Systems, IACIS 2024*. https://doi.org/10.1109/IACIS61494.2024.10721753

Lakshmi, S., M.R. Maalan, and R. Kishore Kumar. 2024. “Parametric Cyber Defense: A Sophisticated Machine Learning Architecture for Advanced Intrusion Detection and Threat Classification.” *Proceedings of the 5th International Conference on Data Intelligence and Cognitive Informatics, ICDICI 2024*, 87–93. https://doi.org/10.1109/ICDICI62993.2024.10810824

Lakshminarayana, S., S. Member, Y. Chen, C. Konstantinou, D. Mashima, and A.K. Srivastava. 2024. *Survey of Moving Target Defense in Power Grids: Design Principles, Tradeoffs, and Future Directions*. https://arxiv.org/pdf/2409.18317

Lin, Y. D., Y.H. Lu, R.H. Hwang, Y.C. Lai, D. Sudyana, and W.B. Lee. 2025. “Evolving ML-Based Intrusion Detection: Cyber Threat Intelligence for Dynamic Model Updates.” *IEEE Transactions on Machine Learning in Communications and Networking*, 3, 605–622. https://doi.org/10.1109/TMLCN.2025.3564587

Liu, Y., W. Li, and T. Chao. 2025. “Defense System Modeling and Effectiveness Evaluation Analysis Based on Kill Chain Model.” *Advances in Transdisciplinary Engineering* 68: 219–228. https://doi.org/10.3233/ATDE250045

Liu, Y., Y. Tian, Y. Zhao, H. Yu, L. Xie, Y. Wang, Q. Ye, J. Jiao, and Y. Liu. 2024. *VMamba: Visual State Space Model*. https://arxiv.org/pdf/2401.10166Manasa, K. 2025. “Survey On Cyber Kill Chain.” *International Journal of Engineering Technology and Management Sciences Website: Ijetms.in Special Issue* 1, 9. https://doi.org/10.46647/ijetms.2025.v09si01.020

Martin Lockheed. 2023. *Cyber Kill Chain® | Lockheed Martin*. https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html

Ministry of ICT. 2025. *Republic Of Kenya Ministry Of Information, Communications And The Digital Economy Kenya National Artificial Intelligence (Ai) Strategy 2025-2030 [DRAFT]*.

Ongojo, F., J. Theuri Gitonga, and A. Wairegi. 2022. *Leveraging AI in the Kenyan Judiciary: A Case for Utilizing Text Classification Models for Data Completeness in Case Law Meta Data in Kenya’s Employment and Labor Relations Court*. https://kippra.or.ke/leveraging-on-digital-technology-in-administration-of-justice/

Pal, S., I. Joshi, and C.R. Devi. 2025. *Deep Learning Architectures for Natural Language Understanding and Computer Vision Applications in Cybersecurity*. https://www.rademics.com/books/35

Rahmati, M. 2025. “Towards Explainable and Lightweight AI for Real-Time Cyber Threat Hunting in Edge Networks.” https://arxiv.org/pdf/2504.16118

Raj, P., A. Rocha, A. Simar, P. Singh, P. Pushan, K. Dutta, and Sundaravadivazhagan Editors, B. 2025. *Building Embodied AI Systems: The Agents, the Architecture Principles, Challenges, and Application Domains.* Springer.

Sani, A. I., and A.I. Sani. 2025. “Cyber Threat Intelligence for Industrial Automation: Al-Powered Strategies.” https://doi.org/10.4018/979-8-3373-3241-3.CH007

SERP. 2025. “ReGLU: GLU Activation Function and Its Variants.” SERP AI. https://serp.ai/posts/reglu/

The Judiciary. 2024. *All Courts Nationwide Go Digital.* News. https://judiciary.go.ke/judiciary-launches-e-filing-in-all-courts-data-tracking-dashboard-and-causelist-portal-portal/

Team, G., M. Riviere, S. Pathak, P.G. Sessa, C. Hardin, S. Bhupatiraju, L. Hussenot, T. Mesnard, B. Shahriari, A. Ramé, J. Ferret, P. Liu, P. Tafti, A. Friesen, M. Casbon, S. Ramos, R. Kumar, C. Lan, S. Le, Jerome, and A. Andreev. 2024. *Gemma 2: Improving Open Language Models at a Practical Size*. https://arxiv.org/pdf/2408.00118

Uddin, M. A., and I.H. Sarker. 2024. *An Explainable Transformer-based Model for Phishing Email Detection: A Large Language Model Approach*. https://arxiv.org/pdf/2402.13871

Vardhan R, V., and V. Kumar. 2025. “AI-Driven Cyber Threat Detection and Log Analysis.” *2025 International Conference on Inventive Computation Technologies (ICICT)*. 676–681. https://doi.org/10.1109/ICICT64420.2025.11004938

White and Case. 2024. *AI Watch: Global Regulatory Tracker - Kenya | White & Case LLP*. News. https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-kenya?utm_source=chatgpt.com

Zhang, N., and Q. Li. 2023. "MTD'23: 10th ACM Workshop on Moving Target Defense." *CCS 2023 - Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 3653–3654. https://doi.org/10.1145/3576915.3624022

Zhao, W. X., K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, Y. Du, C. Yang, Y. Chen, Z. Chen, J. Jiang, R. Ren, Y. Li, X. Tang, Z. Liu, and J.R Wen. 2023. "A Survey of Large Language Models." https://arxiv.org/pdf/2303.18223