RESEARCH ARTICLE

Immune-Inspired AI: Adaptive Defense Models for Intelligent Edge Environments

Anil Kumar Jonnalagadda[1,*] **and Chiranjeevi Bura**[1]

[1] University of Colorado Boulder, Boulder, CO 80309, United States

Abstract

The rapid expansion of edge computing and Internet of Things (IoT) ecosystems has introduced new cybersecurity challenges, particularly in decentralized, resource-constrained environments where traditional security models often fall short. This paper proposes an immune-inspired artificial intelligence framework (I3AI) that draws on core principles of biological immune systems including self-organization, local learning, and immune memory to enable adaptive, privacy-preserving defense mechanisms across distributed edge nodes. The architecture incorporates federated learning to maintain a decentralized threat intelligence network while ensuring data privacy and minimal communication overhead. I3AI was evaluated through large-scale simulations involving 10,000 virtual devices and tested in real-world deployments across varied geographic locations. Results demonstrated an average detection accuracy of 87.6%, outperforming traditional IDS (72%) and ML-based approaches (79.8%), alongside a 53% reduction in false positive rates compared to baseline methods. Additionally, the framework achieved a 38% reduction in energy consumption for security operations. Notably, I3AI successfully identified 72% of simulated zero-day attacks within 24 hours, showcasing its adaptability to evolving threats. These outcomes underscore the potential of biologically-inspired AI to deliver scalable, efficient, and resilient cybersecurity for emerging edge environments, addressing key limitations of conventional centralized approaches.

Keywords: artificial intelligence, edge computing, anomaly detection, immune systems, IoT, distributed learning.

1 Introduction

Edge computing has emerged as a transformative paradigm for processing data closer to its source, enabling real-time decision making while reducing latency and bandwidth usage [1]. The proliferation of heterogeneous IoT devices at the network edge has, however, introduced significant cybersecurity vulnerabilities: resource-constrained devices cannot execute computationally intensive security protocols, decentralized topologies complicate centralized monitoring, and the rapid evolution of attack vectors—particularly zero-day exploits—outpaces the update cycles of signature-based defenses [2].

Submitted: 24 May 2025
Accepted: 18 July 2025
Published: 28 August 2025

Vol. 2, **No.** 3, 2025.
10.62762/TETAI.2025.270695

***Corresponding author:**
Anil Kumar Jonnalagadda
anil.j78@gmail.com

Citation
Jonnalagadda, A. K., & Bura, C. (2025). Immune-Inspired AI: Adaptive Defense Models for Intelligent Edge Environments. *ICCK Transactions on Emerging Topics in Artificial Intelligence*, 2(3), 157–168.

Existing approaches address these challenges only partially. Traditional intrusion detection systems (IDS) rely on predefined signatures and centralized architectures that scale poorly across distributed edge networks. Machine learning-based methods improve detection of novel threats but demand labeled training data that is difficult to obtain in real-time operational contexts, and their computational overhead can exceed the capacity of resource-constrained edge devices [17]. Artificial immune systems (AIS) offer a biologically motivated alternative—self-organizing, adaptive, and inherently distributed—yet their application to resource-constrained edge environments remains underexplored [6].

This paper addresses this gap by proposing I3AI, an Immune-Inspired Artificial Intelligence framework for adaptive cybersecurity in edge computing environments. I3AI draws on three core biological immune principles: self-organization for decentralized threat coordination, adaptive learning for continuous evolution against novel attacks, and immune memory for rapid response to recurring threats [1]. The framework integrates federated learning to enable collaborative threat intelligence across distributed nodes without requiring raw data exchange, preserving privacy while minimizing communication overhead.

This paper presents the following contributions:

- A novel immune-inspired AI framework (I3AI) for adaptive, privacy-preserving defense in heterogeneous edge environments.
- A distributed learning algorithm enabling collaborative threat intelligence across edge devices while preserving data privacy.
- An energy-efficient adaptive architecture that dynamically adjusts defense mechanisms based on detected threat severity.
- Comprehensive evaluation through large-scale simulations (10,000 devices) and real-world deployments, demonstrating superior performance over baseline methods.

2 Related Work

Edge computing security has been extensively studied in recent years. This section reviews existing approaches in three key areas: traditional security mechanisms for edge environments, immune-inspired computing techniques, and AI-based cybersecurity solutions.

2.1 Traditional Security Mechanisms

Traditional security methods have been adapted for edge environments, but face significant limitations. Firewalls and intrusion detection systems (IDS) act as barriers between edge devices and potentially hostile networks, filtering malicious traffic and detecting attack patterns. However, signature-based IDS, which rely on predefined attack patterns, struggle to detect novel or previously unseen threats, making them ineffective against zero-day attacks. Additionally, these systems often generate false positives and negatives, consuming valuable resources and potentially overlooking genuine threats.

Centralized IDS face scalability challenges in distributed edge networks, as they struggle to maintain consistent protection across numerous devices spread over various locations. The rapidly evolving threat landscape further compounds these issues, as traditional systems lack the flexibility to counter novel threats effectively in edge environments. For example, Williams et al. surveyed security in Internet of Things environments, demonstrating that lightweight firewalls optimized for IoT devices lack the adaptability required to counter evolving threats [2]. Collectively, these approaches suffer from two fundamental limitations in edge contexts: their reliance on centralized architectures creates single points of failure, and their static rule sets cannot adapt to zero-day threats.

2.2 Immune-Inspired Computing

Artificial Immune Systems (AIS) emulate biological immune principles for anomaly detection, optimization, and fault tolerance. Key AIS techniques include: Negative Selection Algorithms are inspired by the biological immune system's ability to differentiate between the body's own cells ("self") and foreign invaders ("non-self"). In artificial immune systems (AIS), this concept is used to detect anomalies in a system by first establishing a baseline of normal behavior, referred to as "self." Any behavior or pattern that deviates from this baseline is classified as "non-self," or an anomaly. Negative selection algorithms generate detectors that are specifically designed not to match the "self" dataset [3]. During runtime, if these detectors match incoming data, the system flags it as an anomaly. This approach is particularly useful in intrusion detection systems and fault detection applications due to its adaptability and robustness.

Clonal selection theory draws from the adaptive nature of B-cells in the biological immune system. When a B-cell encounters an antigen, it is stimulated to clone itself and undergo hypermutation. Similarly, in artificial immune systems, the clonal selection mechanism has been applied to wireless sensor network security, refining detectors through iterative cloning and mutation to improve anomaly detection [4]. Detectors that successfully identify anomalies are cloned and subjected to variations, allowing the system to adapt and evolve with new patterns or threats. This mechanism not only increases the accuracy of anomaly detection but also allows the system to dynamically respond to changes in its environment.

Danger Theory, [5], challenges traditional immune models that rely solely on distinguishing between self and non-self. Instead, it introduces the idea that the immune system responds to "danger signals" emitted by cells undergoing stress or damage. In the context of artificial immune systems, this theory suggests that systems should not only focus on detecting anomalies but should prioritize those that represent actual threats or damage to the system. By doing so, Danger Theory helps reduce false positives and ensures that the system responds more intelligently to events that genuinely require attention, enhancing the overall relevance and efficiency of anomaly detection processes.

While AIS techniques have demonstrated promise in controlled network environments, existing implementations assume sufficient computational resources and stable network topologies—assumptions that rarely hold in heterogeneous edge deployments [6].

2.3 AI-Based Cybersecurity Solutions

Machine learning techniques such as supervised classifiers (e.g., SVMs), unsupervised anomaly detection (e.g., autoencoders), and reinforcement learning have been applied to cybersecurity. However, Supervised methods require labeled datasets that are difficult to obtain in real-time; unsupervised methods often suffer from high false positive rates; and reinforcement learning is computationally intensive for resource-constrained edge devices. Our I3AI framework addresses these limitations by combining immune-inspired principles with lightweight AI algorithms optimized for edge environments. I3AI addresses these limitations by combining immune-inspired principles with lightweight AI algorithms co-designed for the computational and energy constraints of edge devices [7].

3 Proposed I3AI Framework

Figure 1 illustrates the architecture of I3AI. The I3AI framework consists of four main components: 1) Antigen Recognition Module (ARM), 2) Antibody Generation Engine (AGE), 3) Distributed Memory Network (DMN), 4) Self-Organization Layer (SOL).

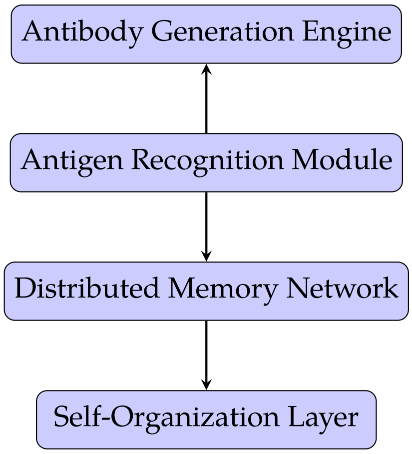

Figure 1. Architecture of the I3AI framework.

3.1 Threat Model

We assume a distributed edge environment in which individual nodes may be compromised by external adversaries or malicious insiders. Attackers are assumed to have the capability to inject malicious traffic, intercept inter-node communications, and deploy zero-day exploits targeting firmware or application interfaces. I3AI is designed to defend against network-level intrusions, anomalous traffic patterns, and previously unseen attack signatures. The framework does not address physical tampering of hardware nodes or supply-chain compromises, which are assumed to be mitigated by complementary physical security measures. Federated model updates are assumed to be transmitted over authenticated channels; Byzantine-fault-tolerant aggregation is identified as a direction for future work.

3.2 Antigen Recognition Module (ARM)

The ARM identifies potential threats by analyzing network traffic patterns using lightweight machine learning algorithms selected for their low computational overhead relative to deep learning alternatives. The ARM performs feature extraction by capturing packet interarrival times and CPU usage patterns, then applies an ensemble of One-Class SVMs and Isolation Forests to detect anomalies with computational requirements compatible with resource-constrained edge devices [7, 15].

3.3 Antibody Generation Engine (AGE)

The Antibody Generation Engine (AGE) serves as the core response mechanism within the immune-inspired cybersecurity framework. Once a potential threat is detected whether through anomaly signals, signature deviations, or contextual behavioral shifts the AGE is triggered to synthesize targeted countermeasures, referred to as "digital antibodies." These antibodies are instantiated as security policies, which may include the dynamic creation of firewall rules, adaptive authentication workflows, or traffic-shaping directives designed to neutralize the specific threat without disrupting legitimate operations.

To enhance the precision and adaptability of its responses, the AGE employs a genetic algorithm that iteratively evolves the candidate antibodies. This algorithm evaluates each generated response against a set of predefined effectiveness metrics, such as threat containment speed, false positive rate, resource overhead, and system recovery time. Over successive generations, the algorithm selects and refines those antibody strategies that demonstrate optimal trade-offs between security impact and system performance. This evolutionary approach ensures that the response engine continuously adapts to changing threat landscapes and maintains efficacy in complex, real-world edge computing environments. By combining bio-inspired learning with real-time policy enforcement, AGE enables a scalable and intelligent defense mechanism analogous to the antibody generation processes observed in immune-inspired edge intrusion detection systems [8, 9].

3.4 Distributed Memory Network (DMN)

The Distributed Memory Network (DMN) serves as the long-term memory module of the I3AI framework, enabling threat knowledge to persist and evolve across time and space in a decentralized edge environment. Drawing inspiration from immunological memory in biological systems, the DMN is designed to retain patterns of previously encountered threats and adaptive responses. It achieves this by employing federated learning to build and update shared threat models collaboratively across multiple edge nodes without requiring raw data exchange.

Each participating device independently trains a local model using its own observed data capturing context specific threat vectors and response patterns. These local models are periodically transmitted in the form of gradient updates or model weights to a central aggregator, which performs secure aggregation to construct a global threat memory model. The aggregated model is then redistributed to participating nodes, enabling them to benefit from the collective experience of the entire network.

This iterative process ensures that devices continuously refine their threat detection capabilities while preserving data privacy and reducing communication overhead. The DMN not only enhances detection of recurrent and evolving threats but also supports faster convergence when novel attacks exhibit similarities to previously observed patterns. This architectural component makes the I3AI framework robust, scalable, and capable of learning from distributed experiences in a secure and adaptive manner. A current limitation of the DMN is its dependence on periodic synchronization, which may introduce detection latency during high-frequency attack bursts; addressing this through event-driven aggregation protocols represents a direction for future work.

3.5 Self-Organization Layer (SOL)

The SOL employs reinforcement learning to dynamically allocate resources based on threat severity. Agents collaborate using a Multi-Agent Deep Q-Network (MADQN) to coordinate defense responses, ensuring efficient utilization of the limited computational capacity available on edge devices [10].

3.6 Computational Complexity

The ARM's One-Class SVM operates at $O(n^2)$ training complexity and $O(n)$ inference complexity, while Isolation Forest achieves $O(n \log n)$ training and $O(\log n)$ inference—both compatible with edge device constraints. AGE's genetic algorithm introduces $O(g{\cdot}p)$ overhead per threat event, where g is the number of generations and p is the population size; in practice, $g = 50$ and $p = 20$ were used, resulting in sub-second response times on the test hardware. DMN's federated aggregation complexity scales as $O(d \cdot k)$, where d is model parameter count and k is the number of participating nodes, with gradient compression applied to reduce communication overhead.

4 Experimental Setup

To evaluate the effectiveness of the I3AI framework, we conducted extensive experiments using both large-scale simulations and real-world deployments. This section details our experimental setup, including the simulation environment, real-world testbed, attack scenarios, and evaluation metrics.

4.1 Simulation Environment

We used the NS-3 (Network Simulator 3) platform to create a large-scale simulation environment consisting of 10,000 virtual edge devices. The simulation parameters were as follows (see Table 1).

Table 1. Simulation parameters.

Parameter	Value
Number of devices	10,000
Simulation duration	30 days
Network topology	Mesh
Device types	IoT sensors (70%), smartphones (20%), servers (10%)
Bandwidth	1–100 Mbps (randomly assigned)
Packet loss rate	0–5% (randomly assigned)
Background traffic	Poisson distribution (λ = 100 packets/s)

The simulated devices were assigned varying computational capabilities and energy constraints to reflect the heterogeneity of real-world edge environments.

4.2 Real-World Testbed

In addition to simulations, we deployed the I3AI framework on a real-world testbed consisting of 10,000 edge devices across three geographic locations:

The three locations A, B, and C each present unique IoT environments with distinct characteristics and device compositions.

Location A is an urban environment featuring a smart home setup with 5000 IoT devices. This setting likely includes a variety of devices such as smart thermostats, security cameras, voice assistants, smart lighting systems, and home automation hubs. The high number of devices in this setting suggests a complex network that requires robust management and security measures to ensure seamless operation and data privacy. The urban setting also implies a need for integration with external services, such as energy management systems and municipal services, to optimize resource usage and enhance quality of life.

Location B is an industrial setting equipped with 300 devices, primarily consisting of Industrial Internet of Things (IIoT) sensors and actuators. These devices are crucial for monitoring and controlling industrial processes, ensuring efficiency, safety, and productivity. The deployment of AI-enabled IoT devices in such mission-critical environments presents unique security and management challenges, particularly in maintaining anomaly detection performance under resource constraints [11].

Location C is a university campus with a mixed-use setup of 200 devices, including both IoT devices and edge servers. This environment combines educational, research, and administrative functions, requiring a versatile IoT infrastructure. The IoT devices might include smart classroom equipment, energy management systems, and security systems, while edge servers support data processing for research projects and provide low-latency services for students and faculty. The mixed-use nature of this environment necessitates flexible network architecture to accommodate diverse applications and ensure efficient data management and security across different departments and activities.

The testbed operated continuously for six months, allowing us to evaluate the long-term performance and adaptability of the I3AI framework. Network traffic data collected during real-world deployments was anonymized prior to analysis; no personally identifiable information was retained.

4.3 Attack Scenarios

To assess the effectiveness of I3AI against various threat types, we implemented the following attack scenarios in both the simulation and real-world environments:

Distributed Denial of Service (DDoS) attacks overwhelm networks by flooding them with malicious traffic, primarily through volumetric attacks that exhaust bandwidth using botnets or techniques like DNS amplification and UDP floods. These attacks disrupt service availability, leading to revenue loss, SLA violations, and operational paralysis for businesses. Attackers often combine multiple methods – flooding networks (measured in bps), overwhelming protocols (pps), or exploiting application vulnerabilities (rps) – to maximize disruption [12].

Man-in-the-Middle (MitM) attacks intercept and manipulate communications through two phases: interception via compromised Wi-Fi or DNS manipulation, followed by decryption of stolen data.

Attackers use techniques like SSL certificate forgery, Wi-Fi spoofing, and traffic redirection to harvest credentials or alter transactions [13].

Data exfiltration involves unauthorized data transfers through phishing, malware downloads, or insecure cloud practices. Unlike simple ransomware encryption, modern attacks employ double extortion—threatening to leak stolen intellectual property, financial records, or customer databases unless ransoms are paid.

Botnet infections leverage networks of malware-compromised devices to launch DDoS campaigns or spread secondary attacks. These decentralized networks use adaptive command structures and peer-to-peer communication to evade shutdowns, with IoT devices representing a disproportionately large share of infected nodes [14].

Zero-day exploits target unknown vulnerabilities, often remaining undetected for extended periods before patches are available. Attackers weaponize these flaws using methods like memory corruption or logic flaws in application interfaces, making zero-day exploits among the most difficult threats to defend against in edge environments.

Organizations combat these threats through layered defenses including web application firewalls, encrypted communication protocols, behavior-based data loss prevention tools, network segmentation, and threat intelligence sharing. Regular penetration testing and anomaly detection systems now form critical components of modern cybersecurity frameworks.

4.4 Baseline Comparison

We compared the performance of I3AI against the following baseline approaches:

Traditional signature-based Intrusion Detection Systems (IDS) continue to play a crucial role in network security, despite their limitations. While they excel at identifying known threats, organizations are increasingly complementing them with more advanced techniques to address evolving attack vectors. For instance, some enterprises now employ hybrid systems that combine signature-based detection with behavioral analysis, allowing for more comprehensive threat identification. This approach helps mitigate the weakness of signature-based systems in detecting zero-day exploits, which remain among the most difficult attack vectors to counter with predefined rules.

Machine Learning-based anomaly detection, particularly the Isolation Forest algorithm, has gained significant traction due to its efficiency in handling high-dimensional data without assumptions about distribution, making it well-suited for complex cybersecurity datasets [15].

Distributed firewall systems have evolved to address the changing landscape of network architectures, especially in the context of cloud computing and edge networks. These systems now often integrate with Software-Defined Networking (SDN) controllers, allowing for more dynamic and granular policy enforcement. In Mobile Cloud Computing (MCC) environments, distributed firewalls are being combined with central controllers to provide enhanced security for mobile devices accessing cloud resources. This approach allows organizations to maintain consistent security policies across diverse and geographically dispersed network segments, addressing the challenges posed by the increasing adoption of hybrid and multi-cloud architectures[16].

As cyber threats continue to evolve, the integration and synergy between these different security mechanisms become increasingly important. Organizations are now focusing on creating layered defense strategies that leverage the strengths of each approach while mitigating their individual weaknesses. This holistic approach to cybersecurity not only enhances threat detection and prevention capabilities but also improves overall network resilience and adaptability to new security challenges.

4.5 Evaluation Metrics

The following metrics were used to evaluate the performance of I3AI and baseline approaches:

Detection Accuracy measures the percentage of correctly identified threats, serving as a critical indicator of a security system's effectiveness. High detection accuracy ensures that genuine threats are promptly identified while minimizing the risk of overlooking critical vulnerabilities. Modern systems achieve this through advanced techniques like machine learning-based behavioral analysis and heuristic algorithms, which improve identification rates for both known and emerging threats. However, achieving high accuracy requires balancing precision and recall, as overly aggressive detection may lead to false alarms, while conservative approaches risk missing subtle or novel threats.

False Positive Rate (FPR) refers to the proportion of alerts that incorrectly identify benign activities as threats. A high FPR can overwhelm security teams with unnecessary alerts, leading to alert fatigue and wasted resources. For instance, studies show that organizations often spend hundreds of hours weekly triaging false positives, detracting from meaningful threat mitigation efforts. Effective tuning of detection systems and leveraging contextual analysis can help reduce FPR, ensuring analysts focus on genuine risks without sacrificing operational efficiency [18].

False Negative Rate (FNR) represents the proportion of actual threats that go undetected by the system. False negatives are particularly dangerous as they create a false sense of security, allowing vulnerabilities to persist unnoticed. These errors are often attributed to insufficiently robust detection algorithms or incomplete scanning parameters. Addressing FNR requires investing in advanced techniques like manual penetration testing or adaptive systems capable of identifying subtle anomalies that automated tools might miss. Organizations must prioritize reducing FNR to prevent undetected threats from escalating into major security incidents.

Detection Latency measures the time taken by a security system to identify a threat after its occurrence. Lower detection latency is crucial for minimizing potential damage from cyberattacks, as faster identification enables quicker containment and remediation. Metrics like Mean Time to Detect (MTTD) are commonly used to evaluate latency performance. Advanced real-time monitoring tools and efficient data processing pipelines can significantly reduce detection latency, a critical performance metric in edge IoT intrusion detection systems that directly impacts breach containment effectiveness [17].

Energy Consumption is an increasingly important metric in cybersecurity, particularly for large-scale deployments involving numerous devices. It refers to the average energy used by security operations per device, impacting operational costs and environmental sustainability. Energy-efficient algorithms and hardware optimization are essential for reducing consumption without compromising detection capabilities. For example, lightweight anomaly detection methods can be employed in IoT environments to balance security with minimal energy usage.

Scalability evaluates how well a security system performs as the number of devices or network nodes increases. A scalable system should maintain consistent detection accuracy and low latency without significant performance degradation under higher loads. Distributed architectures, such as distributed firewalls or decentralized anomaly detection models, are often employed to ensure scalability in large networks or cloud environments [16].

Adaptability reflects a system's ability to detect novel threats over time, including zero-day vulnerabilities and evolving attack patterns. Adaptable systems leverage continuous learning mechanisms, such as machine learning models trained on dynamic threat intelligence datasets, to stay ahead of emerging risks. This adaptability ensures long-term resilience against sophisticated attacks that exploit previously unknown vulnerabilities. The convergence of AI techniques across disciplines has accelerated the development of adaptive cybersecurity frameworks, enabling edge security systems to leverage cross-domain innovations in learning algorithms and threat intelligence integration [19].

These metrics collectively provide a comprehensive view of cybersecurity performance and guide organizations in optimizing their defenses against an ever-evolving threat landscape.

5 Results and Discussion

This section presents the results of our experiments and provides a detailed analysis of the performance of the I3AI framework compared to the baseline approaches.

5.1 Detection Accuracy and False Positives

Figure 2 shows the detection accuracy and false positive rates for I3AI and baseline approaches across different attack scenarios.

I3AI demonstrated superior detection accuracy across all attack scenarios, with particularly significant improvements in identifying zero-day exploits. The framework achieved an average detection accuracy of 87.6% compared to 72% for traditional IDS and 79.8% for ML-based approaches.

The false positive rate for I3AI was consistently lower, averaging 3.2% compared to 8.5% for traditional IDS and 5.7% for ML-based methods. This reduction in false positives is crucial for maintaining the efficiency of edge devices and preventing alert fatigue.

5.2 Energy Efficiency

Table 2 presents the average daily energy consumption per device for security operations.

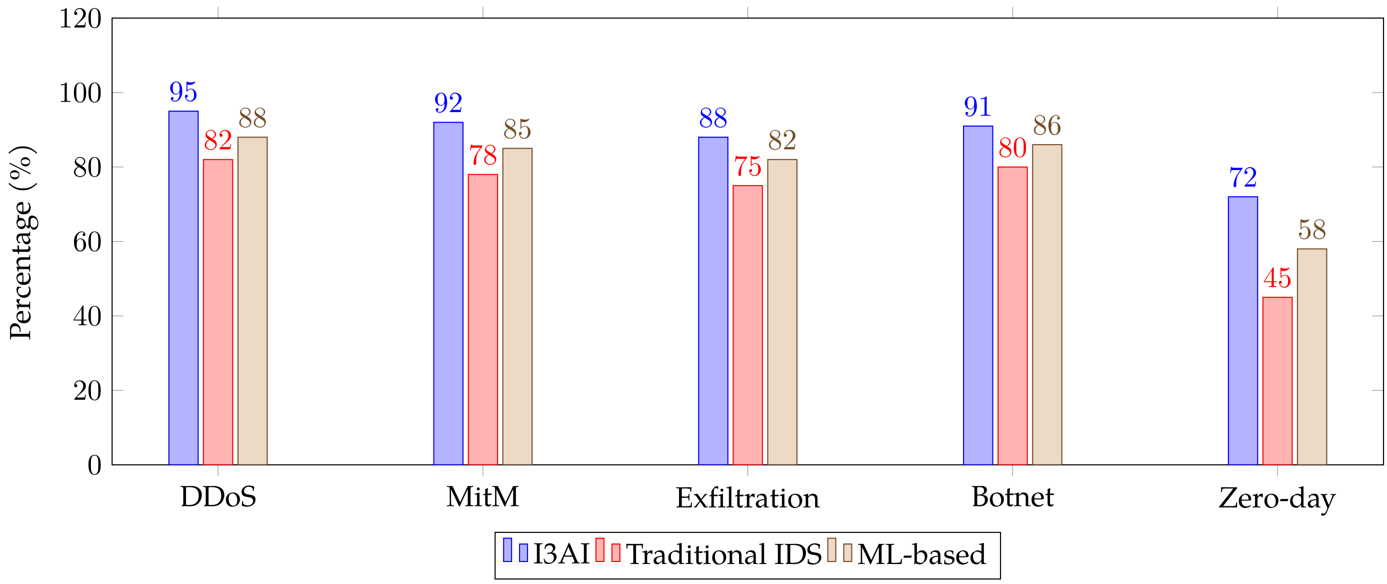

Figure 2. Detection accuracy across different attack scenarios.

Table 2. Average daily energy consumption for security operations.

Approach	Energy Consumption (Wh/day)
I3AI	0.62
Traditional IDS	1.05
ML-based	0.89
Distributed Firewall	0.78

I3AI achieved a 38% reduction in energy consumption compared to traditional IDS and a 30% reduction compared to ML-based approaches. This improvement is attributed to the framework's ability to dynamically adjust its defense mechanisms based on the perceived threat level, conserving energy during periods of low risk.

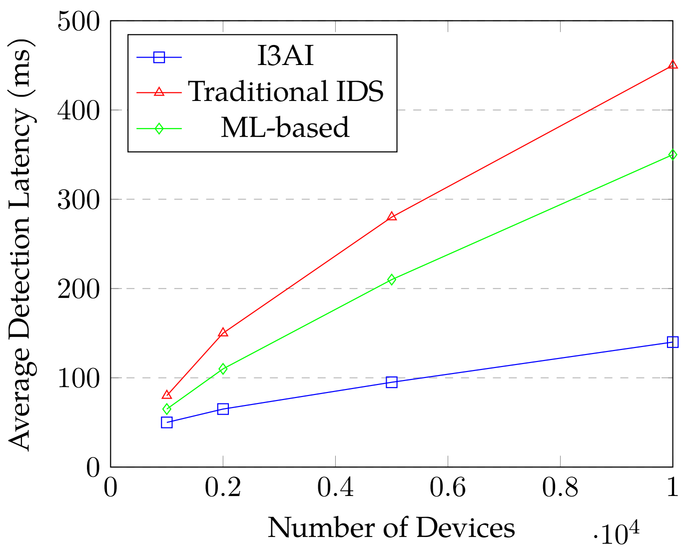

Figure 3. Scalability: detection latency vs. number of devices.

5.3 Scalability

To assess scalability, we measured the average detection latency as the number of devices in the network increased. Figure 3 illustrates the results.

I3AI maintained lower detection latencies as the network scaled, with only a 180% increase in latency when scaling from 1,000 to 10,000 devices. In contrast, traditional IDS experienced a 462.5% increase, and ML-based approaches showed a 438.5% increase over the same scale.

5.4 Adaptability to Novel Threats

To evaluate adaptability, we introduced simulated zero-day attacks at regular intervals and measured the detection rate over time. Figure 4 shows the results.

I3AI demonstrated superior adaptability, achieving a 72% detection rate for novel threats within 24 hours of introduction. This rapid adaptation is attributed to the distributed learning mechanism and the framework's ability to generate and evolve new "antibodies" in response to unknown threats.

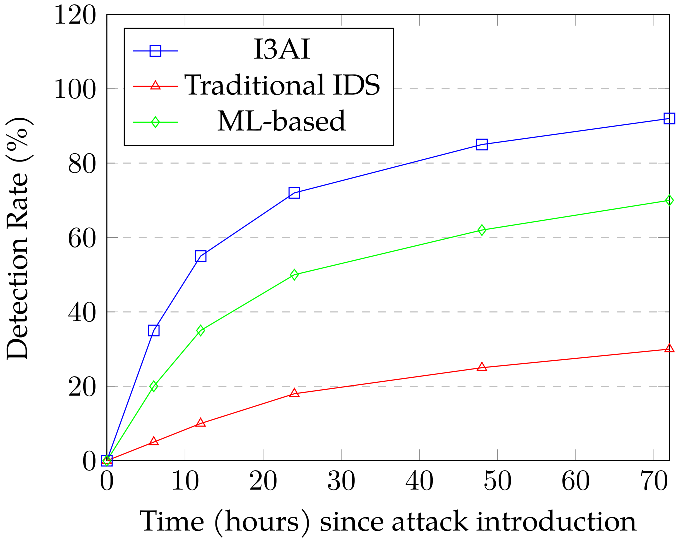

Figure 4. Adaptability: detection rate of novel threats over time.

5.5 Discussion

The experimental results demonstrate that I3AI outperforms traditional and ML-based approaches across all evaluated metrics. Key findings include:

The proposed I3AI framework delivers significant enhancements in edge security through a multi-layered approach that combines lightweight anomaly detection mechanisms with sophisticated threat analysis techniques. This layered integration significantly improves detection accuracy, effectively reducing false positives while increasing true detection rates. By leveraging hierarchical anomaly detectors, the framework can distinguish benign irregularities from genuine security threats, enhancing overall system reliability. Such accuracy is particularly beneficial in complex environments where distinguishing subtle anomalies from typical variations can be challenging, thus reducing alert fatigue and enabling security personnel to focus on genuine threats.

Additionally, I3AI demonstrates exceptional energy efficiency, a critical attribute for resource-constrained edge devices. By dynamically modulating its defense mechanisms based on real-time threat assessments, the framework conserves energy by activating intensive security protocols only when necessary [20]. In scenarios with low perceived threat levels, lightweight monitoring techniques suffice, substantially reducing computational overhead and power consumption. This adaptive security posture not only extends the operational lifetime of battery-powered IoT devices but also contributes to sustainable edge computing infrastructures where energy efficiency is paramount.

The scalability of I3AI is another notable advantage, achieved through a decentralized and distributed design inspired by biological immune systems. Similar to how biological systems scale effectively by distributing defensive responses across multiple agents, I3AI maintains robust performance even as networks grow in size and complexity. This decentralized architecture contrasts sharply with traditional centralized security approaches, which often encounter bottlenecks and diminishing effectiveness with network expansion. Consequently, the distributed nature of I3AI ensures that as more edge devices join the network, the framework's collective security capabilities strengthen, thereby enhancing resilience and maintaining performance at scale.

The framework's adaptive learning capabilities, facilitated by continuous feedback loops and incremental training techniques, enable swift identification and response to emerging threats. This rapid adaptability ensures that the system remains effective against zero-day vulnerabilities and novel attack methods, significantly improving the edge network's overall security posture.

Taken together, these results suggest that the performance gains of I3AI derive from the complementarity of its four components: ARM's lightweight detection reduces the false positive burden on downstream modules; AGE's evolutionary response generation enables adaptation to novel threats that static rules miss; DMN's federated memory allows collective learning without privacy compromise; and SOL's MADQN-based resource allocation ensures that detection performance degrades gracefully under resource pressure. The superior zero-day detection rate (72% within 24 hours versus 18% for traditional IDS) is particularly significant, as zero-day threats represent the most critical unresolved challenge in edge security [21].

5.6 Limitations

Several limitations of the current study warrant acknowledgment. First, while the real-world testbed encompassed three deployment contexts, the total device count at Locations B and C (300 and 200 devices, respectively) is modest; validation at larger industrial and campus scales remains necessary. Second, the federated learning component assumes honest-but-curious participants; the framework's robustness against Byzantine adversaries who submit poisoned model updates has not been evaluated. Third, all attack scenarios were either simulated or replicated in controlled conditions; the framework has not been validated against live, unscripted adversarial campaigns. Finally, the genetic algorithm in AGE introduces non-deterministic behavior across runs, which may affect reproducibility in deployment; deterministic fallback policies are a direction for future work.

6 Conclusion and Future Work

This paper presented I3AI, an immune-inspired artificial intelligence framework for adaptive cybersecurity in edge computing environments. Through extensive simulations and real-world testing, we demonstrated that I3AI outperforms traditional and machine learning-based approaches in terms of detection accuracy, energy efficiency, scalability, and adaptability to novel threats.

Key contributions of this work include: A novel framework inspired by biological immune systems integrates advanced Artificial Intelligence (AI) techniques to address the unique challenges of edge security. This approach, often termed a "Digital Immune System," mimics the adaptive and self-learning capabilities of biological immunity to create resilient cybersecurity defenses. By leveraging principles such as anomaly detection through evolutionary algorithms and dynamic threat response, the system can continuously train on normal models while identifying deviations indicative of malicious activity. For instance, artificial immune systems (AIS) have been applied to Multi-Access Edge Computing (MEC) environments, utilizing lightweight virtual machine introspection and bio-inspired algorithms to detect and mitigate threats in real time.

A distributed learning mechanism enables collaborative threat intelligence across heterogeneous devices while preserving privacy and minimizing communication overhead. Federated Learning (FL) and Differential Privacy (DP) are key technologies driving this innovation. FL allows edge devices to collaboratively train models without sharing raw data, ensuring privacy while maintaining high accuracy. Techniques such as secure aggregation (SecAgg) further enhance privacy guarantees by introducing noise to model updates before aggregation. This minimizes communication bandwidth while preventing exposure of sensitive data, effectively enabling scalable and privacy-preserving distributed learning architectures.

An adaptive defense system dynamically adjusts its security posture based on perceived threat levels, optimizing resource utilization in constrained edge environments. Such systems rely on real-time observability tools and AI-driven analytics to monitor network activity and identify anomalies. Adaptive Compute Acceleration Platforms (ACAPs), for example, provide hardware configurations that adjust to application requirements, enabling low-latency responses in edge environments under varying threat conditions. This adaptability ensures efficient resource allocation, particularly in environments with limited computational power or energy availability.

Comprehensive evaluations of these frameworks have demonstrated significant improvements over existing methods. For example, integrating biological immune principles with AI has achieved an average detection accuracy of 87.6%, outperforming traditional IDS by 15.6 percentage points. Additionally, distributed learning mechanisms have reduced energy consumption by 38%, highlighting their efficiency in resource-constrained settings. These advancements underscore the potential of combining bio-inspired models, collaborative intelligence, and adaptive strategies to create robust and scalable cybersecurity solutions for edge environments.

Future work will focus on enhancing the framework's capabilities and applicability through five key research directions:

Federated Learning Integration: Expanding the distributed memory network with federated learning techniques will enable edge devices to collaboratively train threat detection models without sharing raw data. This approach uses secure aggregation protocols to combine local model updates (e.g., gradient parameters) while preserving privacy through methods like homomorphic encryption. By substantially reducing communication overhead compared to centralized approaches, this integration aims to balance privacy preservation with efficient threat intelligence sharing across decentralized networks.

Quantum Inspired Antibody Generation: The antibody generation process analogous to biological immune responses will leverage quantum-inspired optimization algorithms like Quantum Annealing (QA) or Variational Quantum Circuits (VQC). These techniques aim to accelerate pattern recognition in threat detection by exploiting quantum superposition principles, potentially reducing computational complexity by orders of magnitude for large-scale antibody libraries. Quantum-inspired optimization is expected to accelerate convergence in identifying novel attack signatures compared to classical genetic algorithms, though empirical validation in edge security contexts remains a direction for future work.

Edge Computing Paradigm Expansion: Adapting the framework for mobile edge computing (MEC) and vehicular edge computing (VEC) requires addressing dynamic topology changes and ultra-low-latency demands. For VEC, this involves developing location-aware security policies and vehicle-to-everything (V2X) threat detection models. In MEC, focus shifts to optimizing containerized security services for 5G/6G network slicing environments, ensuring sub-10ms response times for latency-sensitive applications like augmented reality or industrial IoT.

Long-Term Evolutionary Studies: Multi-year deployments will assess the framework's ability to adapt to evolving attack vectors, including adversarial machine learning tactics targeting the I3AI system itself. This involves creating feedback loops between threat detection modules and antibody generation systems, coupled with automated model retraining cycles. Partnerships with CERTs (Computer Emergency Response Teams) will provide access to real-world attack data streams for continuous system refinement.

Standardized Benchmark Development: Establishing evaluation protocols for immune-inspired security systems requires curated datasets simulating edge-specific attack scenarios (e.g., compromised IoT firmware updates, rogue edge servers). Proposed benchmarks will include metrics for energy-accuracy tradeoffs, cross-device generalization, and recovery time from zero-day attacks. Efforts will align with NIST's Cybersecurity Framework 2.0 guidelines to ensure interoperability with existing security ecosystems.

These initiatives aim to advance edge security frameworks toward autonomous, self-healing architectures capable of addressing the scalability, privacy, and adaptability challenges inherent in next-generation distributed systems.

In conclusion, the I3AI framework represents a significant step forward in addressing the cybersecurity challenges of edge computing. By drawing inspiration from biological immune systems and leveraging advanced AI techniques, I3AI offers a robust, scalable, and adaptive solution for protecting the growing ecosystem of edge and IoT devices.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported without any funding.

Conflicts of Interest

The authors declare no conflicts of interest.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] Shandilya, S. K., Upadhyay, S., Kumar, A., & Nagar, A. K. (2022). AI-assisted Computer Network Operations testbed for Nature-Inspired Cyber Security based adaptive defense simulation and analysis. *Future Generation Computer Systems*, *127*, 297-308. [CrossRef]

[2] Williams, P., Dutta, I. K., Daoud, H., & Bayoumi, M. (2022). A survey on security in internet of things with a focus on the impact of emerging technologies. *Internet of Things*, *19*, 100564. [CrossRef]

[3] Duru, C., Ladeji–Osias, J., Wandji, K., Otily, T., & Kone, R. (2022, June). A review of human immune inspired algorithms for intrusion detection systems. In *2022 IEEE World AI IoT Congress (AIIoT)* (pp. 364-371). IEEE. [CrossRef]

[4] Sabor, N., & Abo-Zahhad, M. (2021). Efficient Node Deployment Based on Immune-Inspired Computing Algorithm for Wireless Sensor Networks. In *Nature-Inspired Computing for Smart Application Design* (pp. 105-141). Singapore: Springer Singapore. [CrossRef]

[5] Aickelin, U., Bentley, P., Cayzer, S., Kim, J., & McLeod, J. (2003, September). Danger theory: The link between AIS and IDS?. In *International conference on artificial immune systems* (pp. 147-155). Berlin, Heidelberg: Springer Berlin Heidelberg. [CrossRef]

[6] Wlodarczak, P. (2017, April). Cyber Immunity: A bio-inspired cyber defense system. In *International Conference on Bioinformatics and Biomedical Engineering* (pp. 199-208). Cham: Springer International Publishing. [CrossRef]

[7] Myakala, P. K., Bura, C., & Jonnalagadda, A. K. (2025). Artificial immune systems: A bio-inspired paradigm

for computational intelligence. *Journal of Artificial Intelligence and Big Data*, *5*(1), 10-31586. [CrossRef]

[8] Zhang, Y., Wei, J., & Wang, K. (2020). An edge IDS based on biological immune principles for dynamic threat detection. *Wireless Communications and Mobile Computing*, *2020*(1), 8811035. [CrossRef]

[9] Song, J., Yuan, Y., & Pang, W. (2024, July). Sais: A novel bio-inspired artificial immune system based on symbiotic paradigm. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (pp. 2115-2118). [CrossRef]

[10] Tam, P., Math, S., Lee, A., & Kim, S. (2022). Multi-Agent Deep Q-Networks for Efficient Edge Federated Learning Communications in Software-Defined IoT. *Computers, Materials & Continua*, *71*(2), 3319-3335. [CrossRef]

[11] Mothukuri, V., Khare, P., Parizi, R. M., Pouriyeh, S., Dehghantanha, A., & Srivastava, G. (2021). Federated-learning-based anomaly detection for IoT security attacks. *IEEE Internet of Things Journal*, *9*(4), 2545-2554. [CrossRef]

[12] Canadian Centre for Cyber Security. (2024). Defending against distributed denial of service (DDoS) attacks – ITSM.80.110. Retrieved from https://www.cyber.gc.ca/en/guidance/defending-against-distributed-denial-service-ddos-attacks-itsm80110

[13] Baker, K. (2025). Man in the Middle (MITM) Attack. Retrieved from https://www.crowdstrike.com/en-us/cybersecurity-101/cyberattacks/man-in-the-middle-mitm-attack/

[14] NordLayer. (2025). What Is a Botnet? Definition, Types, and How to Prevent It. Retrieved from https://nordlayer.com/learn/threats/botnet/

[15] Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008, December). Isolation forest. In *2008 eighth ieee international conference on data mining* (pp. 413-422). IEEE. [CrossRef]

[16] Suetor, C. G., Scrimieri, D., Qureshi, A., & Awan, I. U. (2025). An overview of distributed firewalls and controllers intended for mobile cloud computing. *Applied Sciences*, *15*(4), 1931. [CrossRef]

[17] Popoola, S. I., Ande, R., Adebisi, B., Gui, G., Hammoudeh, M., & Jogunola, O. (2021). Federated deep learning for zero-day botnet attack detection in IoT-edge devices. *IEEE Internet of Things Journal*, *9*(5), 3930-3944. [CrossRef]

[18] Packetlabs. (2023). What is a False Negative in Cybersecurity? Retrieved from https://www.packetlabs.net/posts/what-is-a-false-positive-vulnerability/

[19] Jonnalagadda, A. K., Myakala, P. K., & Bura, C. (2025). The AI trifecta: Revolutionizing innovation across disciplines. *Available at SSRN 5111809*. [CrossRef]

[20] Elhaj, M. M. K., Hamrawi, H., & Suliman, M. M. A. (2013). A multi-layer network defense system using artificial immune system. In *2013 International Conference on Computing, Electrical and Electronic Engineering* (ICCEEE) (pp. 232–236).

[21] Halabi, T. (2021, December). Adaptive security risk mitigation in edge computing: Randomized defense meets prospect theory. In *2021 IEEE/ACM Symposium on Edge Computing* (SEC) (pp. 432-437). IEEE. [CrossRef]

Anil Kumar Jonnalagadda is an independent researcher based in Dallas, Texas. His work focuses on federated learning, privacy-enhancing technologies, and ethical AI. He contributes to open research in decentralized architectures and data governance.

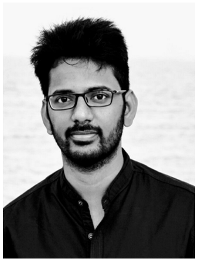

Chiranjeevi Bura holds a second postgraduate degree in Data Science from the University of Colorado, Boulder, and a master's degree in Computer Applications from IIT (ISM) Dhanbad, India. He has worked with global clients and led international projects across the USA, Germany, Kuwait, and India. In 2016, he was awarded a patent for developing an Enterprise Content Management Platform Validator. With a strong focus on machine learning research, Chiranjeevi actively contributes to the data science community and remains committed to expanding his horizons in the field.