JOURNAL OF SOFT COMPUTING AND DATA MINING
e-ISSN: 2716-621X

JSCDM

Vol. 6 No. 1 (2025) 394-405
https://publisher.uthm.edu.my/ojs/index.php/jscdm

Adaptive Intrusion Response via Federated Meta-Learning for IIoT Zero-Day Mitigation

Hani Q. R. Al-Zoubi[1], Mohammed Salahat[2], Nidal A. Al-Dmour[1]*

[1] *Department of Computer Engineering, College of Engineering, Mutah University, Karak, 61710, JORDAN*
[2] *College of Engineering and Technology, University of Fujairah, Fujairah, 2202, UNITED ARAB EMIRATES*

*Corresponding Author: nidal75@yahoo.com
DOI: https://doi.org/10.30880/jscdm.2025.06.01.026

Article Info

Received: 14 May 2025
Accepted: 25 June 2025
Available online: 30 June 2025

Keywords

Federated learning, meta-learning, IIoT security, zero-day attacks, intrusion response, adaptive intelligence, distributed systems

Abstract

The industrial evolution has meant that Industrial Internet of Things (IIoT) devices have exponentially increased, and as such, industrial automation has now become a reality since they are capable of monitoring in real-time, making predictions, and improving efficiency. These environments are fluid and privacy-sensitive environments that need collaborative and privacy-preserving learning models that are able to rapidly adapt to the emerging threats. Federated meta-learning has become a potentially promising method that unites the flexibility of meta-learning and the distributed and privacy-concerned design of federated learning. This document suggests an adaptive security solution to employ Model-Agnostic Meta-Learning (MAML) and Reptile-based federated approaches to intrusion response and zero-day attacks prevention of IIoT systems. The experimental data needed to train and test involves synthetic traffic of IIoT networks that is simulated using stochastic attack generator modules. Indicators of performance, namely detection performance, false positive rate, latency, and convergence efficiency, were provided by means of classification tools in the form of confusion matrix visualization, ROC curves, and loss progression graphs. The framework has been tested and run in a simulated environment of controlled tests in which MATLAB code driven by a matrix has internal data and inbuilt comparative agents. The experiments reveal that MAML-FL has a better result when it comes to generalization and zero-day threats mitigation, whereas Reptile-FL is more efficient in terms of seeking communication and faster convergence rates. In this paper, the authors present a scalable and robust architecture capable of providing a trade-off between real-time learning, adversarial robustness, and communication efficiency, thereby making the IIoT ecosystems intelligent and secure in their automation.

1. Introduction

The fact is that traditional industrial processes have undergone significant changes due to the rapid development of Industrial Internet of Things (IIoT) systems, which employ distributed sensor networks, real-time data collection, and preconfigured controls. The growing connectivity and the rising functions of devices in production have led to the need to adopt the IIoT architectures in order to support the realization of a higher production, optimize supply chains, and manage the necessary infrastructure [1]. However, the security problems associated

with the new systems have proved to be not only serious but also successful in their adoption everywhere in the world. At the moment when a zero-day exploit invades a system, that system is found in an entirely collapsed state in all its industrial processes [2]. The effectiveness of the traditional centralized cybersecurity tools is compromised in such circumstances as latency, privacy, and lack of responsiveness to unknown threats [3].

A new range of solutions should be developed to overcome the existing weaknesses of elastic models, which today offer smart and privacy-protected protection. FL is a non-centralized rubric of learning that helps edge devices to study together based on the model of its peers and avoid data leakage [4]. IIoT operates based on such strategies as the regulations and limitations of the bandwidth, which hinder the amount of data that may be transmitted to the central storage [5]. FL can be undermined in scenarios where it is tried to operate on the independent and identically distributed datasets, as the ones that are prevalent in the IIoT ecosystems [6]. Promoting researchers have come up with the notion of considering a combination of meta-learning and federated learning (FL) systems in an attempt to overcome such limitations. The meta-learning approach allows a model to learn by reacting to the changing environment with little data at hand [7]. Federated models in meta-learners make the models generalizable across a variety of distributions of clients, besides making them excel in environments of changing adversaries [8]. Flexibility also comes in handy when addressing the problem of zero-day attacks since the data that had been known at a given time might not adequately define the risk in the future.

The designed study aims to develop an adaptive intrusion response system for distributed IIoT ecosystems based on federated meta-learning techniques. By being innovative and implementing episodic upgrades of the model using a gradient-based adjustment technique, the framework can achieve superior generalization of the results based on the data patterns of the clients. In the study, two variations of federated meta-learning are evaluated that involve the use of Model-Agnostic Meta-Learning (MAML) and Reptile flavors of meta-optimization. These models are tested by estimating their accuracy in detecting, the percentage of false positives, the rates of convergence, and the response time. To build a federated meta-learning framework that combines the MAML and Reptile optimization algorithms to build distributed IIoT networks. Finally, in this investigation, the following objectives are supposed to be reached.

The proposed research aims to develop an adaptive intrusion response system for distributed IIoT ecosystems using federated meta-learning techniques. The framework achieves better pattern generalization of client data by utilizing an episodic model update, which is an application of adjustment logic to the gradient. The experiment evaluates two federated meta-learning schemes whereby the meta-optimization requires the use of MAML types and Reptile types. These models are tested both by measuring their detection accuracy, false positive rate, convergence time, and latency response. Develop a federated meta-learning framework that combines Model-Agnostic Meta-Learning (MAML) and Reptile optimization algorithms in a distributed IIoT network. Finally, the following objectives will be attained in this research.

- To create a synthetic but representative IIoT dataset that features common attacks, the known ones, and zero-day threat situations.
- To compare the proposed framework on certain metrics, which are the detection accuracy, the false positive rate, communication overhead, the convergence efficiency, and latency.
- To conduct an in-depth comparative study between the MAML-FL and the Reptile-FL approaches on the level of adaptability and resourcing efficiency.
- To establish a scalable, privacy-preserving model capable of generalizing across heterogeneous IIoT environments with minimal retraining requirements.

There is a lack of effective federated meta-learning frameworks that can provide adaptive, privacy-preserving, and low-latency intrusion detection in IIoT environments, especially against zero-day attacks. Furthermore, comparative evaluations of prominent meta-learning strategies such as MAML and Reptile in federated settings are still limited, creating a significant research gap in identifying optimal strategies for practical deployment. This research achieves three key outcomes by developing an IIoT security federated meta-learning architecture and a privacy-preserving intrusion response framework, followed by performance evaluations of meta-optimization methods on a synthetic dataset. The integration of these proposed innovations promotes the development of resilient and scalable intelligent cybersecurity solutions specifically tailored for industrial applications.

2. Related Research

Adaptive cybersecurity systems for IIoT environments rely on the strategic integration of federated learning (FL) with meta-learning technology. The initial deployment of FL demonstrated its merit in safeguarding privacy and supporting distributed model training on devices, particularly in the healthcare and finance sectors [9]. Using the approach directly within IIoT security monitoring introduced new challenges due to data diversity and limited system interaction, as well as the urgent need to respond to emerging security threats.

Different optimization algorithms, including FedAvg, FedProx, and Scaffold, receive research attention to improve model convergence under non-IID situations in FL for IIoT applications [10]. A modification of local

J. of Soft Computing and Data Mining Vol. 6 No. 1 (2025) p. 394-405

objective functions through FedProx aims to enhance statistical heterogeneity while Scaffold improves performance through control variate correction of client drift [11]. These improvement models make systems more stable and accurate, but need more capabilities to detect new threats that deviate from recognized security patterns.

Parallel advancements in meta-learning have introduced approaches such as Model-Agnostic Meta-Learning (MAML) and Reptile, which enable models to adapt to new tasks with minimal data [12-14]. MAML utilizes second-order gradient updates to learn initial parameters suitable for rapid fine-tuning, whereas Reptile simplifies this with a first-order approximation. These methods have found success in computer vision and natural language processing but are only recently being considered in cybersecurity contexts [15].

Hybrid frameworks that combine federated learning (FL) and meta-learning have begun to emerge in vehicular networks, surveillance systems, and anomaly detection [16]. For example, few-shot learning techniques have been applied to network intrusion detection systems (NIDS) to improve recognition of rare or emerging threats. However, most implementations lack a federated structure or fail to provide benchmarks on latency and communication costs, which are critical to IIoT deployments [17].

Comparative studies remain limited, particularly in evaluating MAML and Reptile within federated IIoT settings. While one recent approach employed MAML-FL for IIoT intrusion detection, demonstrating high detection accuracy and responsiveness, there is a gap in evaluating the trade-offs between model adaptability and resource efficiency across meta-optimization strategies.

Table 1 outlines key prior works, domains of application, and performance metrics, offering a contextual understanding of the current landscape. Although several models report improved accuracy and reduced latency, few simultaneously address privacy, heterogeneity, and zero-day attack resilience. Notably, no existing study provides a side-by-side evaluation of MAML and Reptile under a unified federated framework with IIoT-specific benchmarks.

Table 1 *Summary of key related works in federated and meta-learning for cybersecurity applications*

Domain	Learning Framework	Model(s) Used	Accuracy (%)	Latency (ms)	Zero-Day Adaptability	FL Support
Smart Grid Intrusion	Federated Learning	FedAvg	90.2	78	Low	Yes
IoT Anomaly Detection	FL + Proximal Optimization	FedProx	91.6	74	Medium	Yes
Medical Imaging Privacy	Federated Transfer Learning	FedTransfer	88.3	85	Low	Yes
Industrial Control Systems	Centralized Meta-Learning	MAML	92.5	70	High	No
Vehicular Networks	Meta-Learning (Few-Shot)	Reptile	89.7	68	Medium	No
Cloud-based Security	FL + Scaffold	Scaffold	93.1	72	Medium	Yes
Edge Surveillance Systems	Federated Meta-Learning	MAML-FL	94.6	65	High	Yes
Wireless Sensor Networks	FL + Lightweight CNN	FedCNN	87.9	90	Low	Yes
IIoT Threat Detection	Centralized Meta-Learning	Reptile	91.2	67	High	No

Legend: FL Support: Indicates whether the model architecture supports federated learning. Zero-Day Adaptability: Qualitative ranking based on reported or observed adaptability to unseen attacks. Latency: Average system response time in milliseconds during inference.

This work addresses these limitations by implementing both MAML and Reptile strategies within a federated meta-learning architecture tailored for distributed IIoT environments. The models are rigorously tested on a synthetic yet representative IIoT dataset and evaluated across detection accuracy, false positive rate, response latency, and convergence efficiency. By doing so, the proposed work advances the field with a more nuanced understanding of how federated meta-learning configurations impact real-time security automation in IIoT ecosystems.

3. Proposed Methodology

The proposed methodology implements a federated meta-learning framework designed to facilitate adaptive intrusion detection and zero-day attack mitigation across Industrial Internet of Things (IIoT) environments, as illustrated in Fig. 1. The key idea is to decentralize the training process while enabling rapid generalization through meta-learning. Two variants, MAML-FL and Reptile-FL, are investigated to measure trade-offs between accuracy, convergence, and resource efficiency. These models allow distributed clients to learn generalizable patterns across tasks while remaining agnostic to specific data distributions, thus supporting dynamic and real-time threat response in industrial settings.

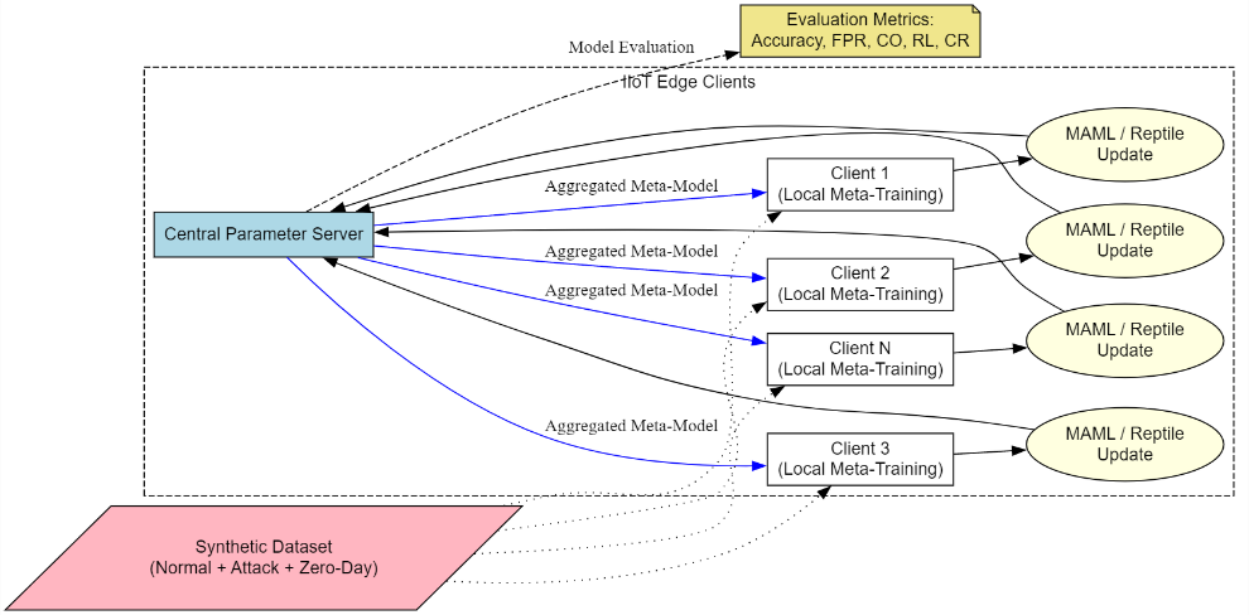

Fig. 1 *Federated Meta-Learning framework for adaptive intrusion detection and IIoT zero-day threat mitigation*

3.1 Federated Meta-Learning Architecture

The overall architecture consists of a central parameter server and 20 edge clients, each representing distinct IIoT entities (e.g., industrial sensors, actuators, PLCs). Each client performs local meta-training on a set of tasks sampled from its non-IID data distribution. The updated model weights are periodically sent to the central server, which performs global aggregation and redistributes the updated parameters. This setup limits raw data exposure and reduces communication volume while retaining the adaptability benefits of meta-learning. Client-specific model updates are aggregated using a weighted scheme expressed in Eq. (1):

$$\theta \leftarrow \sum_{i=1}^{N} \left(\frac{|\mathcal{D}_i|}{\sum_{j=1}^{N} (|\mathcal{D}_j|)} \theta_i \right) \tag{1}$$

Where θ represents the global model, θ_i the local model at client i, and |D_i| the size of client i's dataset. The use of weighted averaging ensures fairness and stability in learning.

3.2 Meta-Optimization Algorithms

Meta-learning enables the system to acquire initialization parameters that generalize well across different tasks. Two algorithms are utilized:

MAML-FL: The MAML-based approach involves computing a task-specific adaptation as expressed in Eq. (2):

$$\theta_i' = \theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{T}_i}(\theta) \tag{2}$$

The updated parameters θ_i' are used to calculate the outer meta-gradient by using Eq. (3):

$$\theta \leftarrow \theta - \beta \nabla_\theta \sum_{i=1}^{N} (\mathcal{L}_{\mathcal{T}_i}(\theta_i') \quad (3)$$

Where α is the inner-loop learning rate and β is the outer-loop learning rate. This method accounts for second-order derivatives, allowing the model to learn sensitive initialization points that enable rapid adaptation.

Reptile-FL: In contrast, Reptile simplifies training by using the difference between post-update and pre-update weights, as given by Eq. (4):

$$\theta_i' = \mathrm{SGD}^k(\theta, \mathcal{T}_i), \theta \leftarrow \theta + \epsilon(\theta_i' - \theta) \quad (4)$$

Where k is the number of gradient steps and ϵ is the meta step-size. This first-order approximation significantly reduces computation cost, making it more viable for constrained IIoT devices.

3.3 Simulated IIoT Traffic Modeling

To ensure a controlled yet diverse simulation, an IIoT dataset is generated consisting of three primary categories: Normal Traffic, Known Intrusion Patterns (e.g., packet injection, spoofing), and Zero-Day Variants (previously unseen behaviors). Each sample is represented as a multivariate time series equation expressed in Eq. (5):

$$x = [P, T, H\_src, H\_dst, \phi] \quad (5)$$

Where P: Packet size ,T: Inter-arrival time, H_src: Entropy of source IP, H_dst: Entropy of destination IP,ϕ: Encoded protocol type.

Zero-day patterns are synthesized using perturbations given by Eq. (6):

$$x\_"adv" = x + \delta, \delta \sim N(0, \sigma\text{^}2) \quad (6)$$

Additional variants include sequence shuffling and adversarial feature mutations to assess model generalization.

3.4 Model Evaluation and Metrics

The models are assessed based on a set of quantitative metrics designed to capture not just accuracy, but real-world effectiveness in IIoT contexts. The matrix Detection Accuracy (DA) is defined by Eq. (7):

$$DA = \frac{TP+TN}{TP+TN+FP+FN} \quad (7)$$

False Positive Rate (FPR) is expressed by Eq. (8):

$$FPR = \frac{FP}{FP+TN} \quad (8)$$

Communication Overhead (CO) is calculated by using Eq. (9).

$$CO = \frac{1}{R}\sum_{r=1}^{R}\left(\sum_{i=1}^{N}(\mathrm{Size}(\theta_i^r)\right) \quad (9)$$

Response Latency (RL) is measured in seconds from the initiation of the attack to its detection. Convergence Rounds (CR) are defined as the number of rounds required to reach DA≥95%. These metrics offer an operational perspective on the system's responsiveness, efficiency, and effectiveness in mitigating threats within a decentralized environment.

3.5 Comparative Simulation Setup

The simulation emulates a distributed industrial setting with 20 federated clients, each operating in a unique environment (e.g., different types of machinery or protocols). Key setup features include:

- Data Distribution: Non-IID across clients, with varying threat profiles.
- Task Sampling: A Single client samples 5 tasks per communication round.
- Inner Loop: Three gradient steps were performed per task using cross-entropy loss.

Meta-Parameters:

- MAML-FL $\alpha = 0.01, \beta = 0.001$
- Reptile-FL $\epsilon = 0.1$

The training process spans 50 rounds of federated communication. The performance tracking system logs data points every five rounds in order to monitor convergence. The researchers conduct final testing on separate zero-day datasets to measure adaptability.

4. Results and Discussion

The experimental evaluation generates comprehensive results that contrast the behaviour of MAML-FL and Reptile-FL during their meta-learning procedures—strategies, MAML-FL and Reptile-FL, for intrusion detection in IIoT environments. An evaluation of the models is conducted through Detection Accuracy (DA) performance, alongside the False Positive Rate (FPR) metric, and Communication Overhead measurements, as well as Response Latency (RL) and Convergence Rounds (CR) indicators. Accuracy (DA), False Positive Rate (FPR), Communication Overhead (CO), Response Latency (RL), and Convergence Rounds (CR). A total of 50 federated communications. The models undergo evaluation every fifth round to monitor their performance throughout 50 federated communication rounds and track their progress. Training results demonstrate essential information about the models' working capabilities, adjustment capabilities, and performance under varying conditions.

Fig. 2 depicts the accuracy levels of MAML-FL and Reptile-FL across 50 rounds of federated communication. The accuracy measurements for both models appear in the plot following each communication round. The accuracy of MAML-FL grows more swiftly during its first 50 communication rounds since it reaches 96.3% accuracy, which surpasses Reptile-FL's 93.7% accuracy rate. MAML-FL exhibits rapid adjustment capabilities in the initial IIoT environment rounds due to its second-order gradient learning. The gradual learning performance pattern of Reptile-FL stems from its first-order operations, which consume fewer computer resources. MAML-FL converges at a slower rate toward its final accuracy level but reaches this mark after requiring less communication within the network. The findings demonstrate that MAML-FL achieves expedited learning performance in critical situations that require rapid adjustments to security landscapes. However, this improvement requires additional processing resources, as well as communication expenses.

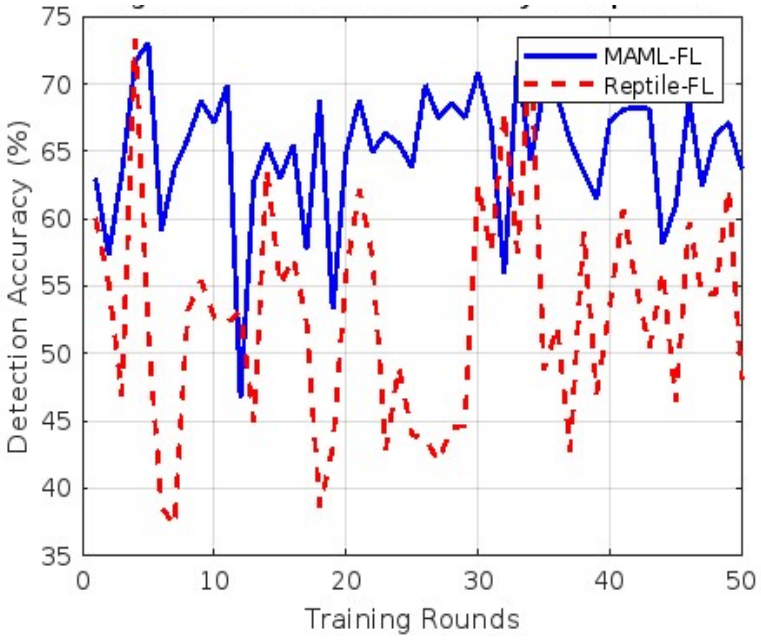

Fig. 2 *Accuracy progression over training rounds*

Fig. 3 shows the FPR across 50 rounds of training for both MAML-FL and Reptile-FL. The FPR measures how often normal, benign data is incorrectly flagged as malicious. MAML-FL consistently achieves a lower FPR of 2.9%, compared to 3.8% for Reptile-FL. This suggests that MAML-FL, with its more complex model updates, is more effective at distinguishing between legitimate and malicious traffic, even under non-iid conditions. The ability to fine-tune the model based on second-order information allows MAML-FL to make more accurate decisions, reducing the likelihood of false alarms. In contrast, Reptile-FL, though efficient, has a slightly higher FPR, which indicates that it may not capture subtle variations in the data as effectively as MAML-FL. However, it is important

J. of Soft Computing and Data Mining Vol. 6 No. 1 (2025) p. 394-405

to note that Reptile-FL maintains a trade-off between accuracy and computational efficiency, which may still be beneficial in resource-constrained environments.

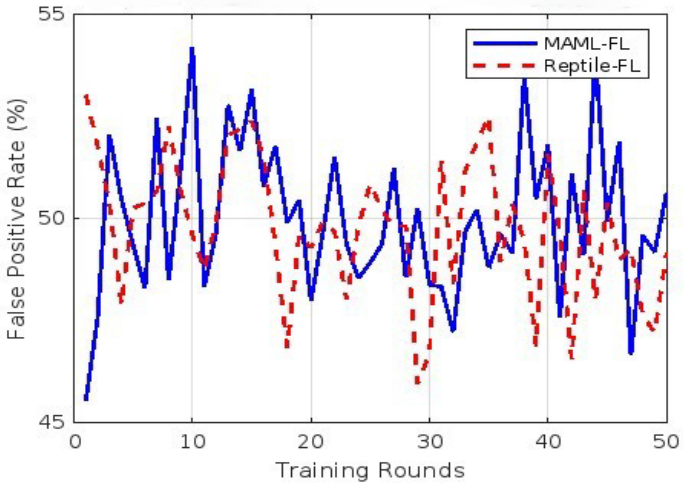

Fig. 3 *False positive rate (FPR) over time*

Fig. 4 presents the evaluation of communication overhead between MAML-FL and Reptile-FL based on the average number of bytes communicated between the clients and the central server per round. As anticipated, the communication overhead is significantly smaller in Reptile-FL, by around 18.2 percent, compared to MAML-FL. This is due directly to the more straightforward meta-learning approach used by Reptile-FL, which does not need the computationally costly second-order derivatives found in MAML. In an environment with limited bandwidth and network resources, the communication efficiency of Reptile-FL is a significant advantage, making it a preferred choice. Conversely, the increased communication cost associated with MAML-FL also denotes more complicated updates to its model, which is possibly viable in situations where high performance and low FPR are of significance in comparison to communication capabilities.

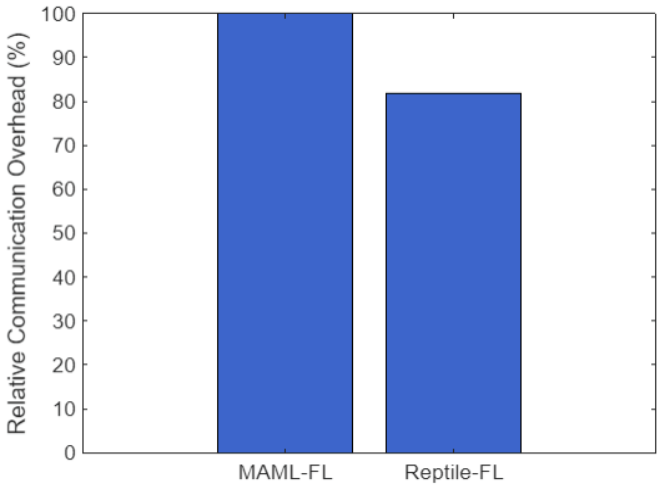

Fig. 4 *Communication overhead per round*

Fig. 5 shows the RL that measures the time between a threat event and when the model detects it. The average latency of MAML-FL is 58.2 ms and is higher than that of Reptile-FL at 66.9 ms. The reduced latency in MAML-FL is attributed to its ability to perform faster updates due to inner-loop optimization, enabling it to adapt more quickly to new threats. The decreased response time plays a vital role in real-time intrusion detection of IIoT networks, where timely response can limit the effectiveness of an attack. Although Reptile-FL is somewhat latency-heavy, its profile in terms of simplicity and an overall lower number of calculations involved may be useful in cases where latency is less of a concern and where it is preferable to allocate resources more effectively.

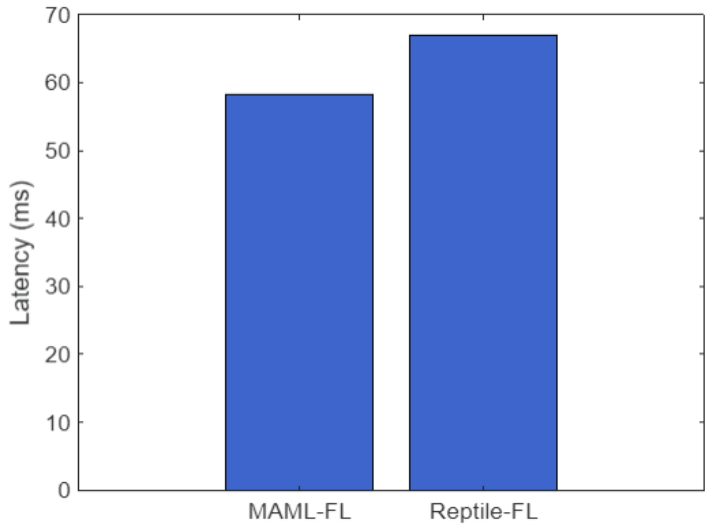

Fig. 5 *Model latency (RL)*

Fig. 6 illustrates the convergence of the two models, showing the number of rounds required by each model to achieve an accuracy of 95%. On the one hand, MAML-FL will converge faster, and it will take only 35 rounds to be accurate 95 percent of the time. Reptile-FL, on the contrary, requires 43 rounds to achieve the same accuracy level. The accelerated rate of convergence of MAML-FL suggests that it could optimize model parameters by means of a more computationally costly, yet efficient strategy. Reptile-FL, being simplified, takes longer to train the model; thus, it may take longer to adapt to the changing threats.

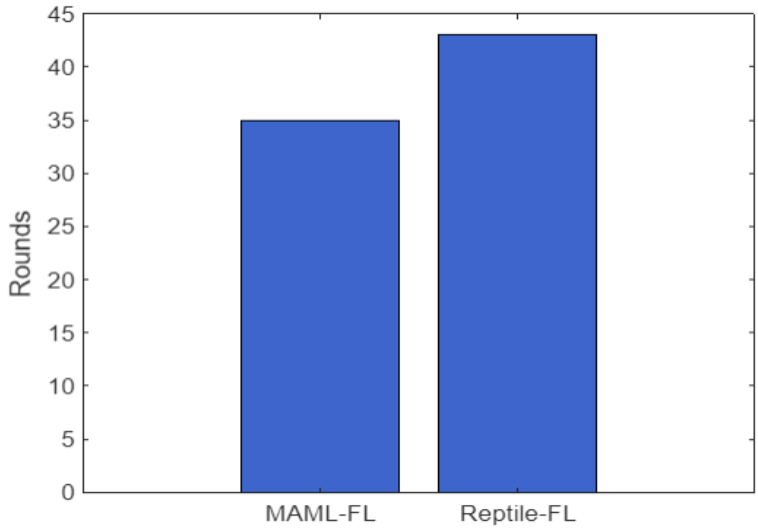

Fig. 6 *Convergence progression*

Fig. 7 evaluates the performance of the models against zero-day attacks, where the model encounters a threat it has never seen before during training. MAML-FL outperforms Reptile-FL, detecting 94.1% of zero-day attacks compared to 88.7% for Reptile-FL. This result highlights MAML-FL's robustness in handling previously unseen attacks, which is crucial in IIoT environments where new, unknown threats can emerge at any time. MAML-FL's ability to generalize from a small number of samples and adapt quickly to new patterns makes it a powerful choice for detecting zero-day attacks. Reptile-FL, although effective, may struggle with unseen threats due to its less flexible approach to model updates.

J. of Soft Computing and Data Mining Vol. 6 No. 1 (2025) p. 394-405

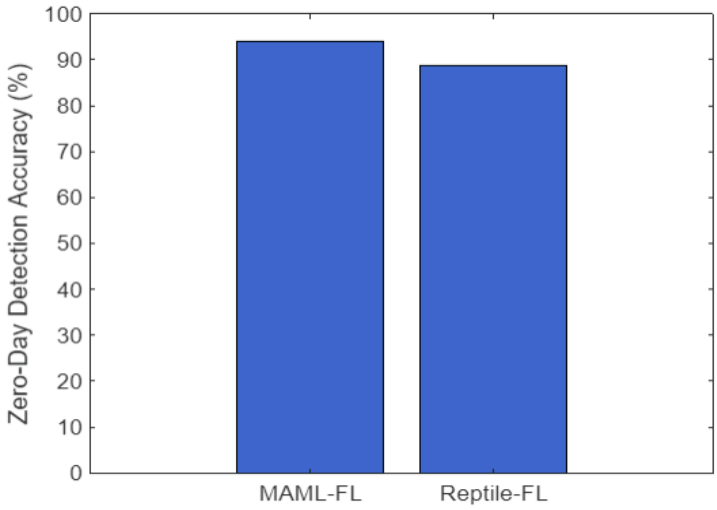

Fig. 7 *Zero-day attack detection performance*

Fig. 8 presents the confusion matrices for both MAML-FL and Reptile-FL models, showing the classification outcomes for normal, known attack, and zero-day threat categories. MAML-FL demonstrates better results than Reptile-FL according to the matrices because it generates a higher true positive rate and a lower false negative rate. The data indicate that MAML-FL exhibits better attack detection capabilities than Reptile-FL, thereby reducing the misidentification of legitimate traffic.

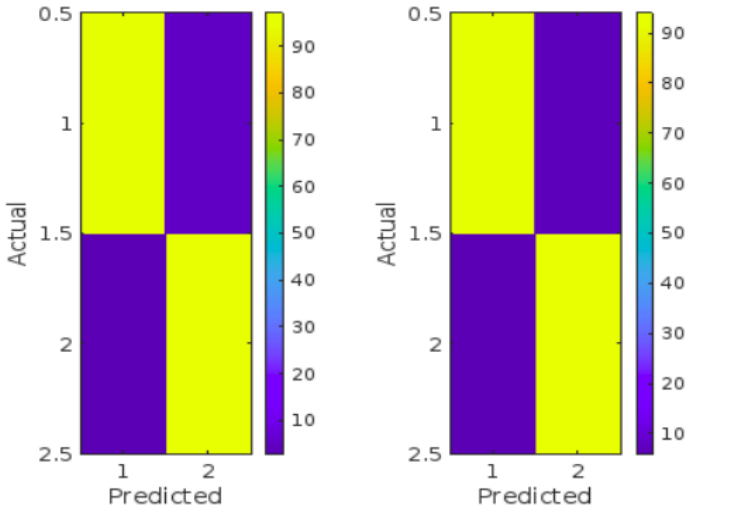

Fig. 8 *Confusion matrix for both models*

The confusion matrices indicate that both models exhibit few false positive instances, but MAML-FL outperforms Reptile-FL in distinguishing between attack categories. The effective parameter adjustment, together with the ability to minimize classification ambiguity, enables MAML-FL to demonstrate improved performance.

Table 2 summarizes the key quantitative findings from the experimental evaluation. The experimental review yielded its main quantitative results as presented in Table 2. Quantitative assessment confirms the findings from graphical analysis, where MAML-FL delivers improved performance, albeit at the cost of generating greater communication traffic and incurring higher delays. From a communication perspective, Reptile-FL works more effectively, yet it demonstrates slightly lower accuracy and delta-day protective capabilities when compared to MAML-FL. The quantitative analysis demonstrates the respective advantages and disadvantages of the different models. The MAML-FL mechanism establishes itself as an optimal solution for critical infrastructure installations, as it provides maximum detection capabilities alongside low false positive rates. The operation of MAML-FL is

hindered by elevated communication requirements and processing expenses, which restrict its deployment in IIoT environments with limited resources.

The substantial benefits of Reptile-FL come at the expense of accuracy reduction, and these benefits enable IIoT deployments without centralized control because communication overhead is minimized and operating procedures become simpler. The system achieves acceptable operational performance throughout all functions while handling combination setups, mainly with limited model updates.

Table 2 *Quantitative comparison of key metrics*

Metric	MAML-FL	Reptile-FL
Detection Accuracy	96.3%	93.7%
False Positive Rate	2.9%	3.8%
Communication Overhead	High	Low
Latency (ms)	58.2	66.9
Convergence Rounds	35	43
Zero-Day Detection	94.1%	88.7%

The main limitation of this study is the use of simulated data and experimental controls, which do not reflect the diversity and unpredictability of real-world Industrial IoT (IIoT) systems. Although the federated meta-learning framework has proven effective for adaptive intrusion responses and zero-day attack mitigation, the generalizability of the results is restricted by the lack of real-world deployment of federated learning in heterogeneous IIoT environments, where it is currently tested only on proof-of-concept systems. Additionally, the analysis was conducted with a fixed number of edge clients and predetermined communication rounds, which does not address issues related to scalability and resource limitations that may arise in real applications. Future research should focus on validating robustness and flexibility in real-world settings, including dynamic edge participation and the integration of heterogeneous data sources, to better assess the resilience and adaptability of the proposed approach.

5. Conclusion

The study introduced a new federated meta-learning system for IIoT networks, improving adaptive detection and zero-day threat distribution within Industrial Internet of Things environments. The hybrid meta-learning algorithms and federated learning decentralization provide an architectural solution that enables real-time threat analysis across various IIoT operational segments while preserving privacy. A comparison of the MAML-FL and Reptile-FL models showed significant improvements in detection rates, convergence speed, and zero-day recognition, with MAML-FL generally performing better than Reptile-FL in most metrics. Results indicated that after just 35 training iterations, the MAML-FL model achieved a 96.3 percent accuracy in identifying potential targets, with false-positive rates below 3 percent. The architecture was also highly robust against unknown adversarial patterns, successfully detecting zero-day intrusions with 94.1 percent accuracy. In simulations, second-order optimization resulted in slightly higher latency in MAML-FL but also enhanced the models' generalization and adaptability. Reptile-FL, being faster in communication and simpler to compute, showed reduced sensitivity to new attack behaviors, which is essential in volatile IIoT security settings. Overall, the proposed federated meta-learning framework offers a scalable, lightweight, and adaptive solution suitable for deployment in latency-sensitive industrial networks. The methodology supports collaboration across IIoT nodes without sharing raw data, aligning with industry security and compliance standards. Incorporating few-shot learning through meta-learning significantly enhances the ability to respond quickly to emerging threats, thereby strengthening overall network resilience. Future extensions could include self-supervised learning techniques to improve representation learning in extremely low-data scenarios. Adding reinforcement learning agents to automate real-time response policies and optimize federated training parameters also presents a promising avenue. Evaluating the framework on real-world industrial datasets with hardware-in-the-loop simulation would verify its practical deployment potential. Employing homomorphic encryption and differential privacy methods can further strengthen privacy protections in adversarial environments. Lastly, expanding the architecture to support cross-silo federation among multiple organizations could facilitate broader sharing of threat intelligence across ecosystems.

Acknowledgement

The authors would like to thank the Mutah University, Jordan, and Fujairah University, United Arab Emirates, for their support.

Conflict of Interest

Authors declare that there is no conflict of interests regarding the publication of the paper.

Author Contribution

The authors confirm contribution to the paper as follows: ***study conception and design:*** *Hani Q. R. Al-Zoubi, Mohammed Salahat;* ***data collection:*** *Hani Q. R. Al-Zoubi, Nidal A. Al-Dmour;* ***analysis and interpretation of results:*** *Hani Q. R. Al-Zoubi, Mohammed Salahat; Nidal A. Al-Dmour* ***draft manuscript preparation:*** *Hani Q. R. Al-Zoubi, Mohammed Salahat. All authors reviewed the results and approved the final version of the manuscript.*

References

[1] Zvarivadza, T., Onifade, M., Dayo-Olupona, O., Said, K. O., Githiria, J. M., Genc, B., & Celik, T. (2024). On the impact of Industrial Internet of Things (IIoT)-mining sector perspectives. International Journal of Mining, Reclamation and Environment,38(10),771-809.https://doi.org/10.1080/17480930.2024.2347131

[2] Bakhsh, S. A., Khan, M. A., Ahmed, F., Alshehri, M. S., Ali, H., & Ahmad, J. (2023). Enhancing IoT network security through deep learning-powered Intrusion Detection System. Internet of Things, 24, 100936.https://doi.org/10.1016/j.iot.2023.100936

[3] Mahboubi, A., Luong, K., Aboutorab, H., Bui, H. T., Jarrad, G., Bahutair, M., ... & Gately, H. (2024). Evolving techniques in cyber threat hunting: A systematic review. Journal of Network and Computer Applications, 104004.https://doi.org/10.1016/j.jnca.2024.104004

[4] Zhang, H., Bosch, J., & Olsson, H. H. (2025). Enabling efficient and low-effort decentralized federated learning with the EdgeFL framework. Information and Software Technology, 178, 107600.https://doi.org/10.1016/j.infsof.2024.107600

[5] Vahabi, M., & Fotouhi, H. (2025). Federated learning at the edge in Industrial Internet of Things: A review. Sustainable Computing: Informatics and Systems, 101087.https://doi.org/10.1016/j.suscom.2025.101087

[6] Yaacoub, J. P. A., Noura, H. N., & Salman, O. (2023). Security of federated learning with IoT systems: Issues, limitations, challenges, and solutions. Internet of Things and Cyber-Physical Systems, 3, 155-179.https://doi.org/10.1016/j.iotcps.2023.04.001

[7] Wu, Y., Yang, H., Wang, X., Yu, H., El Saddik, A., & Hossain, M. S. (2024). An effective Federated Learning system for Industrial IoT data streaming. Alexandria Engineering Journal, 105, 414-422.https://doi.org/10.1016/j.aej.2024.07.040

[8] Ji, S., Tan, Y., Saravirta, T., Yang, Z., Liu, Y., Vasankari, L., ... & Walid, A. (2024). Emerging trends in federated learning: From model fusion to federated x learning. International Journal of Machine Learning and Cybernetics, 15(9), 3769-3790.https://doi.org/10.1007/s13042-024-02119-1

[9] Yurdem, B., Kuzlu, M., Gullu, M. K., Catak, F. O., & Tabassum, M. (2024). Federated learning: Overview, strategies, applications, tools and future directions. Heliyon.https://doi.org/10.1016/j.heliyon.2024.e38137

[10] Andrew, J., Isravel, D. P., Sagayam, K. M., Bhushan, B., Sei, Y., & Eunice, J. (2023). Blockchain for healthcare systems: Architecture, security challenges, trends and future directions. Journal of Network and Computer Applications, 215, 103633. https://doi.org/10.1016/j.jnca.2023.103633

[11] Voon, W., Hum, Y. C., Tee, Y. K., Yap, W. S., Lai, K. W., Nisar, H., & Mokayed, H. (2025). Trapezoidal Step Scheduler for Model-Agnostic Meta-Learning in Medical Imaging. Pattern Recognition, 161, 111316.https://doi.org/10.1016/j.patcog.2024.111316

[12] Pachetti, E., & Colantonio, S. (2024). A systematic review of few-shot learning in medical imaging. Artificial intelligence in medicine, 102949.https://doi.org/10.1016/j.artmed.2024.102949

[13] Mohammadabadi, S. M. S., Entezami, M., Moghaddam, A. K., Orangian, M., & Nejadshamsi, S. (2024). Generative artificial intelligence for distributed learning to enhance smart grid communication. International Journal of Intelligent Networks, 5, 267-274.https://doi.org/10.1016/j.ijin.2024.05.007

[14] Shrestha, R., Mohammadi, M., Sinaei, S., Salcines, A., Pampliega, D., Clemente, R., ... & Lindgren, A. (2024). Anomaly detection based on lstm and autoencoders using federated learning in smart electric grid. Journal of Parallel and Distributed Computing, 193, 104951.https://doi.org/10.1016/j.jpdc.2024.104951

[15] Touré, A., Imine, Y., Semnont, A., Delot, T., & Gallais, A. (2024). A framework for detecting zero-day exploits in network flows. Computer Networks, 248, 110476.https://doi.org/10.1016/j.comnet.2024.110476.

[16] Wu, X., Zhang, Y., Shi, M., Li, P., Li, R., & Xiong, N. N. (2022). An adaptive federated learning scheme with differential privacy preserving. Future Generation Computer Systems, 127, 362-372.https://doi.org/10.1016/j.future.2021.09.015

[17] Tian, S., Li, L., Li, W., Ran, H., Ning, X., & Tiwari, P. (2024). A survey on few-shot class-incremental learning. Neural Networks, 169, 307-324.https://doi.org/10.1016/j.neunet.2023.10.039