GAIA-FL: Generative AI-Augmented Federated Learning for Intrusion Detection System

Ons Aouedi[+], Alexandre Boissel[*], and Kandaraj Piamrat[*]
[+]*SnT, SIGCOM, University of Luxembourg, Luxembourg*
ons.aouedi@uni.lu
[*]*Nantes Université, École Centrale Nantes, CNRS, INRIA, LS2N, UMR 6004, F-44000 Nantes, France*
firstname.lastname@ls2n.fr

***Abstract*—In recent years, federated learning (FL) has gained significant attention as a privacy-preserving solution for distributed machine learning, particularly in cybersecurity applications such as intrusion and attack detection. However, traditional FL models often face challenges related to limited training data diversity, communication overhead, and the ability to adapt to novel or unforeseen attack patterns. At the same time, generative AI (GAI) models have emerged as powerful tools for synthesizing realistic data, enabling enhanced model generalization and robustness. In this work, we propose integrating GAI with FL to create a more effective and adaptive framework for cybersecurity called GAIA-FL (GAI-Augmented FL). GAI can augment FL by synthesizing diverse attack scenarios, enriching local datasets, and addressing data heterogeneity across distributed nodes. We analyze the unique capabilities of GAI, such as data generation, and highlight its potential to improve the performance of FL-based cybersecurity systems. Additionally, we explore the integration of generative models and FL, focusing on their combined ability to detect complex and evolving threats while maintaining data privacy. Unlike existing studies, our work emphasizes the fusion of GAI and FL to tackle the challenges of decentralized intrusion detection and attack prevention. To validate our approach, we present a case study where GAI is used to enhance FL-based network intrusion detection by generating synthetic attack data, improving detection accuracy and robustness. This work demonstrates how this integration can revolutionize cybersecurity in next-generation networks by providing scalable, privacy-preserving, and adaptive solutions for evolving cyber threats.**

***Index Terms*—Federated Learning, Generative AI, Deep Learning, Intrusion Detection System, Cybersecurity.**

I. Introduction

In 2022, the market for GAI in IoT was valued at USD 947.8 million. It is expected to experience substantial growth and is predicted to reach USD 8,952.6 million by 2032 [1]. Experts have projected that the Compound Annual Growth Rate (CAGR) will achieve an impressive 25.9%.[1] Generative AI (GAI), with its ability to generate and synthesize realistic data and model complex patterns, has emerged as a critical enabler for applications requiring adaptability, scalability, and intelligence. In the Internet of Things (IoT) ecosystems, where devices generate diverse and incomplete data, GAI addresses key challenges such as data heterogeneity, scarcity, and quality [2], [3]. For example, GAI can produce synthetic datasets to improve training efficiency and enhance the robustness of AI models. Its application extends to network optimization, anomaly detection, and predictive maintenance, making it a vital tool for advancing IoT and 6G networks.

Simultaneously, Federated Learning (FL) has gained recognition as a privacy-preserving distributed learning paradigm [4]. By enabling collaborative model training across decentralized nodes without sharing raw data, FL maintains data privacy and reduces the risk of data breaches. However, FL suffers from critical limitations, including limited data diversity across nodes, communication overhead, and difficulty in adapting to dynamic scenarios [5]. These limitations hinder its application in complex domains such as intrusion detection and cybersecurity. Therefore, integrating GAI with FL presents a promising direction to overcome these limitations. GAI can synthesize diverse, high-quality datasets locally, enriching training data for FL and addressing data imbalance and scarcity. In cybersecurity, this integration is particularly valuable for intrusion detection systems (IDS), enabling FL models to train on realistic and diverse attack scenarios, thereby improving detection accuracy and adaptability [6]. Moreover, GAI enhances FL's ability to adapt to novel and evolving threats by simulating sophisticated cyberattacks. As we progress toward 6G networks, the complexity and scale of IoT ecosystems are expected to expand significantly, with trillions of connected devices operating across terrestrial and non-terrestrial networks [7]. In this context, cybersecurity becomes a cornerstone of network reliability, and the fusion of GAI and FL provides a scalable, privacy-preserving, and adaptive solution to address these challenges. By combining the generative capabilities of GAI with the decentralized and secure learning of FL, IDS can be equipped to handle emerging threats effectively while ensuring compliance with stringent privacy requirements.

This paper explores the potential of integrating GAI with FL in IDS for IoT and 6G networks and propose GAIA-FL. By leveraging these technologies, we aim to pave the way for intelligent, secure, and resilient networks capable of adapting to evolving cyber threats. To the best of our knowledge, this paper is among the first to explore the integration of GAI with FL specifically for IDS in IoT.

[1]https://marketresearch.biz/report/generative-ai-in-iot-market/

II. Challenges in Federated Learning for Intrusion Detection

FL has gained significant attention for its potential to enable distributed learning while preserving data privacy. However, using FL for IDS presents several challenges that must be addressed to ensure effective and robust performance in real-world scenarios.

A. Non-IID and Imbalance Data

One of the most significant challenges in applying FL for IDS is the non-IID (non-independent and identically distributed) nature of data across devices. Unlike centralized systems, where training data can be balanced and curated, FL relies on data stored locally on devices, which often differs in terms of distribution, sample sizes, and attack diversity [8]. For instance, some devices may primarily encounter Denial of Service (DoS) attacks, while others might log phishing attempts or privilege escalation incidents. This heterogeneity can result in a global model that overfits frequently encountered attack types but performs poorly on rare or unseen attacks. Moreover, the data imbalance is a critical issue. Devices with fewer data samples or less diverse datasets may contribute less to the global model, causing their unique attack patterns to be underrepresented. In extreme cases, these devices may even fail to participate in the training process due to insufficient data, reducing the overall effectiveness of the system. Advanced aggregation techniques such as weighted averaging or data augmentation strategies could help mitigate these issues.

B. Privacy and Security Risks

While FL is designed to enhance privacy by ensuring that raw data never leaves local devices, the system is not immune to privacy and security vulnerabilities. Malicious participants can exploit the FL process to conduct data poisoning or backdoor attacks where they inject compromised updates into the training pipeline. These attacks can lead to a global model that misclassifies certain inputs, undermining the IDS's ability to detect intrusions effectively. Additionally, advanced adversarial techniques like model inversion and gradient leakage allow attackers to infer sensitive information from shared model updates. For instance, an attacker could reconstruct details about the training data by analyzing gradients shared during the FL process. These vulnerabilities directly challenge the core premise of FL as a privacy-preserving framework. To counter these risks, robust defense mechanisms are required. Techniques such as secure aggregation, differential privacy, and adversarial training can help mitigate these vulnerabilities. Secure aggregation ensures that individual updates are encrypted during transmission, preventing attackers from accessing raw gradients. Differential privacy introduces noise into the updates, reducing the likelihood of sensitive data reconstruction. However, implementing these measures introduces additional computational and communication costs, which must be optimized for large-scale deployments.

C. Limited adaptability to unseen attack scenarios

FL-based IDS systems often struggle with limited adaptability to novel and unseen attack scenarios, such as zero-day exploits or advanced persistent threats (APTs). These types of attacks deviate significantly from known patterns and signatures, making them particularly challenging to detect with models trained on historical data. In an FL setup, each device trains locally on data specific to its environment, which may reflect only a narrow subset of the global attack landscape. Consequently, the global model tends to overfit to frequently observed attack types, leaving it poorly equipped to generalize to emerging or unseen threats. The issue is further exacerbated by the reliance on static datasets that fail to capture the dynamic and evolving nature of cybersecurity threats. For example, a dataset collected from IoT devices last year might not contain samples of newly developed malware or exploit techniques. Without access to diverse and up-to-date attack data, FL-based IDS systems risk becoming obsolete in the face of rapidly changing cyberattack vectors. Addressing this challenge requires innovative strategies such as continual learning, where the model is periodically updated with new attack data, and adversarial training, where synthetic attack samples are generated to simulate emerging threats. Another promising avenue is to incorporate meta-learning techniques, enabling models to adapt quickly to new environments with minimal training data

III. GAIA-FL

A. Generative AI for FL in IDS

GAI has revolutionized machine learning by introducing models capable of creating realistic and diverse data. Techniques like Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and Diffusion Models (DMs) have proven invaluable in addressing data scarcity and heterogeneity challenges across various domains [3]. In the context of FL-based IDS, GAI plays a pivotal role in enriching local datasets, improving model generalization, and enhancing the adaptability of FL frameworks. A key strength of GAI is its ability to synthesize realistic attack scenarios that go beyond the patterns observed in historical data. By generating diverse datasets that include unseen attacks, generative models can compensate for the limited representation of such attacks in localized FL training environments. Additionally, GAI enhances data diversity by mitigating the effects of non-IID data across FL devices. This is particularly crucial in distributed systems where localized datasets often lack variety and fail to represent global attack patterns. By generating synthetic data that captures complex attack behaviors, GAI ensures that the global FL model benefits from a more balanced and representative training set. Beyond data generation, GAI can also be used to simulate future network behaviors, predict potential vulnerabilities, and provide proactive insights for IDS. This predictive capability makes GAI an invaluable tool for strengthening FL-based IDS, ensuring the system remains robust and adaptive in dynamic environments.

TABLE I
COMPARISON OF GENERATIVE AI-ENHANCED FL MODELS AND TRADITIONAL FL MODELS FOR IDS

Models	Benefits	Concerns
Generative AI-enhanced FL	• **Enhanced data quality**: Generative AI synthesizes realistic attack scenarios, enriching training datasets and addressing biases caused by non-IID data. • **Improved adaptability**: By simulating novel and rare attack scenarios (e.g., zero-day threats), generative AI enhances FL's ability to adapt to evolving threats. • **Privacy reinforcement**: Synthetic data enables FL systems to avoid sharing sensitive or raw data, further enhancing privacy beyond standard FL practices. • **Reduced reliance on communication**: Local data augmentation reduces the frequency and volume of global model updates, saving bandwidth and improving efficiency. • **Detection of rare attacks**: Synthetic data augmentation enables better representation of low-frequency but high-impact threats, strengthening the global model's anomaly detection capabilities.	• **Computational burden**: Training generative models alongside FL models increases computation and memory usage, particularly in resource-constrained environments. • **Security risks**: Poisoning of synthetic data can compromise the reliability of both local and global models. • **Integration complexity**: Combining generative AI with FL requires significant expertise, increasing the complexity of deployment and maintenance.
Traditional FL models	• **Simpler implementation**: Traditional FL is easier to deploy, requiring less infrastructure and fewer specialized tools. • **Low computational requirements**: Compared to generative AI, FL models require fewer resources, making them better suited for deployment in constrained environments like IoT nodes. • **Privacy-preserving design**: By ensuring raw data remains local, FL reduces the risks of data breaches during training.	• **Limited adaptability**: Models trained on static and localized data struggle to handle novel and rare attack patterns. • **Data biases**: Localized, non-IID datasets lead to biased global models that perform poorly on diverse or uncommon attack patterns. • **Communication bottlenecks**: Frequent global model updates create bandwidth and latency challenges, particularly in large-scale deployments. • **Limited data diversity**: Without synthetic data generation, the global model relies solely on local datasets, which may lack coverage of critical attack scenarios.

B. FL for Generative AI in IDS

FL provides a privacy-preserving framework for training machine learning models across distributed datasets without centralizing raw data. While FL is traditionally used for collaborative learning, it can significantly enhance the development of GAI models, particularly in the context of IDS. One of the key challenges in training GAI models, such as GANs, DMs, and VAEs, is acquiring sufficient high-quality and diverse data. In traditional centralized training, collecting this data involves privacy risks, latency, and communication overhead, especially in sensitive applications like intrusion detection. FL mitigates this challenge by allowing distributed devices to collaboratively train generative models while preserving the privacy of raw data. By sharing model updates rather than the data itself, FL enables the development of robust generative models that learn from diverse, decentralized datasets. Moreover, FL enhances GAI by reducing the risk of model bias. Since generative models often rely heavily on the quality and diversity of training data, leveraging FL ensures exposure to a broader range of attack scenarios distributed across nodes. This improves the ability of GAI to create realistic and diverse synthetic data that reflect a wide array of cyber threats. In addition, FL provides a scalable framework for training GAI models in resource-constrained environments. Devices with limited computational power can participate in the collaborative training of lightweight generative models, enabling their deployment in environments like IoT or edge devices. By combining FL with GAI, IDS frameworks can benefit from synthetic data generation at scale, ensuring enhanced generalization, privacy preservation, and adaptability to evolving threats.

IV. CASE STUDY AND RESULTS

A. Methodology

Our methodology consists of two main components: using a Deep Neural Network (DNN) with FL for IDS tasks and employing a DM with FL for generating synthetic data locally at each device. DM, specifically Latent DM (LDM), is a type of generative model that predicts and removes noise from input features. In particular, LDM is advanced generative models that operate in a compressed latent space to enhance computational efficiency while maintaining high-quality generative capabilities. Using a pre-trained encoder, the input data (e.g., images or time series) is transformed into a lower-dimensional latent representation, where the diffusion process adds Gaussian noise over multiple timesteps. The model is then trained to reverse this process by iteratively removing noise using a neural network, ultimately reconstructing the original latent representation. The reconstructed latent data is decoded back into the original data space using the corresponding decoder. By working in the latent space, LDM significantly reduce the computational demands of traditional diffusion models

while preserving critical data features. Their training objective minimizes the prediction error of the noise added during the forward process, ensuring robust denoising.

Our proposition offers several advantages, including robustness to high data heterogeneity and substantial performance improvements. The process of our proposition involves multiple rounds and globally follows the phases below:

- **Step 1:** The FL server initializes the global DNN model for IDS and disseminates it to all participating clients (IoT devices). This initialization allows devices to start the local training process with a pre-defined global model structure.
- **Step 2:** Each client trains the global DNN model locally using its initial dataset. The training process leverages the unique, possibly non-IID, data available at each client to personalize the model for the specific environment of the device.
- **Step 3:** Once local training is complete, the clients upload their locally updated DNN model weights to the FL server. This step enables collaborative learning by aggregating updates from all clients to improve the global model.
- **Step 4:** To address the limitations of non-IID data distributions, each client uses its local dataset to train the autoencoder component of the Latent Diffusion Model (LDM). The LDM operates in a lower-dimensional latent space, ensuring efficient representation and generation of synthetic data.
- **Step 5:** After local training, the clients upload their trained autoencoder models to the FL server. The server aggregates these models to construct a global LDM, enhancing the consistency and quality of synthetic data generation across all clients.
- **Step 6:** The global LDM, generated at the FL server, is sent back to the clients. Each client uses the global LDM to generate augmented datasets locally, addressing data imbalance and diversity issues caused by non-IID distributions. These augmented datasets enrich the training process for the DNN model.
- **Step 7-8:** The process of training the DNN model with both the initial and augmented datasets, followed by aggregation and LDM updates, is iteratively repeated over multiple communication rounds. As seen in the combined loop in Fig. 1, this iterative process ensures continuous improvement in the global DNN model's ability to detect intrusions by leveraging the combined power of real and synthetic data.
- **Step 9** The FL server aggregates both models received from the different clients. For the aggregation, the FederatedAveraging (FedAvg) algorithm is used.

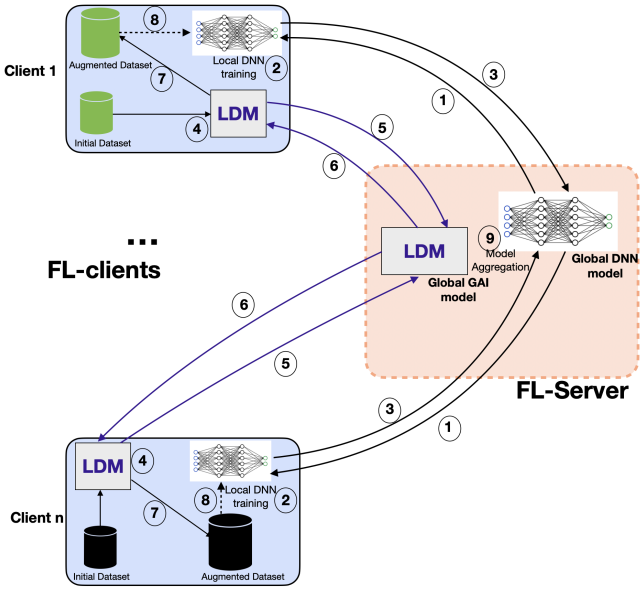

Fig. 1. The process of our GAI-FL for IDS

B. Experiment Setup

Experiments were conducted using the CIC-IDS2017 dataset, which aims to provide realistic attack samples and up-to-date attack categories [9]. The dataset was preprocessed following the steps outlined in [10]. 10% of the dataset was reserved as a test set through a random split, and the remaining data was randomly distributed among 100 client datasets constructed to produce non-iid data distributions. For each category, we determine the proportion of total samples to be attributed to each of the participating devices' datasets by sampling from a symmetric Dirichlet distribution, where the concentration parameter $\alpha \in [0.1, 1]$ can be adjusted to simulate different levels of non-iidness (lower α results in a stronger non-iid effect). The IDS architectures used for experiments are based on the Deep Neural Network (DNN) implementation used in [10]. We ran experiments on both Binary and multiclass classification tasks, thus two possible output dimensions for the IDS are listed.

C. Results

We present the results of binary and multi-class classification tasks. We investigate the impact of non-iidness (using different values of α) and data generation amounts (from 0 to 100,000 generated samples) on model performance.

- **Binary classification results**

Table II highlights the substantial accuracy gains achieved through synthetic data generation, as illustrated in Figure 2. For instance, in the binary task with $\alpha = 0.1$, the inclusion of 1,000 synthetic samples increased accuracy from 0.8612 to 0.9453, representing a significant leap in performance. This underscores the ability of GAI to address data heterogeneity and scarcity by creating a more balanced representation of attack scenarios across distributed nodes. However, as the amount of generated data increased to 10,000 and 100,000 samples, the accuracy slightly decreased to 0.9418 and 0.9216, respectively. This suggests that excessive synthetic data can introduce noise or redundancy, highlighting the need for careful calibration of data generation.

For higher α values, where local data distributions are closer to iid, the benefits of synthetic data diminish. For example, with $\alpha = 0.9$, the accuracy improvement is minimal, increasing from 0.9796 to 0.9806 with 1,000 generated samples, as can be seen in Table II. This indicates that in scenarios where the local datasets are already representative of the global distribution, additional synthetic data has limited impact.

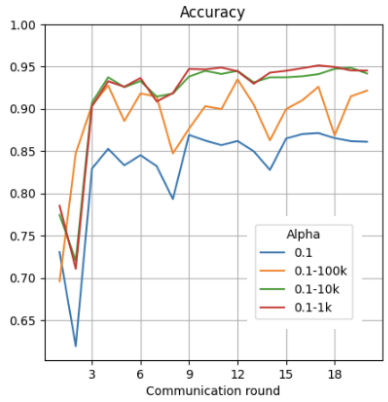

Fig. 2. Accuracy while varying generated samples

TABLE II
BINARY CLASSIFICATION RESULTS

Non iidness (α)	Generation amount	Accuracy
0.1	0	0.8612
	1000	0.9453
	10000	0.9418
	100000	0.9216
0.2	0	0.9615
	1000	0.9581
0.3	0	0.9641
	1000	0.9589
0.9	0	0.9796
	1000	0.9806
	10000	0.9790

- **Multi-classification results**

In the multi-class task (Table III), the benefits of synthetic data are similarly shown in highly non-iid settings. For $\alpha = 0.1$, generating 100,000 synthetic samples boosted accuracy from 0.8903 to 0.9270, highlighting the importance of addressing data imbalance and heterogeneity. The trend remains consistent across other non-iid scenarios, with notable improvements for $\alpha = 0.2$ and $\alpha = 0.3$. For example, with $\alpha = 0.2$, accuracy increased from 0.9479 to 0.9561 when 100,000 synthetic samples were generated. This demonstrates how GAI enhances FL by mitigating the adverse effects of skewed data distributions.

However, in scenarios closer to iid (e.g., $\alpha = 0.9$), the improvement is less significant. Accuracy only increased from 0.9800 to 0.9809 with 1,000 generated samples, indicating that when data is already well-balanced and representative, the need for synthetic augmentation is minimal. These results suggest that the value of GAI for FL is most pronounced in highly heterogeneous environments.

TABLE III
MULTI-CLASSIFICATION RESULTS

Non iidness (α)	Generation amount	Accuracy
0.1	0	0.8903
	1000	0.9001
	10000	0.9023
	100000	0.9270
0.2	0	0.9479
	1000	0.9534
	10000	0.9484
	100000	0.9561
0.3	0	0.9547
	1000	0.9556
	10000	0.9561
	100000	0.9636
0.5	0	0.9744
	1000	0.9774
	10000	0.9783
	100000	0.9776
0.9	0	0.9800
	1000	0.9809

Therefore, based on these results, we can conclude that the integration of GAI with FL provides a scalable solution for enhancing IDS performance in distributed and heterogeneous environments. By addressing the core challenges of non-iid data distributions and data scarcity, our proposition enables FL to achieve higher accuracy and adaptability in real-world cybersecurity applications. The findings also highlight the need for future research to optimize synthetic data generation techniques to balance performance gains with computational efficiency.

V. CHALLENGES AND FUTURE DIRECTIONS

A. Security and Privacy

Although FL provides privacy by keeping raw data localized on devices, it is still vulnerable to attacks such as model inversion, gradient leakage, and data poisoning [11]. These vulnerabilities can compromise privacy and the integrity of the global model, especially in systems integrating GAI, where synthetic datasets may inadvertently introduce new attack surfaces. To mitigate these risks, advanced privacy-preserving techniques such as differential privacy and homomorphic encryption should be integrated into the GAI-FL pipeline. Differential privacy can obscure sensitive information in shared updates, while homomorphic encryption ensures secure computations on encrypted data. Meta-model approaches that operate on metadata instead of raw data can also enhance privacy by reducing exposure to sensitive information [12]. Additionally, reputation-based mechanisms that reward honest behavior among nodes and detect malicious participants can strengthen security and trust in distributed systems.

B. Data Heterogeneity (non-IID)

Data heterogeneity, or the non-IID nature of local datasets across devices, is a critical challenge in FL that affects the generalizability and convergence of the global model [13]. Variations in local data distributions, attack patterns, and sample sizes can lead to biased models that fail to detect

rare or novel threats. GAI presents a promising solution by generating synthetic data to augment local datasets and mitigate the effects of non-IID distributions. However, ensuring the quality, diversity, and representativeness of synthetic data is vital to avoid introducing noise or biases. Future research should explore adaptive generative techniques, such as dynamic GANs or DM, to create synthetic data tailored to local environments. Moreover, federated data augmentation techniques could enable nodes to collaboratively enhance data diversity without compromising privacy. These efforts can ensure that FL systems perform effectively in heterogeneous and dynamic environments.

C. Energy Consumption

The integration of GAI and FL introduces significant computational and energy demands, particularly for IoT devices and edge nodes with limited power resources. Training GAI models such as GANs or DM, along with iterative FL updates, can impact battery-operated devices rapidly. Energy-efficient strategies, including model compression, quantization, and pruning, can reduce computational overhead while maintaining model performance. Collaborative edge-cloud architectures that offload heavy computational tasks to cloud servers can also alleviate the strain on low-power devices [14]. Additionally, adaptive scheduling algorithms that prioritize devices with higher energy reserves can improve the sustainability of the system. Future research should focus on lightweight GAI models and energy-aware frameworks to balance performance with power efficiency. Incorporating renewable energy sources or energy-harvesting technologies into IoT devices could further enhance the sustainability of GAI-FL systems in real-world deployments.

D. Deployment Considerations

The deployment of GAI models within FL frameworks involves crucial design decisions that affect performance, privacy, and efficiency. Centralized deployment of GAI models on global servers allows for high-performance training and synthetic data generation but introduces latency, communication overhead, and privacy risks [15]. In contrast, decentralized deployment of lightweight GAI models on edge devices reduces communication costs and enhances privacy by ensuring that sensitive data remains localized, though it demands optimized models for resource-constrained environments. A hybrid approach, where GAI models are centrally trained and distributed for localized fine-tuning, offers a balanced trade-off between privacy, scalability, and efficiency. Future research should investigate dynamic deployment strategies that adapt to network conditions and device capabilities. Collaborative edge-cloud architectures and energy-efficient frameworks tailored to the heterogeneous requirements of IoT and 6G networks can further support scalable and privacy-preserving deployments.

VI. Conclusion

This study has presented a comprehensive overview of integrating Generative Artificial Intelligence (GAI) and Federated Learning (FL) for Intrusion Detection Systems (IDS) in IoT and 6G networks. The integration of GAI and FL offers several advantages, such as enhancing data diversity, improving adaptability to novel threats, and preserving privacy in distributed environments. By combining the complementary strengths of GAI and FL, this framework addresses key challenges, including non-IID data distributions, limited data diversity, and the dynamic nature of cyber threats. Through conceptual discussions, use cases, and performance comparisons, the potential of this integration has been demonstrated. To further advance this framework and overcome current limitations, we outline several critical future research directions below.

VII. Acknowledgement

This work was supported by the ANR CHIST-ERA project Di4SPDS-Distributed Intelligence for Enhancing Security and Privacy of Decentralised and Distributed Systems.

References

[1] S. Sai, M. Kanadia, and V. Chamola, "Empowering iot with generative ai: Applications, case studies, and limitations," *IEEE Internet of Things Magazine*, vol. 7, no. 3, pp. 38–43, 2024.

[2] X. Wang, Z. Wan, A. Hekmati, M. Zong, S. Alam, M. Zhang, and B. Krishnamachari, "Iot in the era of generative ai: Vision and challenges," *arXiv preprint arXiv:2401.01923*, 2024.

[3] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.

[4] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.

[5] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings *et al.*, "Advances and open problems in federated learning," *Foundations and trends® in machine learning*, vol. 14, no. 1–2, pp. 1–210, 2021.

[6] O. Aouedi, K. Piamrat, G. Muller, and K. Singh, "Federated semisupervised learning for attack detection in industrial internet of things," *IEEE Transactions on Industrial Informatics*, vol. 19, no. 1, pp. 286–295, 2022.

[7] W. Saad, M. Bennis, and M. Chen, "A vision of 6g wireless systems: Applications, trends, technologies, and open research problems," *IEEE network*, vol. 34, no. 3, pp. 134–142, 2019.

[8] S. Agrawal, S. Sarkar, O. Aouedi, G. Yenduri, K. Piamrat, M. Alazab, S. Bhattacharya, P. K. R. Maddikunta, and T. R. Gadekallu, "Federated learning for intrusion detection system: Concepts, challenges and future directions," *Computer Communications*, vol. 195, pp. 346–361, 2022.

[9] I. Sharafaldin, A. H. Lashkari, A. A. Ghorbani *et al.*, "Toward generating a new intrusion detection dataset and intrusion traffic characterization." *ICISSp*, vol. 1, pp. 108–116, 2018.

[10] D. Javeed, M. S. Saeed, M. Adil, P. Kumar, and A. Jolfaei, "A federated learning-based zero trust intrusion detection system for internet of things," *Ad Hoc Networks*, vol. 162, p. 103540, 2024.

[11] V. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantanha, and G. Srivastava, "A survey on security and privacy of federated learning," *Future Generation Computer Systems*, vol. 115, pp. 619–640, 2021.

[12] O. Aouedi and K. Piamrat, "F-bids: Federated-blending based intrusion detection system," *Pervasive and Mobile Computing*, vol. 89, p. 101750, 2023.

[13] ——, "Surfs: Sustainable intrusion detection with hierarchical federated spiking neural networks," in *ICC 2024-IEEE International Conference on Communications*. IEEE, 2024, pp. 2173–2178.

[14] O. Aouedi, "Towards a scalable and energy-efficient framework for industrial cloud-edge-iot continuum," *IEEE Internet of Things Magazine*, 2024.

[15] J. Konečnỳ, H. B. McMahan, D. Ramage, and P. Richtárik, "Federated optimization: Distributed machine learning for on-device intelligence," *arXiv preprint arXiv:1610.02527*, 2016.