2025 International Conference on Emerging Smart Computing and Informatics (ESCI)
AISSMS Institute of Information Technology, Pune, India. Mar 5-7, 2025

Intrusion Detection using Federated Learning with Application of Federated Dropout

Swati Shinde
Department of Computer Engineering
Pimpri Chinchwad College of Engineering
Pune, India
swati.shinde@pccoepune.org

Prof. Bal Virdee
Center for Communication Technology
London Metropolitan University
London, United Kingdom
b.virdee@londonmet.ac.uk

Ashish Khanna
Department of Computer Science
Maharaja Agrasen Institute of Technology
Delhi, India
ashishkhanna@mait.ac.in

Saurabh Dhakite
Department of Computer Engineering
Pimpri Chinchwad College of Engineering
Pune, India
dhakitesaurabh@gmail.com

Tushar Badlani
Department of Computer Engineering
Pimpri Chinchwad College of Engineering
Pune, India
tusharbadlani0@gmail.com

2025 International Conference on Emerging Smart Computing and Informatics (ESCI) | 979-8-3315-1568-3/25/$31.00 ©2025 IEEE | DOI: 10.1109/ESCI63694.2025.10988162

***Abstract*— The massive proliferation of IoT devices in recent years creates new challenges in securing distributed systems, especially when it comes to intrusion detection and mitigation in real-time. Most traditional cloud-based intrusion detection approaches encounter issues ranging from increased latency to bandwidths and even centralized vulnerabilities. This paper introduces an innovative IDS framework based on federated learning (FL) combined with a novel federated dropout mechanism and edge computing. Federated dropout updates model parameters selectively during communication rounds, minimizing bandwidth usage considerably without compromising model performance. This work resolves key challenges such as communication overhead and encrypted traffic processing, opening the door to strong, real-time, and decentralized security solutions.**

Keywords— Edge Computing, Intrusion Detection, IoT, Blockchain, Real-time Security

I. Introduction

The high deployment of the IoT devices into the gigantic IoT network causes increased connection among the heterogeneous elements, which enables intelligent automation, surveillance, and regulation in the sectors. However, such a heavy array of devices concurrently introduces an increasing risk towards security because they predominantly become vulnerable to cyber threats such as unauthorized access, data compromise, and DoS attacks. Conventional IDS depends largely on centralized cloud computing with restrictions on latency, bandwidth utilization, and protection of privacy in data.

As illustrated in Fig. 1, edge computing, which processes data at the network's periphery, near the data source, has emerged as a viable alternative. It provides real-time data processing, reduced latency, and increased scalability for IoT networks. However, with the diversity and scale of IoT networks, developing an efficient and accurate IDS remains a challenge [1].

Federated learning therefore gives way to a new approach for intrusion detection as the model can be trained in a decentralized manner across multiple devices or edge nodes such that raw data does not need to be sent to a central server for training [2]. This involves every edge node building an IDS model based on its locally collected data and sending only updates to the central server for aggregation. This decentralized learning process increases the level of data privacy and bandwidth efficiency and is more apt for IoT environments.

Combining federated learning with edge computing, the intrusion detection models can draw on distributed data processing without latency and data privacy violation. This approach would eliminate some centralized IDS-related vulnerabilities in terms of bottlenecks and single points of failure, in addition to enhancing network attack resilience in a system [3].

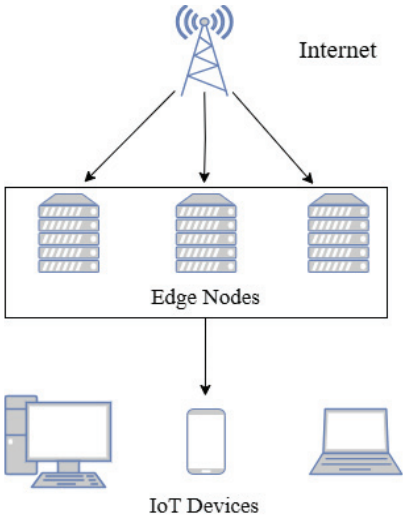

Fig. 1. Edge Computing Network

II. Related Work

Federated learning and edge computing convergence has turned into a promising approach to enhance intrusion detection in IoT networks. The current literature defines various techniques and approaches that address different

aspects of intrusion detection in distributed systems, considering privacy, efficiency, and scalability issues.

The Mun and Lee [3] proposed an internet traffic classification framework that relies on federated learning in balancing both data privacy and computational efficiency. In this regard, their work showed the possibility of federated learning in approaching the performance of centralized models with regards to precision metrics compared to a decentralized training process on client devices instead of submitting their data to centralized servers. This methodology aligns well with the requirement of privacy-preserving techniques in intrusion detection systems, particularly in IoT networks, as data is being processed locally on edge nodes, thus minimizing risks related to centralized aggregation of data. Results highlight the local processing significantly for improvement in privacy and do not study communication overhead in large federated learning deployments. Liu *et al.* [2] presented the collaborative intrusion detection framework, which relies on blockchain and federated learning in vehicular edge computing environments.

They tend to integrate blockchain to secure the process of model updates with such a model; that way, the integrity and reliability of data can be assured through non-tamper ability of model updates. FL combined with blockchain achieves data privacy and security issues, which can significantly be posed in vehicular networks requiring strong intrusion detection mechanisms. However, such increased complexity in blockchain requires higher communication and computational overheads, which would be demanding for implementation in highly dynamic IoT settings where rapid communication and low latency are critical. Nie *et al.* [1] proposed an IDS model for social IoT using collaborative edge computing and GANs. Using GANs to train the IDS on a number of edge nodes, it advances the detection of threats while emphasizing secure and decentralized learning, enhancing the model's capacity to detect new and evolving intrusion patterns.

GANs can model complex intrusion behaviours because of their generative capability. It can improve the detection accuracy across very diverse IoT environments. However, the high computational requirement of GANs may limit its effectiveness in resource-constrained IoT devices and thus strict in the deployments of this on edge-based IDS real-world application scenarios. Singh, Chatterjee, and Satapathy propose a hybrid IDS framework developed with the help of edge computing and techniques specific to machine learning for mobile edge computing (MEC) networks [4]. Their approach takes advantage of edge-based detection to minimize latency, thereby allowing for real-time intrusion detection in mobile environments. The hybrid setup combines various machine learning models suited to the fluctuating nature of MEC, where resources are mostly limited, thus providing an all-season solution for intrusion detection.

While they are able to reduce latency, dependency on time-interval-based updates of the models might lead to communication overheads, which continues to be challenging for scalability in IoT settings that are distributed. Gyamfi and Jurcut published a comprehensive review of intrusion detection methods targeted for IoT network environments, focusing especially on MEC, machine learning methodologies, and relevant datasets [5]. This review points out various benefits of edge computing for IDS, like that data locality offered by MEC enables low latency response crucial to implement real-time intrusion detection systems in IoT systems. In this regard, the authors emphasize the point that although MEC may reduce some privacy and scalability concerns, improving communication efficiency in federated learning remains a significant challenge in resource-constrained IoT networks. Spadaccino and Cuomo [6] discussed the potential of combining machine learning with edge computing for IoT IDS, identifying both opportunities and challenges.

In addition, borrowing from Hemanth et al. [7], our model includes intelligent edge computing to harmonize the interaction of edge and cloud resources. Li et al. [8] proved the scalability and applicability of federated learning in various fields, validating its feasibility in IoT IDS applications. Mohy-Eddine et al. [9] also underlined the significance of ensemble learning in augmenting IDS performance, a concept included in our federated dropout approach to enhance model robustness and accuracy.

Rahman et al. [10] and Rashid et al. [11] pointed out the importance of using privacy-preserving methods, particularly in encrypted environments, where both centralized and decentralized strategies are confronted with inherent limitations. Based on this observation, our federated dropout framework meets scalability and privacy challenges by focusing parameter updates on essential components while lowering communication expenses without compromising accuracy in detection.

Lastly, Fedorchenko et al. [12] recognized the requirement of federated learning-based IDS to be adaptive towards heterogeneous IoT devices along with efficient management of communication overhead. Our design is compliant with these requirements, providing a scalable, privacy-enhancing, and resource-conserving IDS solution suitable for real-world IoT scenarios.

Their work identifies some of the pressing issues, including model generalization over heterogeneous IoT devices, resource constraints, and the requirement for real-time processing. They argue that FL would aid in handling these issues by providing adaptive and distributed learning models. However, the study also reveals that high communication overhead could be a limiting factor when interacting with clients and the server to update the models repeatedly, thereby restricting scalability in massive IoT networks. Our analysis of these works reveals that in federated learning, privacy preservation is a strong necessity while dealing with the encrypted traffic and the overhead of communication incurred by frequent model updates in federated learning. These are mainly confronted challenges in IoT networks, which are most commonly devices in resource constraint networks, where significant proportions of traffic will be encrypted to protect user privacy.

To address these shortcomings, we propose federated dropout: a way to update selectively only a fraction of model parameters for each communication round. This mechanism reduces the volume of data to be exchanged between clients and the central server, thereby reducing communications costs while preserving privacy by directly operating on the encrypted traffic data without decryption requirements. Federated dropout thus concentrates only on critical parameter modification for a given iteration, thereby efficiently alleviating communication constraints while maintaining model accuracy and robustness. Further, our approach uses edge computing fundamentals that enable intrusion detection

close to data origins to minimize latency in the system's responsiveness in real time to emerging threats. Our proposed framework combines federated dropout with encrypted traffic handling, and it aims for a scalable, privacy-preserving, and resource-efficient IDS for IoT networks, under real-world IoT environments' constraints.

This approach not only increases confidentiality as well as effectiveness in the intrusion detection methodology but also enhances its adaptability to the unique challenges of different and ever-changing IoT environments.

III. Methodology

The methodology encompasses data preprocessing, design of model architecture, training strategies, and metrics employed for performance evaluation. There were six different configurations; these configurations were run, including ANN and CNN, within centralized, federated, and federated dropout frameworks. Each configuration was tested against a specified set of performance criteria to ascertain its effectiveness as an intrusion detection mechanism specifically in the context of resource-constrained Internet of Things networks..

In order to keep uniformity among all models, the primary source of data was maintained to be KDD Cup 1999, coupled with extensive preprocessing and standardized metrics of evaluation. Dataset Creation.

A. Dataset

It uses the KDD Cup 1999 dataset, which is very well known as a benchmark for network intrusion detection system evaluation. Data preparation involves loading the dataset and converting it to what the model requires. Some of the most important features summarized in the dataset include duration, protocol type, service, and flags that represent the most important aspects of network traffic, which were encoded to transform categorical variables into their corresponding numerical values. This transformation was required to make it machine learning framework-compatible and to ensure proper usage of categorical data during training of the model. After doing encoding, this dataset was split into a training set and test set at the ratio 80:20. This would enable giving much more extensive training to the model and its testing against the test data. Each data sample was reshaped into two-dimensional format to include in the CNN-based architecture allowing for spatial features in the dataset, and in addition, the training data were preprocessed as batches to permit effective memory usage and computational speed during model training.

B. Architectural Models

To address the different intrusion needs in IoT networks, two major architectures are used: ANN and CNN. This model consists of dense layers and implements fully connected layers with activation functions of ReLU. This architecture is optimized using a binary cross-entropy loss function, where it allows the network to create the minimum errors in the classification of the intrusions. This architecture uses two convolutional layers combined with max-pooling operations, which were implanted to capture spatial features intrinsic to the intrusion detection. It further processes the output of convolutional layers through fully connected layers culminating in a sigmoid-activated output layer that allows for binary classification as shown in Fig. 2.

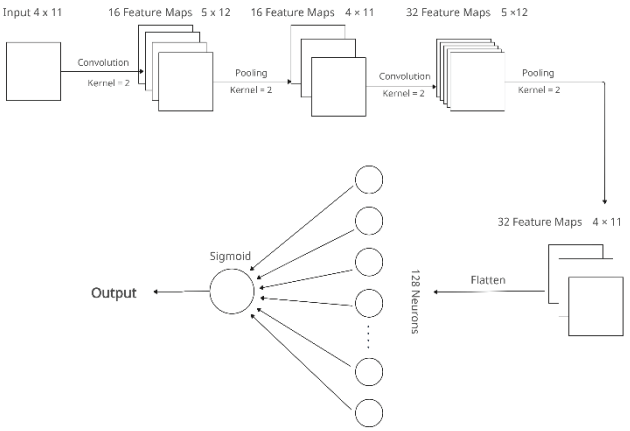

Fig. 2. CNN Architecture

In the case of ANN as well as CNN architectures, three of the training paradigms are used: centralized training, federated training, and federated training with dropout. Centralized training is a way of traditional training, wherein the entire dataset is trained on a single model, whereas federated learning distributes the training process across multiple clients, providing decentralized processing of data. Federated dropout yet brings a new regularization mechanism: a portion of the neurons in the model are randomly disabled in every round of communication. This targeted dropout would reduce the costs of the associated communication and would induce redundancy in neuron connection to overfitting.

C. Training Phase

Each of the three types of model configurations-centralized, federated, and federated with dropout-was learned procedurally. In centralized learning, the dataset's overall data was used directly to update the model's weights through iterations of epochs. The models in this architecture utilized iterative back-propagation to iteratively fine-tune parameters and converge towards optimal intrusion detection solutions.

By distributing the data across multiple clients, federated learning is typical. In this approach, each client independently trains a local model with its allocated subset of data. Once set numbers of epochs of training are accomplished, the weights of the model of each client are averaged to produce an update of the global model. This deconstruction methodology ensures preserving data privacy because it keeps the data on the client side rather than on a central server. In the third training configuration, federated dropout had been applied wherein all the neurons in each of the client models had been selectively turned off. This would reduce the communication overhead typical to federated learning while decreasing the parameters transmitted between the clients and the central server. Through this dropout mechanism, only a fraction of parameters in the model was updated during each communication round, thus allowing for less data transmission yet an improvement in terms of scalability in bandwidth-limited environments.

IV. Evaluation and Results

The assessment of each model configuration was carried out utilizing four key metrics—accuracy, precision, recall, and F1-score—on a designated test set that was not used during training. Accuracy functioned as a broad performance indicator, representing the percentage of instances that were

classified correctly. Precision evaluated the proportion of true positives relative to all instances predicted as positives, thereby gauging the model's capability to accurately identify intrusions while minimizing false alarms. Recall quantified the ratio of actual intrusions identified by the model, meaning its potential in distinguishing the anomalous networking behavior. The F1-score, which was computed as the harmonic mean of precision and recall, reflected the overall measure that considered the false positive and false negatives rates-very much conducive for measuring models in security-related application areas like intrusion detection.

A. Precision and Recall:

Precision assesses the proportion of true positives among predicted positive detections, while recall calculates the proportion of true positives identified among all actual positives. Both metrics are critical for understanding the balance between detection and false alarm rates.

$$Precision = \frac{TP}{TP+FP} \quad (1)$$

$$Recall = \frac{TP}{TP+FN} \quad (2)$$

B. Accuracy:

This metric measures the ratio of correctly classified instances to the total number of instances, providing a general indication of model performance.

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN} \quad (3)$$

C. F1-Score:

The F1-score is the harmonic mean of precision and recall, representing a balanced metric that considers both false positives and false negatives.

$$F1\ Score = 2 * \frac{Precision * Recall}{Precsion + Recall} \quad (4)$$

D. Results

The performance evaluation metrics used in the present study-accuracy, precision, recall, and F1 score-all point to better effectiveness of the ANN and CNN frameworks in intrusion detection contexts, centralized and federated alike. For the centralized ANN model, the accuracy rate is impressive at 99.91%, while precision, recall, and F1 score experience significant increases. Although accuracy falls slightly to 98.37% when trained in the federated learning environment, precision, recall, and F1 score remain at relatively high values of 98.39, 99.59, and 98.99, respectively as per TABLE I. . This slight fall in accuracy and precision is reflective of some fundamental inherent flaw in federated learning; aggregation of model parameters by each distributed client in its local environment without having access to the actual data reveals a possibility of slightly affecting model performance. However, both recall and F1 score remain high, which means a significant percentage of positive cases is still discovered by the model, its main desired property in intrusion detection-related tasks.

TABLE I. MODEL PERFORMANCE

Algorithm	*Accuracy*	*Precision*	*Recall*	*F1 Score*
ANN	99.91	99.99	99.89	99.94
ANN with Federated Learning	98.37	98.39	99.59	98.99
ANN with Federated Dropout	88.54	91.48	94.50	92.96
CNN	99.87	99.95	99.88	99.92
CNN with Federated Learning	99.57	99.96	99.51	99.73
CNN with Federated Dropout	79.60	80.61	98.15	88.52

When federated dropout is implemented, which involves randomly dropping particular neurons while training to lower the information shared between the clients, the ANN model suffers a significant drop in performance to 88.54% accuracy. Precision and recall also decline, whereby the model achieves an accuracy of 91.48%, recall of 94.50%, and an F1 score of 92.96%, respectively. This decrease is expected because federated dropout adds an additional sparsity layer to the model that leads to lower accuracy associated with information loss. However, the technique efficiently reduces communication overhead. Federated dropout decreases the communication burden as it is sending fewer parameters or updates during model aggregation, which is a major concern in IoT networks with bandwidth and energy constraints. In so doing, it therefore leads to this trade-off: slightly reduced detection efficacy but definitely significant improvements in communication efficiency and a trade-off that may be justified in bandwidth-limited contexts where bandwidth conservation is paramount.

The CNN model follows the same scenario. On a centralized configuration, the CNN achieves accuracy of 99.87%, with high values of precision of 99.95%, recall of 99.88%, and F1 score of 99.92%. In federated learning, the accuracy might slightly decrease to 99.57% though the performance metrics remain robust with an F1 score of 99.73%. This consistency in performance shows that CNN, due to its feature extraction capability, federated learning is achieved with less degradation in performance when compared with the ANN model. However, federated dropout in CNN results in an immense decline in accuracy by 79.60 % and precision by 80.61%, though recall goes up to 98.15%, which makes an F1 score of 88.52%. The findings suggest that the federated dropout, although it affects the precision and overall accuracy of the CNN, remains capable of maintaining a considerable percentage of true positive instances—a requirement very important for the reduction of false negatives in security-related applications.

In summary, although federated dropout introduces performance compromises, it plays a vital role toward reducing the communication overhead. Federated dropout reduces the frequency and size of updates needed between the server and clients for showing significant bandwidth savings. This efficiency is especially useful in IoT and Edge Computing environments where devices are constrained in connectivity and power resources. Although this means a slight drop in model accuracy, communication costs are reduced and scalability increases, meaning that more devices can participate in the federated learning framework without overloading network capacities.

V. Conclusion

This paper explores the integration of federated learning with edge computing as a promising approach towards enhancing intrusion detection systems for IoT networks with important balances to privacy, efficiency, and overheads in communication. Centralised traditional IDS approaches are characterized by latency due to bandwidth limitation and privacy concerns since they demand centralization of data that is unfavourable for expansive as well as dynamic IoT environments. By enabling federated learning, the framework allows for very localized model training on dispersed edge devices and safeguards data privacy through localization of raw data, reducing thereby the need to transfer the data to a central server.

Our findings indicate that Federated Dropout integrated with edge computing leads to higher adaptability and scalability of IDS in the IoT context, especially for traffics encrypted. The framework satisfies the specific requirements of heterogeneous IoT deployments, creating intrusion detection from end towards the services with reduced latency and cost for communication in real-time with privacy preservation which comes at the cost of slightly degraded accuracy. Future research endeavours may aim at refining the Federated Dropout methodology to reduce accuracy degradation further, while simultaneously investigating sophisticated privacy-preserving strategies and enhancing model generalization to bolster robustness across diverse IoT devices. This investigation highlights the promise of federated and edge-based methodologies to markedly enhance secure, scalable, and privacy-conscious management of IoT networks.

References

[1] Nie, L., Wu, Y., Wang, X., Guo, L., Wang, G., Gao, X., & Li, S. (2021). Intrusion detection for secure social internet of things based on collaborative edge computing: a generative adversarial Network-Based approach. *IEEE Transactions on Computational Social Systems*, *9*(1), 134–145.

[2] Liu, H., Zhang, S., Zhang, P., Zhou, X., Shao, X., Pu, G., & Zhang, Y. (2021). Blockchain and federated learning for collaborative intrusion detection in vehicular edge computing. *IEEE Transactions on Vehicular Technology*, *70*(6), 6073–6084.

[3] Mun, H., & Lee, Y. (2020). Internet traffic classification with federated learning. Electronics, 10(1), 27.

[4] Singh, A., Chatterjee, K., & Satapathy, S. C. (2021). An edge based hybrid intrusion detection framework for mobile edge computing. *Complex & Intelligent Systems*, *8*(5), 3719–3746.

[5] Gyamfi, E., & Jurcut, A. (2022). Intrusion detection in internet of things systems: a review on design approaches leveraging multi-access edge computing, machine learning, and datasets. Sensors, 22(10), 3744.

[6] Spadaccino, P., & Cuomo, F. (2020). Intrusion Detection Systems for IoT: opportunities and challenges offered by Edge Computing and Machine Learning. arXiv preprint arXiv:2012.01174.

[7] Hemanth, D. J., Gupta, B. B., Elhoseny, M., & Shinde, S. V. (Eds.). (2023). Intelligent Edge Computing for Cyber Physical Applications.

[8] Li, L., Fan, Y., Tse, M., & Lin, K. Y. (2020). A review of applications in federated learning. *Computers & Industrial Engineering*, *149*, 106854.

[9] Mohy-Eddine, M., Guezzaz, A., Benkirane, S., & Azrour, M. (2023). An effective intrusion detection approach based on ensemble learning for IIoT edge computing. *Journal of Computer Virology and Hacking Techniques*, *19*(4), 469-481.

[10] Rahman, S. A., Tout, H., Talhi, C., & Mourad, A. (2020). Internet of things intrusion detection: Centralized, on-device, or federated learning?. *IEEE Network*, *34*(6), 310-317.

[11] Rashid, M. M., Khan, S. U., Eusufzai, F., Redwan, M. A., Sabuj, S. R., & Elsharief, M. (2023). A federated learning-based approach for improving intrusion detection in industrial internet of things networks. *Network*, *3*(1), 158-179.

[12] Fedorchenko, E., Novikova, E., & Shulepov, A. (2022). Comparative review of the intrusion detection systems based on federated learning: Advantages and open challenges. *Algorithms*, *15*(7), 247.