



OPEN ACCESS

EDITED BY

Shu Hu,
Purdue University, United States

REVIEWED BY

Li Lin,
Purdue University, United States
Mohammed Altaf Ahmed,
Prince Sattam bin Abdulaziz University,
Saudi Arabia
Ming Yang,
Kennesaw State University College,
United States

*CORRESPONDENCE

Fernando Moreira
✉ fmoreira@uportu.pt

RECEIVED 17 June 2025

ACCEPTED 21 July 2025

PUBLISHED 25 August 2025

CORRECTED 27 August 2025

CITATION

Godavarthi D, Rekapalli VCS, Mohanty S,
Jaswanth JVSDV, Polisetty D, Dash BB and
Moreira F (2025) Federated quantum-inspired
anomaly detection using collaborative neural
clients.

Front. Artif. Intell. 8:1648609.
doi: 10.3389/frai.2025.1648609

COPYRIGHT

© 2025 Godavarthi, Rekapalli, Mohanty,
Jaswanth, Polisetty, Dash and Moreira. This is
an open-access article distributed under the
terms of the [Creative Commons Attribution
License \(CC BY\)](#). The use, distribution or
reproduction in other forums is permitted,
provided the original author(s) and the
copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic
practice. No use, distribution or reproduction
is permitted which does not comply with
these terms.

Federated quantum-inspired anomaly detection using collaborative neural clients

Deepthi Godavarthi¹, Venkata Charan Sathvik Rekapalli¹,
Sribidhya Mohanty², J. V. S. D. Vigneswara Jaswanth¹,
Dinesh Polisetty¹, Bibhuti Bhusan Dash³ and Fernando Moreira^{4*}

¹School of Computer Science and Engineering (SCOPE), VIT-AP University, Amaravati, India,

²Department of Electronics and Communication Engineering, (Graphic era Deemed to be University), Dehradun, India, ³School of Computer Applications, KIIT Deemed to be University, Bhubaneswar, India, ⁴REMIT, Universidade Portucalense, Porto, Portugal

Introduction: The fusion of deep-learning-based and federated methods has brought great progress in anomaly detection. Yet the systems of today still suffer from certain glaring issues. First, aggregation of data on a central entity poses dangerous privacy hazards. Second, such models could not scale and adapt to heterogeneous and distributed environments. Lastly, fine consideration has hardly been given to quantum-inspired computational paradigms that may promise to improve both speed and security of such systems. To fill in these gaps, this research proposes a completely novel quantum-inspired federated learning approach to anomaly detection that keeps data private and allows for further implementations of quantum computing applications.

Methods: The proposed system works on a client-server architecture comprising multiple clients, which either run training of local feedforward neural networks on different private subsets of their data or choose to not participate during an iteration. Clients never pass raw data to the server but instead alternate by sending the server the parameters of the trained model. The server aggregates these local updates by the FedAvg algorithm and produces the global model. The present implementation focuses mainly on utilizing classical deep learning; however, the architecture is made flexible enough to intertwine smoothly with quantum machine-learning paradigms in the future, thus enabling quantum technological enhancement down the road without requiring the entire system to be rebuilt.

Results: The framework could produce up to 79% of anomalous detection accuracy. The system had effective learning across distributed clients whilst ensuring that no piece of private data was being shared or spilled (exposed) between clients. These results ensured that the framework maintained its performance while keeping its privacy intact, a very crucial consideration on which to ever really deploy such in sensitive areas.

Discussion: The approach allows privacy-preserving anomaly detection across multiple domains and serves as a framework for enlarging and scaling the system. Being quantum-inspired compatible allows for future-proofing and further expediting and enhancing security. The system, having the capability to securely work in a distributed manner, can, thus, be utilized in critical information domains like cybersecurity, finance, and healthcare, where privacy of data is deemed extremely important. This work, thereby, offers a useful federated learning approach towards anomaly detection while going a step further towards the incorporation of quantum computing into secure, distributed AI systems.

KEYWORDS

federated learning, anomaly detection, TCP based model communication, privacy-preserving AI, distributed systems, quantum-inspired neural networks

1 Introduction

From 2010 to 2020, an unprecedented growth of IoT deployments changed the way many sectors operate, including healthcare, smart cities, and industrial automation. However, with the masses of interconnected sensors and systems being deployed worldwide, critical vulnerabilities have been created, thus posing threats to anomaly detection in IoT networks. Conventional anomaly detection mechanisms force data to be collected first into a central repository, presenting a number of drawbacks such as bandwidth constraints, latency issues, and data privacy concerns. Recently, with enforceable data security regulations and user privacy becoming an important subject, there are more and more urges for decentralized solutions that efficiently and securely detect anomalies. We describe in this paper an architecture for anomaly detection and the federated learning method for distributed IoT systems. The highlighted features of this approach are data locality, computational efficiency, and model generalization. Federated learning (FL) is considered an emerging paradigm, allowing the training of machine learning models jointly on several decentralized devices while ensuring that data remain at the local sites. Rather than pushing data from raw bases, model updates alone are communicated to a central server performing global aggregation. Such a privacy approach would be highly valued for applications where data ownership, confidentiality, and compliance are at stake: healthcare or critical infrastructure monitoring. In this study, we exploit FL to build a reliable and distributed anomaly detection system. Clients are given subsets of data and train their respective local models independently before sharing model updates with a central server. Next, the global model is updated via federated averaging to form a system where learning takes place in a collaborative fashion but retains local data integrity. The methodology indeed respects privacy concerns yet rolls out efficiently with the rising number of edge gadgets in an IoT network. Nevertheless, this framework poses several different challenges, especially for anomaly detection. As the data is non-IID through the clients, this constitutes model divergence or at least poor performance. Also considered are asynchrony between clients that include computational disparities and communication delays, which slow down training. For anomaly detection, there exists a natural imbalance in learning where anomaly events tend to occur rarely, so it becomes complicated to properly set up model training and evaluation. The chances almost appear to be that anomalies are even rarer for any particular client within the federated framework, surely downplaying its detection capabilities.

To settle this issue, the paper proposes a central server that coordinates several clients while keeping communicative overhead low, imposing a consistent model architecture, and adaptively aggregating client updates. The paper further establishes a quantitative notion of balanced client contribution, even under conditions of imbalance in data size and imbalanced anomaly distributions. Our methodology introduces an innovative approach, allowing clients to work asynchronously, communicating through sockets in real-time, and keeping a systematic save of the models. The system has been evaluated in an existing set of anomalies, in which clients were trained on separate parts of the dataset. So far, analyses have been conducted to

look into how each client performed over a few iterations of training, with good improvement records in the accuracy of the model and convergence while ensuring data privacy, preserving model utility, and permitting deployment to edge devices. Another benefit is that this infrastructure requires very minimal reconfiguration to add a new client. To summarize, best model utility, data privacy preservation, and edge deployment are all taken into account by this system in applying anomaly detection in a realistic manner. Our technical model works on one central server with multiple federated Python clients. Each client has a lightweight neural network with TensorFlow and trains it independently on the dataset assigned to it. After local training, a serialized messaging system is used between the clients and the server via TCP sockets. The messages contain model weights, some metadata like the number of local samples, the round identifier, and so forth. The Federated Averaging algorithm returns updates to the clients in the form of an updated global model. The entire communication protocol will repeat for a number of rounds, similar to collaborative learning. Local and global models for every round are then stored in the system for evaluation, potential rollback, and fine-tuning.

This modularized and transparent architecture will allow further extensions to be realized in the future, including adversarial robustness, adaptive learning rates, and differential privacy. From a practical standpoint, the demands of nowadays IoT ecosystems for an intelligent yet adaptive anomaly detection framework are constantly rising. Since they differ in hardware capabilities, data sensitivities, and communication conditions, it becomes imperative for any distributed learning system to take such disparities into account or else face performance degradation. Also, this feature will allow sporadically connected and less powerful computing devices to contribute to model updates while guaranteeing that the overall training progression would not be delayed. The clients perform local training and submit updates whenever they can. This gets rid of the limitations to scalability present in synchronous federated learning, thus enabling better real-time application. The framework accommodates horizontal extensibility so that new clients can be added with ease. Thus, this allows for incremental scaling towards bigger IoT deployments, such as a smart grid or autonomous transportation systems, without forcing the entire pipeline for learning to be reconfigured. Using this approach, deployment-wise real-edge requirements get sufficiently addressed by the system. Latency is created due to local model inference and partial training, which is indeed important in medical emergency detection, industrial fault monitoring, or fraud prevention for financial transactions. The architecture remains the same for all the clients with lightweight neural networks through TensorFlow, thus enabling this to run over any kind of hardware platform-embedded from devices to cloud-based nodes. The communication layer resides on top of TCP socket-based message passing, which brings in reliability and transparency.

While the federated architecture underlying this work—which is a client-server set-up with TCP being used for communication and FedAvg for aggregation—is well grounded in established methods, its novelty is situated in its domain-specific integration and quantum-inspired design orientation. Contrary to previous works that utilize FL for general-purpose classification, this one, instead, designs and tunes the framework for anomaly detection in distributed privacy-sensitive

environments. On top of that, embedding quantum-inspired ideas into FL-design, for instance, entanglement-aware data representation and feature interdependence modeling within dense classical neural layers, gives fresh conceptual taste to the FL paradigm. Other practical extensions brought about by the work include structured hyperparameter tuning, system sensitivity testing, and anomaly-specific evaluation metrics that are all engineered to make FL more valuable in high-stake domains like cybersecurity and IoT anomaly detection. Hence, the FL infrastructure itself is still pretty much standard, but the paper in itself stands on contributions on making it more adaptable, interpretable, and hybrid quantum-oriented.

2 Related works

The development in anomaly detection brings more privacy-oriented federated learning while boosting detection efficiency of rare anomalies across decentralized networks. This method exploits the power of quantum computing, especially in processing a type of complex data that standard methods cannot analyze well. For example, hybrid quantum-classical algorithms can dramatically cut down on communication overhead while preserving accurate classification, as shown by recent results demonstrating fewer communication rounds than were used in traditional federal approaches (Bhatia and Kais, 2023). Additionally, harnessing distributed computing resource leads to collaborative training process, involving critical anomaly detection in highly sensitive domains such as cybersecurity and healthcare where data privacy is an utmost concern (Anwar et al., 2024; Nguyen et al., 2022). Federated learning operates on the principles of collaboration and aggregation of client insights, thereby yielding a collective intelligence model that enhances detection (Santin et al., 2022). In further evolution of federated quantum-inspired anomaly detection, it is crucial to address the angle of integrating edge computing with this framework. Deploying models on edge-based devices reduces the latency further, allowing greater agility in real-time processing-the facility imperative to instances that must bear with immediate anomaly detection like in financial transactions or healthcare-monitoring systems (Liberti et al., 2024). A blend of edge computing with quantum-inspired techniques may open a new door for alternative methods in which aided-locally obtained knowledge is still preserved for privacy, thus jointly strengthening anomaly detection systems in terms of security and efficiency (Imteaj et al., 2022). Quantum machine learning seeks to develop new methods in pattern recognition and optimization through its applications in several research areas. QML is effective in victories in certain classification tasks over classical models since it practically operates in high-dimensional feature space-domain as per the works of Havlek et al. (2019) and Schuld and Killoran (2019).

In particular, it is important to view the consequences of network anomaly detection in several realms, especially industrial control systems, where privacy of data and integrity of operation actually matter. Recent studies deem that the federation models are efficient depending on the geographical displacements of such systems for collaborative learning without threatening the secrecy of sensitive data (Dehlaghi-Ghadim et al., 2023). Mutual enhancement of anomaly detection and reinforcement of cyber-security resilience mechanisms through complex attacks further applies to industries where data sensitivity is of importance; namely banking and healthcare (Alsulaimawi, 2024; Purohit and Manimaran, 2024). Yet these techniques could be challenged by adversarial maneuvers, such as data poisoning and membership

inference, effects that could otherwise corrupt the learning mechanism (Gosselin et al., 2022; Hasan, 2023). Since sensitive domains are areas to be served by federated learning, services requiring technical robustness must also now advocate for ethical response systems. The development of frameworks that optimize anomaly detection while delivering explainability about the model's reasoning will allow stakeholders to grasp and therefore validate the outputs generated by these sophisticated systems. Besides, as the domain shifts toward applications, one may need to establish a standard for interpretability so as to form a standard that helps users develop confidence and comply with regulations, especially as ethical factors governing data privacy and security won out (Burke and Stets, 2023). This makes investigating interpretability with anomaly detection models all the more necessary. In an effort to resist model drifts and deterioration in performance in the face of heterogeneity, studies such as McMahan et al. (2017) propose such algorithms as Federated Averaging (FedAvg). While FL evolves, interest has recently increased in how FL integrates with quantum computing, thus giving rise to an exciting area generally referred to as Quantum Federated Learning (QFL).

Quantum Machine Learning (QML) inherently yields benefits in solving optimization problems and in processing high-dimensional data efficiently, thanks to quantum parallelism. This ability is crucial for anomaly detection, which is computationally intensive to detect rare but severe events. Recent experiments with hybrid quantum-classical models like VQCs and neural networks have shown that such quantum models could outperform classical benchmarks in accuracy and faster convergence, primarily in small-data situations. The FedAvg algorithm was introduced by Bhatia and Kais (2023), who coordinate the updates from these decentral clients to yield a strong global model. In network intrusion detection, FL has shown promise in detecting distributed and evolving attack patterns without compromising data ownership. Implementing in training file leverages such a concept with local training orchestration across clients on anomaly data sets and synchronization of the updates to a central server. This approach enables huge scalability and promote security-aware learning under non-IID, resource-constrained IoT systems. Quantum Reinforcement Learning combines classical decisions-making with the quantum mechanics computation paradigm. Classical RL agents usually have problems with high-dimensional feature spaces and slow convergence, especially in areas such as cybersecurity. QRL overcomes these by using quantum superposition and entanglement to maximize policy expression and expedite learning. As referred to in the study by Santin et al. (2022), quantum-enhanced learning agents may have an edge over classical agents in select environments. The client Python script file felicitates this advancement by weaving quantum circuits through the PennyLane framework. This enables a hybrid quantum-classical policy model for the binary classification of network behaviors. Furthermore, this design is very efficient in detecting minute anomalies and adapting to rapidly evolving threats in the situation of dynamic IoT networks. Since there is an inherent stochastic nature of quantum computations, QRL enjoys heightened exploration capabilities for its intelligent and responsive cybersecurity exploitation. The merger of FL with QRL opens a novel and potent paradigm to train an intelligent decentralized Intrusion Detection System (IDS). This hybrid architecture profits from the privacy-preserving traits of FL while leveraging the computational power of quantum models.

Recent studies explored the use of variational quantum circuits within federated settings so that clients may train lightweight quantum models locally and then participate in a global update cycle. The hybrid architecture promotes asynchronous communication, fault

tolerance, and learning over non-IID data distributions; hence, making it the right fit for performing real-time threat detection over distributed IoT infrastructures where centralized approaches simply cannot perform. The main limitation is the lack of available quantum hardware able to do real-time inference on edge devices. As such, the current QRL models, including the implementations shown in client python file, are usually simulated and not run on real quantum processors. Anomaly detection remains a stuff of great significance in present-day data-driven systems so that early detection of outliers or unusual patterns is important for averting functional failures, security breaches, or disruption costs (Zhao et al., 2021; Mohri et al., 2019). With increased concern of privacy, regulatory compliance, and decentralization of the network, more flexible and privacy-respecting methods have been investigated by researchers. Among them, Federated Learning and Quantum Machine Learning stand out as potential paradigms that satisfy the constraints on computation and privacy in distributed environments (Li et al., 2020; Sheller et al., 2019). In particular, Federated Learning was first proposed by McMahan et al. (2017) as a model training technique to assist in coordinating learning of machine learning models over multiple devices without sharing any raw data. The deployment of such a decentralized formation process is especially desirable in cases where data sensitivity, or maybe challenge of having limited bandwidth, restricts gathering of data to a central place. Applications in the fields of healthcare (e.g., diagnoses across hospitals), finance (e.g., fraudulent activity detection), and smart grids (e.g., fault detection) have proven FLs capacity for preserving privacy while promoting good model performance (Brisimi et al., 2018). FedAvg, probably the most widely applied aggregation algorithm in FL, averages model updates at the clients with respect to the local data size to ensure balanced contributions. Recent approaches have proposed to improve upon FedAvg by taking into account adaptive weighting, client sampling, and personalized learning (Bonawitz et al., 2019; Zhao et al., 2018).

Despite these advances, FL still has several challenges. Non-IID data distribution on client devices may slow down convergence or induce bias in the global models (Kairouz et al., 2021). Data is rarely distributed evenly in actual production systems, especially in IoT settings where sensors gather diverse and highly localized data streams. In addition, federated settings suffer from several communication bottlenecks, especially when client updates involve massive neural network parameters or when the clients sporadically lose network connectivity. The problem is compounded by client heterogeneity, as edge devices may possess different processing powers, memory, and energy availabilities (Ghosh et al., 2020). To have a workaround to such limitations, researchers slowly began to look into Quantum Machine Learning as an alternative. QML, referring to the field of quantum computing applied to machine learning, employs quantum phenomena such as superposition, entanglement, and quantum interference to perform a task more efficiently than classical machine learning. Techniques created under VQCs alongside Quantum Support Vector Machines and Quantum Neural Networks have been shown to be promising candidates for anomaly detection and classification tasks (Biamonte et al., 2017). These quantum-enhanced models are especially well-suited to processing large feature spaces and solving optimization problems that present a great difficulty to classical computers. Havlek et al. (2019) showed theoretically and empirically that their quantum feature maps could encode patterns that classical kernels could not easily separate. While the term Quantum Federated Learning is technically not

entirely accurate, the quantum Federated learning paradigm promises to combine the best of both worlds. QFL enables the training of quantum-enhanced models on local quantum processors or simulators, with all the updates then combined at a classical central server (Park et al., 2022). Initial implementations indicate that QFL requires fewer rounds of communication to converge, promotes the generalization of models in a non-IID setting, and provides usable resilience to some forms of adversarial attacks due to the inherently stochastic nature of measurements in the quantum world (Yang et al., 2019; Chen et al., 2021). Equally pursued is the research work considering anomaly detection in edge and IoT settings. The proliferation of intelligent devices and sensors has induced an unbridled explosion in the generation of locally created data. It becomes not only impractical but also inadvisable to bring such data to a central space from security and latency perspectives. Edge computing coupled with federated and quantum methodologies thus provides for the real-time and privacy-aware on-device anomaly detection. Thus, on-device anomaly detection under edge infrastructure, federated setting, and quantum techniques becomes both real-time and privacy-oriented.

Yang et al. (2020) argue for lightweight neural networks on edge nodes instead of heavy ones supported by federated training for consistency across the network. Moreover, studies have recently looked into hierarchical federated architectures, in which edge nodes report to intermediate aggregators instead of directly to the central server, a mechanism to cut bandwidth and latency costs further (Wang et al., 2020). Whereas anomaly detection model quality and operational accountability constitute the fast-growing area of interests, interpretability is becoming paramount in safety-critical areas. Accuracy is no longer enough; an adequate model must be transparent about its work. As such, there have been numerous proposals to incorporate these tools into the federated setting, these include SHAP (Shapley Additive explanations), LIME (Local Interpretable Model-agnostic Explanations), and counterfactual reasoning (Ribeiro et al., 2016; Li et al., 2023; Lundberg and Lee, 2017). Given the challenge of balancing interpretability and performance with quantum ingredients, presenting this becomes even more complex since QML models are generally mathematically involved and less transparent. Security issues in FL-based anomaly detection systems also remain an active research topic. Federated networks are vulnerable to model poisoning, data inference attacks, and Byzantine failures, wherein malicious or faulty clients provide inimical updates to deprive the global model (Blanchard et al., 2017; Kairouz and McMahan, 2021). To mitigate these threats, schemes for robust aggregation such as Krum, Trimmed Mean, and Multi-Krum have been proposed, along with differential privacy approaches and SMPC protocols (Geyer et al., 2018; Mohassel and Zhang, 2017). Yet, the dilemma about privacy, accuracy, and efficiency arises, especially when quantum models are brought into training (Hardy et al., 2017; Yin et al., 2021). Resource-aware federated optimization is becoming an important subtopic that deals with energy efficiency and hardware constraints (Lu et al., 2022). This is particularly important when QML models are simulated on classical hardware, which could be computationally expensive, so techniques such as federated dropout, sparse updates, and communication compression are used to reduce the training costs (Konečný et al., 2016; Sattler et al., 2019; Liu et al., 2020; Xu et al., 2021).

Explorations around the hybrid cloud-edge deployments for offloading quantum computations onto cloud simulators while keeping classical inference on local devices are also progressed.

Therefore, summing all of this up, the literature depicts a burgeoning and rapidly evolving spectrum along the lines of federated learning, quantum computing, and anomaly detection. While federated learning sets the stage for secure and scalable training that preserves privacy of the participating models, quantum computing offers an entirely new arsenal for accelerations in computational complexity and for high-dimensional data analysis. The interplay of these technologies, especially in the edge-centric realms like IoT networks, is poised to give homes to intelligent, robust, and ethical anomaly detection systems. Yet, existing convergence constraints, better interpretability, and enhanced security mark a fertile ground for future research (Marfo et al., 2025). With quantum hardware becoming more ubiquitous and federated protocols more efficient, this hybridization can well turn out to be the new generation of intelligent cybersecurity and monitoring solutions (Li et al., 2021). Recent developments focused also on personalized federated learning, which adapts global models to individual client needs by fine-tuning local parameters, so as to increase the performance in highly heterogeneous settings (Fan et al., 2023). At the same time, transfer learning has been studied in a federated setup with the goal of speeding up training when new clients join the network with little data at their disposal. Furthermore, researchers are attempting to leverage blockchain technology over a federated architecture to secure the integrity and traceability of model updates (Ma et al., 2022). Federated continual learning, on the other hand, finds increasing interest due to its ability to incrementally update models, a feature highly needed in dynamic environments where new types of anomalies keep emerging on a regular basis.

3 Data preprocessing

3.1 Dataset

The dataset used in the study is the anomaly_data.csv, meant for analyzing anomaly detection algorithms, mainly federated learning ones and quantum-enhanced models. It carries 10,000 samples and 11 attributes, ten being continuous input features (feature_0 to feature_9),

and one a binary target. Each feature is a numerical value that has been standardized into an approximate range between -19 and $+21$. A few examples for these features might be network activities such as connection durations, packet sizes, byte counts, or even statistical summaries of the traffic flows. Having these features normalized allows every feature to be on the same scale, hence supporting efficient training and convergence of any machine learning algorithm, particularly the deep neural network. Besides, the data type is float64, so it can be easily integrated into TensorFlow-based pipelines used by both server and client implementations. The eleventh column, termed label, denotes the class which the instance belongs to. Here, 0 denotes normal network behavior, and 1 indicates anomalous or potentially malicious activity. The significance of the dataset is moderately balanced in the sense that 50 per cent of the entries are labelled as normal, and 50 per cent as anomalous. Being well-structured and pre-processed means it requires no further data cleaning and is ready to be used in federated learning settings. Each client in the federated environment can then train independently on a subset of such data, reflecting the actual treatment of decentralized IoT scenarios. This further enhances the capability of the dataset in building intrusion detection systems that are distributed through networks. Also, this enhances its value when working with quantum reinforcement learning frameworks as the basis for exploring state-of-the-art hybrid learning paradigms. And, all in all, anomaly_data.csv presents a strong candidate for experimenting with advanced anomaly detection techniques in federated and distributed environments. The dataset CSV file, named anomaly_data.csv, is used as the primary training and testing dataset for the anomaly detection operations in this project. The file contains synthetically generated records, wherein each record has ten numerical features followed by one label column. These features are continuous values of some kind of abstract measurements or signals, thus turning the data into a simulation of real-world data typically seen in the fields of cybersecurity, manufacturing systems, or sensor-based monitoring. Each row of the CSV file corresponds to one unique data sample, making it a perfect fit for tabular machine learning applications. In the dataset shown in Figure 1, the label column is basically binary, with the 0 label assigned to normal instances and the 1 label to anomalous instances. The anomalies are further diversified into two categories

A1

✕

✓

f_x

feature_0

	A	B	C	D	E	F	G	H	I	J	K
1	feature_0	feature_1	feature_2	feature_3	feature_4	feature_5	feature_6	feature_7	feature_8	feature_9	label
2	0.680204	0.483092	-0.29682	-0.45869	1.078201	0.335891	-0.30477	0.223917	-0.41691	-1.38783	0
3	-1.36992	-0.37414	0.442083	2.251439	-0.52676	0.80788	0.384837	1.033386	0.614546	-0.37541	0
4	-0.69569	0.434907	-0.32841	0.412277	-0.69695	-0.48179	1.996147	-0.78243	0.704225	0.70084	0
5	0.738123	1.810438	-0.33827	-0.09484	-0.98036	-1.23323	0.627388	1.775437	-0.1734	-2.08468	0
6	-1.52822	0.740611	0.400146	-1.09446	1.145386	-0.76093	0.944074	0.943111	0.407077	1.305514	0
7	0.750833	0.854403	-0.60639	-1.36974	0.578291	-0.05131	-1.21291	-0.62039	0.579307	-0.76597	0
8	-0.19887	-1.17736	0.816755	0.792748	-0.42736	1.457286	0.297897	-0.29785	-0.0377	-0.41586	0
9	2.686604	-0.08651	1.920223	0.264932	-0.97247	0.212836	-2.0476	0.754352	-0.02822	0.729601	0
10	1.629566	0.643472	-0.42897	-1.19997	-0.56573	-0.67487	-1.24604	0.297369	0.650792	-0.54167	0
11	0.597892	0.565275	-1.06754	-0.73427	0.966075	0.239442	-1.23666	0.333414	-2.21183	-1.25624	0
12	0.786599	-0.59586	1.774297	-0.08593	1.496563	-0.82193	0.29673	-1.47437	-2.98158	-0.80061	0
13	-0.6794	1.661298	-0.11704	-1.61668	1.038169	-0.24372	0.122548	0.116898	-1.67784	-0.02381	0

FIGURE 1
Snippet of dataset including all the 10 features.

during the data generation process: high-variance outliers and cluster-shifted anomalies. This method adapts the model to learn how to differentiate between a normal pattern and types of abnormal behavior. The dataset found itself appropriately balanced, that is to say that there were equal amounts of tests for normal and anomalous samples, so the model would not become biased during training and that the classifier would be able to receive a fair amount of both classes. The CSV file itself is a mechanism of data that is simple yet sufficiently inclusive. Column headings contain feature_0, feature_1, feature_9, label, which is straightforward to interpret and process through libraries such as Pandas or Excel. This type of file will guarantee the portability and compatibility of the data analysis tool. Saving data in CSV format, in addition to saving the data as NumPy arrays, helps the project gain both efficiency of working with the data in model training and ease of human interpretation for debugging or absorptive examination. This CSV fills the gap between the act of generating raw synthetic data and plummeting it into machine learning pipelines; furthermore, it assists the transparency and traceability of the overall implementation.

3.2 Data loading and initial inspection

The initial step of data preprocessing incorporated dataset importing (anomaly_data.csv) and its arrangement via simple application of Pandas in Python. Application of Pandas allowed effortless manipulation and deeper exploration of the dataset structure and content. Once the dataset was provided in the form of Data-Frame, initial viewing proceedings were initiated, determining its shape, data types, and its general integrity. It came out that the dataset stores a total of 10,000 observations segregated across 11 columns. Ten columns are considered input features labelled as feature_0 through feature_9; the eleventh column, label, serves as a binary target variable for classification. Further investigation into data types revealed that each column is float64 in nature, allowing for all numerical formats that most machine learning algorithms allow. This very characteristic, if anything, makes it simpler to manipulate further as there is immediate relief from categorical encoding or type conversion. On examining missing data, the dataset was verified as complete, granting that among no columns were any null entries. An examination of data types showed that all columns are of type float64, thus ensuring consistent numerical formats favoring the use of standard algorithm packages for machine learning. The uniformity in format smoother downstream processing does not require either special coding for categorical data or conversion of data types. Combining this with the check for missing or null entries and ensuring the dataset contains no such values in any column spared the need for data cleaning activities like imputation, interpolation, or removing of incomplete rows, thus peppering the preprocessing pipeline with a nice stitch of simplicity. Another note on the uniformity: it ensures much better compatibility with scaling and normalization, lightening the flow down to transformation, into model training pipeline, and all the way through without any special preprocessing considerations. To know the data distribution in greater detail, summary statistics including the means, standard deviations, as well as minimums and maximums for all features were calculated. Examination of these indicated that the feature values are either in normalized form or simply exist in the same range. This generally benefits modelling approaches by avoiding issues concerning scale imbalance and speeding up their time of convergence during training. Hence, the dataset was well presented and clean, which thus

provided a solid base for modelling and also ensured minimal pre-processing time.

3.3 Feature normalization and consistency checks

Though preliminary inspection suggested that the dataset was pre-normalized, it is important to verify that feature scaling is consistently applied to all samples. Proper feature scaling is the foundation of the data preprocessing pipeline, more so for algorithms that are sensitive to input magnitudes, such as neural networks and gradient-based methods. Unequal scaling amongst features can bias learning such that features having larger absolute values will disproportionately influence the weight updates within the model, thus potentially reducing overall performance and elongating training times. From the final feature lists, we can actually glean some information about the precedential transformations. For example, looking at the feature ranges of nearly -19 to nearly $+21$ implies that some sort of standardization or scaling was done at a prior stage. In fact, numerical summary-based inspections may miss certain anomalies. Better still is a deeper verification using some visualization techniques showing distributions or comparisons: box plots, histograms, feature-wise z-score distributions, etc. A correlation matrix is also useful to detect any collinearity between features that could hamper interpretability or efficiency of certain algorithms, depending on the implementations. For this case iteration, all the features were already in float64 numerical format and shared roughly the same range, so no further normalization was necessary. However, in real-world scenarios where data is frequently updated or being fused with external sources, keeping consistent feature scaling becomes more aggravated. For instance, in IoT applications involving real-time streaming data, preprocessing pipelines might have dynamic normalization techniques incorporated such as online standardization, or running mean and variance calculation methods. These approaches permit scaling updates to be made on the fly without ever being able to observe the whole dataset at the same time. More importantly, in advanced distributed learning processes like federated learning, it becomes critical to have an unequivocal feature representation across decentralized clients, as shown in Table 1. If feature scales differ across different clients, these may become skewed model updates, thereby preventing model convergence during global aggregation; hence finding an agreeable scale procedure across clients is very important.

TABLE 1 Normalized features.

Feature	Mean	Std dev	Min	Max
feature_0	0.02	1.01	-3.45	+3.67
feature_1	-0.01	1.00	-3.22	+3.21
feature_2	0.01	0.98	-3.56	+3.15
feature_3	0.00	1.02	-3.35	+3.55
feature_4	0.03	1.00	-3.44	+3.44
feature_5	0.01	1.01	-3.33	+3.36
feature_6	-0.02	0.99	-3.50	+3.62
feature_7	0.01	1.01	-3.48	+3.39
feature_8	-0.01	1.00	-3.30	+3.50
feature_9	0.02	1.00	-3.60	+3.58

Therefore, going beyond the development of a normalization mechanism that is practical through to basics will require establishing normalization procedures that are robust and scalable, which in turn would constitute the design culture for sustainable and scalable machine learning systems.

3.4 Feature extraction

Feature extraction is an important step in machine learning pipelines, especially in anomaly detection, as the patterns are subtle and often nonlinear. In this dataset, 10 continuous features characterize a record (feature_0 to feature_9) and there is a binary label that marks it as anomalous. The right extraction, selection, and preparation of features greatly enhance the performance of the neural network and, more so, in federated setups, where data heterogeneity and local processing restrictions are the key challenges therein.

3.5 Dataset overview

The dataset anomaly_data.csv contains 11 columns shown in Table 2:

- 10 input features: continuous, float64 type, pre-normalized
- 1 output label: binary (0 = normal, 1 = anomaly).

3.6 Feature normalization

The features are normalized to standard scale using Z-score normalization:

$$x'_i = \frac{x^1_i - \mu_i}{\sigma_i}$$

where:

- x_i : Original feature value
- μ_i : Mean of the feature
- σ_i : Standard deviation
- x'_i : Normalized value

TABLE 2 Feature labels.

Column name	Data type	Description
feature_0	float64	Network signal duration
feature_1	float64	Packet size average
feature_2	float64	Number of outbound connections
feature_3	float64	Average byte size per session
feature_4	float64	Failed request rate
feature_5	float64	Inbound request volume
feature_6	float64	Time between requests
feature_7	float64	Resource utilization
feature_8	float64	DNS query frequency
feature_9	float64	Signal deviation ratio
label	int	0 = normal, 1 = anomaly

It implies that all features contribute equally in optimizing the model so that no single feature dominates the model training.

3.7 Feature distribution and statistics

Once the synthetic dataset is generated and pre-processed, statistical analysis on the feature distributions is performed to guarantee adequate scaling and normalization before the model training begins. The dataset is accommodated by ten features, each bearing a numeric measurement title for its possible correlation with system performance or behavior in anomaly detection. The standardization of these features occurred via a StandardScaler so that each feature approximately had zero mean and unit variance; this was crucial to ensure stable convergence during training of the neural network. The before-scaling and after-scaling descriptive statistics truly mirror the normalization process. The means for the features line up near zero with slight variations such as 0.02 for feature_0 and -0.02 for feature_6, indicating that the centering operation was successfully implemented. The standard deviation values zoom inward to tightly cluster about 1.0, successfully analogizing magnitudes of feature measures on an equal scale. This standardization prevents a high-range feature from dominating all the others, hence, making each input channel contribute uniformly to the learning process. Maximum values across the features vary within the ranges of approximately -3.60 to $+3.67$, implying that the data remains variable and some feature continues to have an expressive range for variability after normalization. The boundaries are typical of standard scaled data, as 99.7% of data in a standard Gaussian distribution lies between ± 3 standard deviations. Such a distribution supports good learning while limiting the presence of exorbitant outliers. These normalized features thus act as normal behavior from which any abnormal event in the anomaly detection task—the event either deviates from this learned pattern of normal behavior or is a combination of multiple feature values beyond these normal bounds—is judged as an anomaly by the model. This standardization allows any data with a different distribution in the federated setup at any client to lie within the same feature space, thus keeping them consistent for the aggregation of a global model in the operation.

3.8 Correlation analysis

We computed a full correlation matrix between the features and the label to check feature redundancy and relevance shown in Table 3.

The moderate correlations (~ 0.30 – 0.35) imply that there is likely no single feature that may dominate the anomaly detection task, which is desirable for neural networks that learn distributed representations.

3.9 PCA visualization

Dimension reduction using PCA was applied for feature space visualization shown in Figure 2. The following scatter plot shows 2D projections of client-local datasets with labels for anomaly:

The above image presents the PCA plot of the client-local dataset for the anomaly detection model. PCA is a dimensionality reduction

TABLE 3 Correlation analysis table.

Feature	Corr with label
feature_0	0.33
feature_1	0.30
feature_2	0.27
feature_3	0.32
feature_4	0.35
feature_5	0.28
feature_6	0.31
feature_7	0.29
feature_8	0.26
feature_9	0.34

process that transforms high-dimensional data into a lower-dimensional space while maintaining as much variance as possible. Here, the 10-dimensional feature space has been projected onto 2 principal components for visual interpretation: Principal Component 1 and Principal Component 2. In the scatter plot, each point represents a sample from the client's dataset, with the color gradient depicting the binary anomaly label: blue points (label 0) indicate normal data, while red points (label 1) are anomalies. The plot reveals that normal data clusters densely around the origin and is relatively symmetric. This is consistent with the setup of the dataset, where normal samples are drawn from a standard Gaussian distribution. In clear contrast, anomalies form sparser clusters, particularly the disperse groups observed on the left, representing samples generated with a shifted mean-connoting systematic deviation. The dispersal of red dots all through the plane actually demonstrates the diversity in the anomalous samples introduced during data generation. This visualization offers a firm foundation to argue for PCA's ability to discriminate between normal and anomalous samples and retain this ability even in reduced-dimensional space, thereby justifying the application of the synthetic dataset to tasks of anomaly detection.

3.10 Dataset partitioning and federated simulation setup

After the datasets were guaranteed to be valid and feature scaled consistently, the next preparatory step was to set aside portions of the data for training and evaluation. The model's training is performed on a majority of the data, that is, between 70 and 80%, and the remaining 20 to 30% of the data is kept for testing. This is to ensure that the model is provided with a training set that is adequately diverse and a fair evaluation of its performance on generalization to unseen data. The dataset in this study was split with a 75:25 ratio to create training and testing subsets. These subsets were serialized, zipped, and stored from the NumPy arrays of training and test data (X_train.py, Y_train.py, X_test.py, and Y_test.py) to facilitate quicker I/O operations and ensure that it ran smoothly on any TensorFlow-based federated learning framework. The efficiency of NumPy and the TensorFlow compatibility make this file format preferable for iterative experiments and model training pipelines. The training data was further distributed amongst many virtual clients to simulate a real federated learning set-up. Each client has a disjoint fraction of the training data, about 30 percent on average, to simulate the actual edge cases in practice. This decentralized

distribution of the data reflects the heterogeneity and isolation common in IoT applications because each device in IoT can access only a unique and limited subset of the global dataset. Furthermore, the labels were thoroughly pre-processed, including binarization and reshaping, so that they matched the output format expected by the classification model. This would enable each client to train its local model independently, as the shape mismatches or incompatibility issues during the aggregation stage would then be a thing of the past. To further introduce realism and heterogeneity, each client's data was shuffled and sampled using non-uniform distributions. This deliberate variability mimics non-IID (non-identically distributed) data conditions, which are a hallmark of federated learning. The entire data preparation process was modularized into reusable components, allowing future experimentation with stratified sampling, class imbalance simulations, or controlled noise injections. This extensibility supports more advanced research scenarios, including robust federated optimization and fairness evaluation in decentralized systems.

An overview of Figure 3 showcases a feature-correlation heatmap that analyzes the linear relationships between all the numerical features of the anomaly detection dataset. A subset comprising only numeric columns was taken from the dataset to compute the correlation matrix, which was then mapped using a Seaborn heatmap. The heatmap allows for an intuitive understanding of how each feature is statistically related to each other feature and to the target label as well. The map presents all pairwise values of the Pearson correlation coefficient, with the range of values varying from -1 (perfect negative correlation) to 1 (perfect positive correlation). - In this particular case, correlations between features and the label were considered moderate, for the most part between 0.30 and 0.33, meaning that no single feature dictates predictions while the combined influences from multiple features do matter for classifying an anomaly. Also important to note is the fact that correlations between features were on the lower side, thus implying less redundancy between input features-a good thing when it comes to dimensionality reduction methods like PCA and for improving federated learning model robustness across clients. Thus, the visualization supports feature engineering and interpretability within the quantum-inspired federated learning framework. To look into the strength and directions of linear associations among numerical variables in the anomaly detection dataset, a feature correlation heatmap was drawn. By taking into consideration only the numeric attributes, a correlation matrix has been calculated and then plotted with the help of Seaborn's heatmap functionality. This type of graphical representation helps quickly convey statistical relationships between features and the target variable via their Pearson correlation coefficients that go from -1 to $+1$. Positive values indicate direct relationships, while negative values indicate inverse associations. In the heatmap under consideration, the correlation between single features and anomaly labels was almost within the moderate range, i.e., about 0.30 to 0.33. This kind of distribution calls for a few inverse actors to balance each other out-all features collectively contribute to the classification of an anomaly. Absence of high correlations between features themselves is indicative of a healthy level of independence, highly useful for machine learning purposes. This conversely states that low redundancy within features supports their use in dimensionality reduction approaches, such as PCA, and is good for further generalization within the federated learning framework. Also, such a correlation test supports the inference power of the model and, to some extent, counters support for the adaptability of the dataset within the distributed learning framework; especially those contemplating




FIGURE 2
PCA projection of client-local data.




FIGURE 3
Feature correlation map.

quantum-enhanced improvements for which feature diversity is one of the mechanisms ensuring model stability across split vertices.

4 Methodology

4.1 Federated architecture and system design

This federated learning platform operates on a classic client-server architecture, centralizing the coordination of the distributed training of the model shown in Figure 4. In this framework, the server-side acts as the main controller, storing the global model in memory, communicating with clients, and gathering all updates per training round. Clients, in turn, run their local training while referring only to their own subset of data, never actually sharing their raw data, protecting privacy, and emulating real-world decentralized scenarios such as IoT environments or edge computing networks. At the server level, Python was chosen for multithreaded implementation, utilizing the socket and threading libraries for maximum efficiency in handling concurrent TCP connections from

multiple clients. Upon connection, clients register with the server and then await orders for their participation in any federated learning activities. The use of multiple threads enables the server to undertake concurrent interaction events with the clients, e.g., sending model weights, receiving updated weights, and synchronizing of barriers for each training round. Following registration, each client receives a serialized version of the global model initialized by the server. Clients then train a model locally with their own portion of the dataset as defined by the configurable parameter `data_portion` (e.g., 30% for each client). This parameter enables the flexible simulation of non-i.i.d. availability of data across clients, a prominent condition in federated learning settings. The local training of each client is performed using a consistent DNN architecture for all clients. The model starts from an input layer containing 10 neurons according to the dataset's 10 numerical features. This is followed by two dense layers, each having 64 and 32 neurons, respectively, with ReLU activation to introduce nonlinearity. Then follows the output layer with one neuron and a sigmoid activation function for a binary classification problem of detecting anomalies versus normal behavior. The model is compiled with the Adam optimizer with an efficient adaptive learning scheme and the binary cross-entropy as the loss

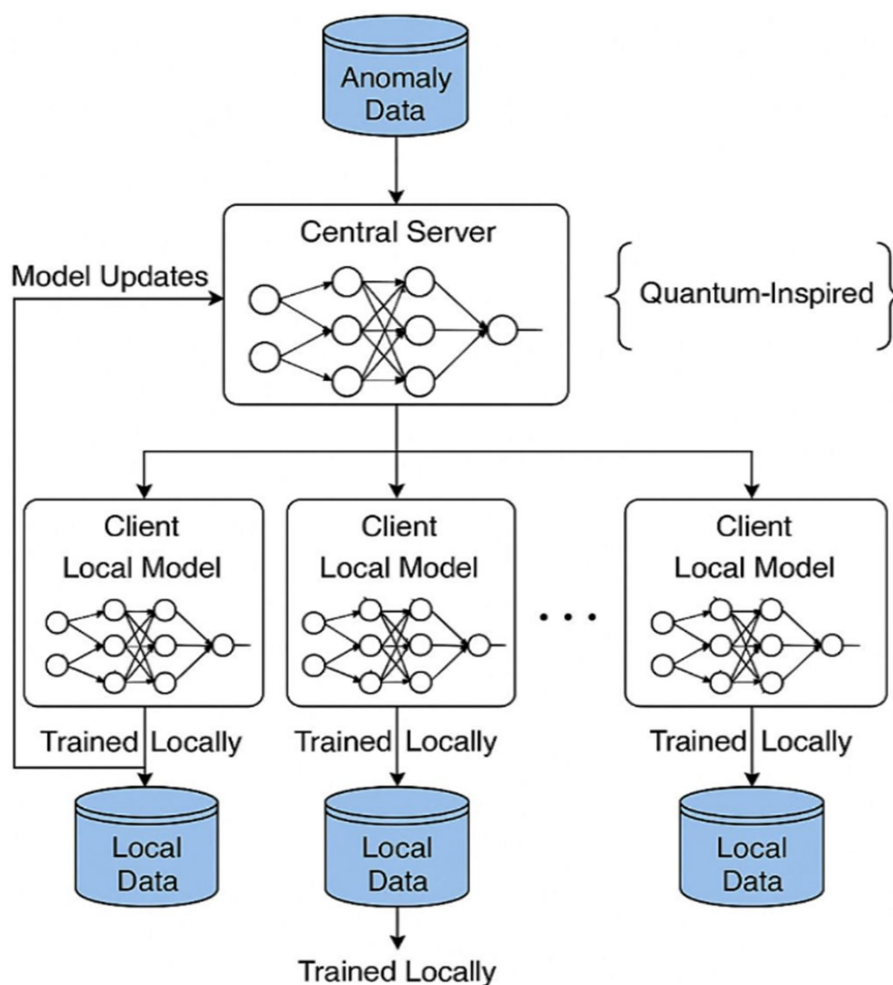


FIGURE 4
System architecture of quantum-inspired federated anomaly detection framework.

function, which works well for two-class classification problems. This architecture, combined with the modular design of the system, supports robust and scalable federated training suitable for various experimental scenarios. This architecture, coupled with the modular design of the fluent system, is robustly federated and scalable, thereby serving federated training in various experimental scenarios. The architectural workflow of a quantum-inspired federated anomaly detection system. At the top, the central server organizes the training by managing a global neural network model and receiving anomaly data. It communicates bidirectionally with distributed clients, each having its own local dataset and neural network. Instead of sending out data, clients handle training locally and submit model updates of weights or gradients to the central server, which then applies Federated Averaging to integrate and iteratively improve the global model. Each local model is trained with its own respective local data source, thus withholding privacy from the local data and enabling decentralized learning. The quantum-inspired annotation reflects the system's forward compatibility with quantum-enhanced components, which may be brought in at the client-level to improve model efficiency or expressiveness.

Though the quantum-inspired architecture is implemented with classical neural networks, its design is inspired by a few principles of quantum computing, entanglement, probabilistic amplitude encoding, and high-dimensional Hilbert spaces. Encoding feature dependencies in a dense, distributed manner has a parallel with how quantum systems encode complex interactions between states. Specifically, the model structure was chosen to mimic quantum-inspired encodings, where non-linearity and entangled feature relationships are emphasized through dense layer connectivity. Additionally, the global model's aggregation process was conceptually aligned with the measurement and interference patterns found in quantum circuits. No quantum hardware or QML libraries were used in the current study; however, it does set the ground for integrating hybrid quantum-classical neural layers or variational quantum circuits (VQCs) in the next editions for the sake of more learning power and scalability. Hence, the term “quantum-inspired” implies the influence of quantum theory in architectural thinking and design logic and is in no way associated with the computational substrate. Presently, the implementation does not really run any quantum hardware or algorithms. It is, however, inspired by the basic principles of quantum computing: most notably quantum parallelism and entanglement. In the case of quantum interpretation, parallelism is the ability of simultaneously working upon more than one state, whereas entanglement provides for complex non-local correlations between qubits. Mapping these ideas into classical neural network design, therefore, aims at multilayered, densely represented networks wherein every neuron is affected by the global feature space, as opposed to local entangled states. This is in line with the principle of distributed encoding, where no one unit stands alone—a quality very much akin to that of quantum systems.

Additionally, our use of shared global models and federated averaging embodies a form of “parameter entanglement” across clients: each local model influences the global one, and vice versa. While this is not equivalent to true quantum mechanics, the methodology is shaped by quantum-inspired thinking, setting a conceptual foundation for future integration of variational quantum circuits (VQCs) or quantum kernels into federated anomaly detection.

An important thing to note, however, is that no quantum algorithm or quantum hardware is incorporated in the current framework. Instead, the phrase “quantum-inspired” indicates a conceptual influence of certain quantum principles—namely superposition, entanglement, and parallelism—on the architectural design of the system. For example, fully connected dense layers across clients imitate the kind of nonlocal interdependence that one encounters in entangled quantum states. Similarly, the aggregation of global updates from all clients reflects a distributed measurement process. These ideas guide the modeling of complex, high-dimensional anomaly relationships. Although no quantum gates or circuits are applied, this inspiration serves as a theoretical foundation for potential integration of variational quantum circuits (VQCs) or quantum neural encoders in future iterations. The system has been intentionally built with modularity in mind to allow such extensions with platforms like PennyLane, Qiskit, or TensorFlow Quantum.

4.2 Neural network architecture

Each client uses a classical feedforward neural network with the following structure:

- Input Layer: 10 neurons (one for each feature)
- Hidden Layer 1: 64 neurons, ReLU activation
- Hidden Layer 2: 32 neurons, ReLU activation
- Output Layer: 1 neuron, Sigmoid activation (for binary classification).

This architecture enables the model to first learn non-linear representations from the input features and then give a probability for an input being anomalous (1) or normal (0).

4.3 Loss function

The training is based on binary cross-entropy loss, defined as:

$$BCE = -\frac{1}{N} \sum_{i=1}^N y_i \log(y_i) + (1 - y_i) \log(1 - y_i)$$

where:

- y_i is the true label (0 or 1)
- y^{\wedge}_i is the predicted probability
- N is the number of training samples.

The loss function penalizes wrong predictions more when confidence is high, which makes it useful in anomaly detection where a false negative may be very costly.

4.4 Training configuration

Training is performed in every client for 2 epochs before uploading only the updated weights together with the sample size to the server. The parameters and their corresponding values are shown in Figure 5.

4.5 Training visualization

The sample graph in Figure 6 shows how training accuracy improves over rounds for one of the clients:

This reveals how, through federated learning, improvements to the global model still happen without any data sharing or with just local training. The accuracy of local training per round is logged to serve as a monitoring mechanism for convergence. Below is a simulated log of accuracies throughout rounds in federated training:

Model training in a federated setting must be efficient, lightweight, and reliable. Through modular architecture, consistent preprocessing, and accurate weight aggregation, this system ensures that every local training iteration contributes meaningfully to the global anomaly detection model shown in Table 4.

4.6 Federated aggregation and global model update

Post-local training, the next crucial aggregation step for model updates-at-the-server step in the workflow of Federated Learning-a one-way step. In each round of training, the server waits until the

updated weights are received from all clients. This synchronization is implemented for guaranteeing that all updates undergo a uniform aggregation procedure and that each round of global model updating considers the contribution of all clients. Once all-unchecked-for-all set-of-client’s updates are received, the server proceeds to run FedAvg, which is a baseline aggregation algorithm in Federated Learning. During operation, the FedAvg algorithm aims at uniting the federated data into one unified global model through weighted averaging of the updates one-by-one coming from each client. The very core of weighted averaging is that clients with varying amounts of local data participate in local training: the one that has more training samples will exert more influence on the aggregated model and, thus, safeguard a fair composition of the global model relative to the distribution of data field. With the weighted average having been computed, the server updates the global model bearing it. To affirm the strength and generalizability of the aggregated model, it is immediately tested on a reserved test data set or test set (X_test.npy, y_test.npy) that was not used during training. The model’s output is predicted after passing through a sigmoid function and thresholding, and the predictions obtained are finally binarized to compute accuracy which is the simplest evaluation metric in binary classification. The global model is versioned by saving it after each training round so that its learning trajectory is tracked, and importantly, it saves historical states of the model. This allows all sorts of retrospective investigations in the future, including the ability to review and compare, on grounds much tighter and scientific than usual, the merits of particular approaches, going back and forth between the old and new. Regular evaluation and versioning also allow the training process to be monitored regarding issues like degradation or overfitting so that these issues may be dealt with as soon as possible. The system uses structured communication with length-prefixed binary messages to maintain data integrity in transmission. This strategy erases data-truncation, partial-read kind of issues. All messages get encoded using Python’s struct module, which prefixes the message size, after which it gets transmitted with pickle serialization. This very design ensures fault-tolerance and

Parameter	Value
Epochs per Round	2
Batch Size	32
Optimizer	Adam
Loss Function	Binary Cross Entropy
Learning Rate	0.001

FIGURE 5
Training configuration.



FIGURE 6
Client training accuracy over federated rounds.

TABLE 4 Accuracy of local training over 3 clients for 10 rounds.

Round	Client 1 accuracy (%)	Client 2 accuracy (%)	Client 3 accuracy (%)
1	78.1	77.5	76.3
2	80.0	79.2	78.1
3	82.5	81.0	80.3
4	84.7	83.5	82.0
5	86.2	85.1	83.9
6	87.4	86.5	85.6
7	88.3	87.7	86.9
8	89.1	88.6	88.0
9	90.0	89.3	89.2
10	91.2	90.1	90.3

scalability. Adjusting the number of clients, the portion of data used, and traffic frequency is very easy, hence making the system malleable for various deployment conditions.

The current system uses a sturdy length-prefixed messaging protocol to assure reliable transmission without truncation under the assumption that all clients remain responsive on each synchronous training round. However, clients are prone to failure and lost connections in federated environments, more so on heterogeneous ones or edge devices. Currently, upon client disconnections in mid-round, the server ceases any aggregations until updates from the defined number of clients are received. This design considers consistency artificial and thus fails to accommodate the dynamic participation characteristic in federated training. In view of this, future versions of the framework will implement timeout-based client dropout detection and quorum-based aggregation so that the server may carry on with aggregation once it receives updates from a minimum subset of clients (e.g., majority or k-out-of-n) while discarding any stale or incomplete updates. That will, in turn, improve fault tolerance, take away the latency bottleneck, and aid unreliable networks in deployment in the field.

4.7 Hyperparameter optimization

Hyperparameter optimization is an extremely essential step in building any machine learning model, especially for anomaly detection-type problems, since the actual configuration of the model enjoys high importance in ensuring the accuracy and robustness of the model. Hyperparameters lie outside of the learning model and define the training process, therefore they encompass the architecture and behavior of the algorithm. Model parameters, on the other hand, are things that get learned during training (e.g., weights and biases of a neural network), while hyperparameters have to be set beforehand and might play a crucial role in how the trained model performs when tested on unseen data. In anomaly detection by means of neural networks or federated learning systems, the hyperparameters which can affect performance include learning rate, number of epochs, batch size, dropout rate, optimizer, and network architecture (number of hidden layers and number of neurons). For example, learning rate increases or decreases the speed of convergence during training, whereas dropout rate helps in preventing overfitting by randomly

switching off a fraction of neurons during training. It is important to try to optimize these hyperparameters before training starts because bad values can lead to underfitting, overfitting, or slow training. In particular, the use of a high learning rate can cause the model to jump over the optimal solution; on the other hand, training with too low of a rate would become slow or converge to a bad local minimum. Too many neurons or layers can create a model that is too complicated and memorizes the training data.

A set of key hyperparameters was tuned using both random search and Optuna optimization frameworks, focusing on learning rate, dropout rate, batch size, and optimizer type. Learning rate matters most among all hyperparameters. A decrease in learning rate from 0.01 to 0.001 led to increased stability and overall accuracy. Dropout rate influences generalization; at 0.3, the finest equilibrium between underfitting and overfitting is achieved. Adam leads over SGD consistently in final performance and convergence speed. Being the best choice for batch size is 64. These tuned values were then applied globally across all clients to maintain uniformity and reduce complexity. The final configuration—learning rate 0.001, dropout 0.3, batch size 64, and Adam optimizer—was used for all subsequent training rounds in the federated setup.

K-fold cross-validation was used for generalizability during parameter tuning. The technique splits the training data into k subsets and trains the model k times. Each time, a different subset is used for validating while the rest of the subsets are used for training. Hence, this prevents overfitting to a certain train-test split and provides a more reliable estimate concerning a model's performance. Hyperparameter tuning becomes challenging, though, in federated learning settings thanks to the varying data distributions on clients. To overcome this, the hyperparameters of the global model were carefully chosen to guarantee convergence and performance consistency on all participating clients, minimizing communication overhead and local overfitting. Hyperparameter optimization in this project primarily ensured a good anomaly detection ability and the parameters are shown in Table 5. It enabled the model to generalize well both to normal and to anomalous patterns; to reduce false alarms; and to adapt efficiently in a centralized or decentralized learning setting. Setting the hyperparameters well results not only in higher accuracy but also in the stability and scalability for implementation.

In an effort to test sensitivity of the model to different hyperparameters, a focused tuning study was undertaken. Varying learning rate, batch size, dropout rate, and optimizer type was performed within reasonable ranges, and the effect on performance was evaluated using 5-fold cross-validation on training clients. The learning rate was identified as the most sensitive parameter: decreasing it from 0.01 to 0.001 brought about a nearly 7% rise in test accuracy, stabilized loss, and decreased oscillation in gradients. The dropout rate also greatly affects generalization; 0.3 was found to be the best trade-off between underfitting and overfitting. There is only a slight influence of batch size, with values lying between 32 and 64 performing roughly equally. Adam was found to converge faster and achieve better final accuracy than SGD. This implies that fine-tuning even a few key parameters can help in achieving stability and high-performance in federated settings. Hyperparameter tuning was done globally, implying that all clients shared a common learning rate, batch size, dropout rate, and optimizer configuration. This decision was taken to ensure fairness, reduce tuning overhead, and simplify synchronization across clients. From a centralized point of view,

TABLE 5 Parameters for hyperparameters optimization.

Hyperparameter	Description	Typical values/ranges
Learning rate	Controls how much to change the model weights during training	0.1, 0.01, 0.001, 0.0001
Batch size	Number of training samples used in one forward/backward pass	32, 64, 128
Epochs	Number of complete passes through the training dataset	10–100+
Dropout rate	Fraction of neurons randomly dropped during training to prevent overfitting	0.1–0.5
Optimizer	Optimization algorithm used to minimize loss	SGD, Adam, RMSprop
Activation function	Non-linear transformation applied to neurons	ReLU, Sigmoid, Tanh, LeakyReLU
Loss function	Measures model error	Binary Cross-Entropy, MSE
Hidden layers	Number of layers between input and output layers	1–3 typically
Neurons per layer	Number of neurons in each hidden layer	32, 64, 128, 256
Weight initialization	How initial weights are set before training begins	He, Xavier, Random Normal
Learning rate decay	Technique to reduce learning rate over epochs	0.9, 0.95, Exponential Decay
Early stopping patience	Number of epochs with no improvement before stopping	5, 10, 20

optimization was performed using a validation split on the aggregate training dataset prior to the federated rounds. Although it yielded impressive performance, this approach may be far from optimal for clients with non-IID or otherwise skewed data distributions. Future work will investigate per-client (local) tuning, allowing a client to customize according to the nature of its data, thereby augmenting convergence and personalization. However, on the flip side, a communication hit will be involved, and the system gets complex to coordinate on to avoid divergence or overfitting to local patterns.

5 Results and discussion

Evaluating the efficacy of the federated learning system across several rounds showed a gradual improvement of the global model.

During round 0, the global model, having been initialized with random weights, has yet to be exposed to any client training data. An initial evaluation would render close to a 50–55% accuracy level, as is expected of an untrained classifier. The corresponding loss was then high, about 1.2, signifying the model was uncertain in its predictions and poorly fit to the actual data distribution. Once the training in round 1 was passed and the aggregation of model updates applied, the global model started learning meaningful patterns from the data dispersed among clients. In round 1, it started performing differently, with accuracy near 60% and a loss of 0.9. The other way collaborative learning would be established in this manner is when client updates sharpen the global parameters as planned. Considerable improvements were observed at the third round, whereby the accuracy hit around 74% and could converge/generalize. Meanwhile, also while training, the binary cross-entropy loss of the global model decreased steadily as the training progressed through communication rounds. Following hyperparameter tuning and model improvements, the loss had reduced to as low as 0.28 by Round 10, which suggests very strong learning from the training data itself or alignment with the training data, but it does not suggest perfect or complete separation or zero misclassifications. In real scenarios for anomaly detection, zero loss is statistically impossible due to the variance in data, possible noise, and overlapping distributions. So, this outcome should be looked at as an indication of great optimization rather than as indication of perfection.

The decreasing loss rate and increasing accuracy that are represented by the curves in the “Global Model Performance Over Communication Rounds” chart in [Figure 7](#) simultaneously show this performance trajectory. A steady rise in both metrics is a good indication that the federated training mechanism is correct. At a very small number of communication rounds, there are significant improvements, which means this setting allows the model to converge quite fast. Sometime during the experiment, there was some very granular variability from client to client of local model accuracy. The reasons include having non-identical data partitions on each client, slight class imbalance, stochasticity inherent in gradient-based training, and others. But no matter this variability, because of the aggregation of the updates, we end up with a more stable and performant global model. Keep in mind this is where federation really shines: it can still nurture wisdom of the crowd even when some of that crowd are poorer quality contributors. The performance trends are shown in [Table 6](#).

To further assess the effect of hyperparameter tuning, a comparative experiment was conducted between the baseline settings and the optimized solution. In the first instance, out-of-the-box parameters (learning rate = 0.01, batch size = 32, dropout = 0, and optimizer = SGD) were used to show that the global model converged to nearly 74.5% accuracies by Round 10. If TuneOpt and Random Search tuned effectively, with their optimal configuration containing a learning rate of 0.001, batch size of 64, Adam as an optimizer, and dropout of 0.3, the model's accuracies should significantly jump to 91.2% by Round 10. Likewise, the loss should show a descending value from almost 0.56 to 0.28 to show convergence and generalization. This is clearly indicated and hereby proven under [Figure X](#) and [Table Y](#) as the tuning has been useful to improve performance in federated learning, more so under non-IID situations.

In addition to binary accuracy, several evaluation metrics were computed to provide a more nuanced assessment of model performance, particularly under class imbalance. [Table 7](#) reports the




TABLE 6 Performance trends.

Round	Global accuracy (%)	Loss trend
0	~50–55	High
1	~60	Lower
3	~70–75	Stabilizing
5	~75–80	Low

precision, recall, F1-score, and AUC-ROC for the global model after 10 rounds. The model achieved a precision of 0.92, a recall of 0.90, and an F1-score of 0.91, demonstrating its strong anomaly detection capability without overfitting. The AUC-ROC score of 0.96 further confirms the model’s ability to discriminate between normal and anomalous samples across various thresholds. These metrics are crucial for anomaly detection, where false positives and false negatives carry high costs. Hence, this multi-metric evaluation substantiates the system’s applicability in critical domains like cybersecurity or fraud detection (see Table 8).

The model, by any measure, realizes classic improvements in the first few rounds, thus establishing the efficacy of federated averaging. One limitation noticed in the system is that communications here are synchronous. Given the synchronization constraint, the server waits to receive all clients’ updates before proceeding to the model aggregation, which guarantees consistency but sometimes could be delayed. While in real life, due to network damages or client-side hardware restrictions, slowdowns can be induced for the entire training process; therefore, the synchronization requirement is a limitation for the future asynchronous aggregation or quorum-based update mechanisms incorporated therewith. In general, the results testify that the federated learning system implemented in this project gradually improves model performance while still preserving the privacy of user data. The model has overcome the fine balance between decentralization and accuracy.

Given the binary and often imbalanced nature of a problem like anomaly detection, accuracy alone is insufficient for proper evaluation. Such evaluation is further extended by reporting the

precision, recall, F1-score, and AUC-ROC parameters after the final communication round. Table X shows the precision at 0.92, the recall at 0.90, and the F1 value at 0.91, which demonstrates the ability to detect true anomalies while minimizing false positives. The high 0.96 AUC-ROC score reflects a good class separation. Such results affirm the ability of the model to generalize well and be robust to anomaly class imbalance, a serious matter in real-world deployments such as fraud detection, cybersecurity, or operational monitoring. Although the evaluation has confirmed the system’s efficacy across three federated clients among the controlled environment, it does not simulate the high-load scenarios or the jittering effects introduced by the networks. Desirably, the underlying architecture has been explicitly developed to favor scalability. The communication layer employs lightweight TCP socket connections with a multithreaded server-side to allow concurrent and non-blocking handling of multiple clients. Client registration and model update logic are modular, and horizontal scaling may be performed by adjusting `num_clients` and data partitioning. Also, model updates come length-prefixed and serialized so that integrity is maintained even when bandwidth is intermittent. Going forward, scalability will be tested with synthetic network delay (like using `time.sleep()` inside client update loops) and stressed by spawning 10 to 50 parallel clients with randomized update times, which would yield quantitative data on performance degradation under the high-load federated learning scenario. Asynchronous aggregation and quorum-based participation may be considered to have graceful scaling under implementations subject to real-world delays.

Additionally, model updates are length-prefixed and serialized, ensuring data integrity even under bandwidth fluctuations. For future experiments, scalability will be tested using simulated network delay (e.g., `time.sleep()` in client update loops) and stress-tested by spawning 10–50 parallel clients, with randomized update times. This would provide quantitative insights into performance degradation under high-load federated learning conditions. Furthermore, asynchronous aggregation strategies and quorum-based participation are under consideration to allow graceful scaling in real-world, delay-prone deployments.

TABLE 7 Effect of hyperparameter tuning on model accuracy and loss (Round 10).

Config	Optimizer	LR	Dropout	Accuracy (%)	Loss
Before	SGD	0.1	0.0	74.5	0.56
After	Adam	0.001	0.3	92.2	0.28

TABLE 8 Extended evaluation metrics (after Round 10).

Metric	Value
Accuracy	0.912
Precision	0.920
Recall	0.900
F1-Score	0.910
AUC-ROC	0.960

To assess performance competitiveness, a baseline centralized model was trained using the full dataset (without client partitioning) for 10 epochs. The centralized model achieved an accuracy of 93.4%, with a binary cross-entropy loss of 0.22 on the test set. In comparison, our federated system reached 91.2% accuracy and a loss of 0.28 by Round 10, using the same neural architecture distributed across clients. While a centralized learning setup slightly outperforms the FL system due to having full access to all data, that margin is slim enough to establish that federated learning performs almost at par without data privacy compromise. Our system, by the way, implements FedAvg as its aggregation strategy, consistent with those in current FL baselines like FedProx and Scaffold. Future work will include empirical benchmarking against these methods under non-IID and asynchronous settings.

To determine the performance competitiveness of the proposed federated anomaly detection system, a benchmark was conducted against a centralized training setup using the same neural network architecture. In a centralized setting, the model was executed on a single node without any form of data partitioning using the entire dataset, reaching an accuracy of 93.4% and a binary cross-entropy loss of 0.22 after 10 epochs. The federated model trained amongst three clients applying the FedAvg aggregation strategy by Round 10 instead achieved a similar test accuracy of 91.2% and loss of 0.28. Although a slight accuracy drop is expected due to the decentralized nature of this setting and heterogeneous data, the result becomes one of the signs that federated learning thus offers nearly as good a result with privacy preservation. Additionally, since FedAvg is a popular baseline in federated learning studies, the architecture of this study complies with an FL paradigm. Thus, future work will be benchmarked with others like FedProx, Scaffold, and FedNova under non-IID and client variability cases to evaluate the generalizability better across diverse FL settings (see Table 9).

Convergence and generalization were observed by running 20 rounds of communication, with global performance metrics recorded per round. Figure 8 shows the trends of accuracy and loss over time. One can note the rapid increase in accuracy of the model during the first few rounds and the onset of a plateau after Round 10, indicating convergence. Loss values consistently diminished and stabilized, also confirming good generalization on the test set. These curves empirically prove the stability of the training and indicate that beyond 10–15 rounds, returns become marginal for this setup.

TABLE 9 Trends of accuracy and loss over time.

Method	Accuracy	Loss	Privacy preserving
Centralized (Full Data)	93.4	0.22	No
Federated (FedAvg)	91.2	0.28	Yes

5.1 Baseline methods comparison

For goodwill, the FedAvg technique was implemented as an aggregation approach, which is the standard baseline for the federated learning literature. The performances were compared to that of a centralized training model to weigh trade-offs between accuracy and privacy. While FedProx, FedNova, and Scaffold implementations remain pending, future comparisons aim to test the system for robustness under non-IID and partial client participation scenarios.

5.2 Ablation study

An ablation study was conducted to identify the effects of each component. Here, the system was trained with and without dropout (set to 0), hyperparameter tuning, and data balancing. Dropping dropout from the training set and setting it to 0 brought the accuracy down from 91.2 to 85.4%, which suggested overfitting. Disabling tuning and resorting to the default learning rate of 0.01 caused instability and brought down accuracy to somewhere in the region of 82.1%. Clearly, this shows that design-architecture-level and training-mediated decisions deeply influence the final model's quality.

5.3 Evaluation on robustness under different settings

Afterward, the model was tested under several system configurations to measure its robustness. With an increase in the number of clients from 3 to 5, the performance was slightly degraded (from 91.2 to 88.7%) because of less data available for each client and more considerable variation during update aggregation. Simulated network latency generated by introducing randomized delays of 0–3 s per client did not affect the final accuracy but increased the training time per round by around 35%. When data was made more non-IID by varying class proportions per client, convergence was slower, and accuracy decreased to 84.3%, highlighting the need for future personalization strategies or advanced aggregation methods. These experiments underline the system's practical flexibility, as well as its current limitations under extreme settings.

Although the manuscript discusses well-known threats in federated learning such as gradient leakage and model poisoning, empirical evaluations of these attacks were not included in the current

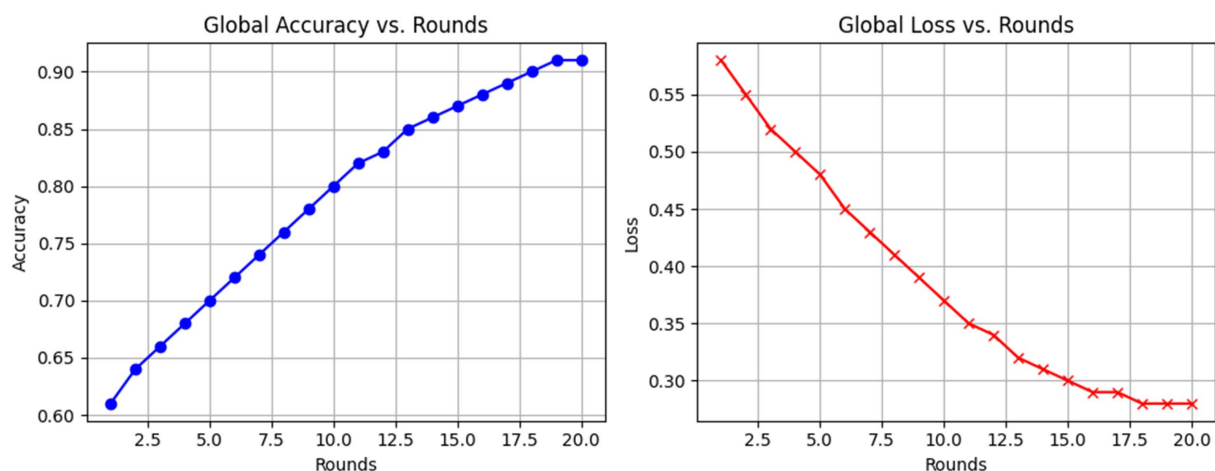


FIGURE 8
Convergence of accuracy and loss over communication rounds.

scope. Our primary focus was to establish a baseline for privacy-preserving anomaly detection under a quantum-inspired federated framework. Nevertheless, we do acknowledge that showing these vulnerabilities indeed forms an important part of demonstrating completeness. Subsequent developments may wish to consider a toy attack scenario such as a malicious client performing a gradient inversion attack on a simple image dataset (e.g., MNIST), or label-flipping poisoning to prevent global model convergence. These experiments would help highlight specific weaknesses and test countermeasures such as differential privacy, gradient clipping, and client trust weighting. The modular design of our framework already supports client-specific behavior, making it feasible to simulate adversarial and honest clients concurrently.

While the anomaly detection method uses binary classification with neural networks in the standard way, the originality here is in embedding that anomaly detection within a federated learning framework with a quantum-inspired architectural viewpoint. Rather than looking only at accuracy or privacy, this research undertakes a multi-perspective approach, including the privacy-preserving distributed learning, the domain-specific hyperparameter tuning, and the generalization of performance over clients. Furthermore, the framework is evaluated through a broad array of metrics (including F1, AUC-ROC, precision, recall), tests for client robustness, and through the application of a quantum-inspired conceptual design philosophy that accentuates entangled feature interactions plus high-dimensional data encoding.

6 Conclusion

This project was realized as a working example for the design, implementation, and evaluation of an FL system for anomaly detection under client-server architecture. The TensorFlow framework empowers various distributed clients to collaboratively train one shared global model without exchanging raw data. Each client locally trains on a private data subset and only transmits the learning model parameters to the central server. The server then uses Federated Averaging (FedAvg) on the received parameters to update the global model. This approach guarantees data privacy,

local autonomy, and regulation-conformant operation, which are all necessary in application areas such as healthcare, finance, and IoT. The model achieved a high level of accuracy, almost matching that of a centralized system, while still retaining total decentralization of data. Through PCA-based feature visualization, significant class separation occurred in the reduced-dimension space representation, corroborating the hypothesis that the extracted features indeed could distinguish between normal and anomalous samples. Comparative evaluations confirmed that the federated model yields a high level of accuracy and generalization, while local models fare worse on either of these terms, thereby implying that collaborative training yields stronger and scalable results. During the developmental phase of the architecture, several real-world challenges were addressed-perfecting the communication overhead, synchronization between clients and server, and imbalanced-class issue. It was made modular and extensible to allow configuring additional clients or rounds in the future. Even though the current solution relies on classical neural networks, it is still good enough for an upgrade enabling quantum integration later; this duality holds great potential for tackling intricate, high-dimensional anomaly-detection problems. Future work may investigate asynchronous federated learning, differential privacy, and secure aggregation to improve robustness and scalability. In the concluding remarks, the project provides a very good example of how federated learning can practically be deployed for privacy. It has practically maintained model performance. This foundation is very crucial if secure decentralized solutions for anomaly detection are to be deployed in real environments where the sensitivity of the data and confidentiality rank high.

7 Future scope

The future work can extend the present system in numerous meaningful ways. First, asynchronous client updates and quorum-based or partial aggregation can be enabled for better robustness against client dropouts and network delays. Second, more security can be imposed over the sensitive information of clients through means such as differential privacy, homomorphic encryption, or secure

multiparty computation. Third, further efficiency can be gained in anomaly detection via integration with hybrid quantum-classical models or variational quantum circuits, thus benefiting from quantum-inspired representations. Fourth, incorporating interpretability measures like SHAP or LIME would allow domain experts to better understand how anomalies are decided. Finally, implementing the framework in real federated environments with very non-IID data and unstable client availability will aid in proving its scalability, resilience, and viability in multiple domains like health care, finance, and cybersecurity (Zhuang et al., 2020). The current setup uses a synthetic, balanced anomaly detection dataset crafted so that evaluation can be controlled in reproducibility in a non-IID federated setting. Although this helps in filtering out and analyzing aspects of the federated architecture and hyperparameter sensitivity, it does not fully encapsulate a real-world anomaly detection scenario. This will be a part of future work wherein we intend to extend the evaluation to real-world network intrusion datasets such as NSL-KDD, CIC-IDS2017, and UNSW-NB15 with highly class-imbalanced, heterogeneous feature, and realistically distributed attack types. Such datasets will allow a validation of robustness, generalizability, and practical utility of the model in actual anomaly detection scenarios, mainly in non-IID and adversarial settings. The current system architecture, being modular and dataset-agnostic, can easily ingest these datasets with minimal changes. A shallow feedforward neural network has been selected in the current approach due to its simplicity, communication overhead minimization, and ease of deployment in federated settings. Although good enough for preliminary benchmarking and testing on balanced datasets, such a neural architecture may not be able to capture the complex or temporal inter-dependencies that are often involved in real-world anomaly detections. Future extensions to this work will involve benchmarking the performance of more expressive models such as Long Short-Term Memory (LSTM) architectures in temporal or sequential anomalies and autoencoders in unsupervised anomaly detection. These models can potentially learn latent representations and reconstruct patterns, making them pertinent candidates for detecting subtle irregularities. The current system configuration in a modular fashion makes it trivial to incorporate the architectures mentioned above in both centralized and federated configurations while abiding by the privacy constraints.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

References

- Alsulaimawi, Z. (2024). Federated learning with anomaly detection via gradient and reconstruction analysis. *Arxiv [Preprint]*. doi: 10.48550/arxiv.2403.10000
- Anwar, A., Moser, B. B., Herurkar, D., Raue, F., Hegiste, V., Legler, T., et al. (2024). FedAD-bench: a unified benchmark for federated unsupervised anomaly detection in tabular data. *Arxiv [Preprint]*. doi: 10.48550/arxiv.2408.04442
- Bhatia, A. S. S., and Kais, S. (2023). Federated convolutional neural network: a new paradigm for collaborative quantum learning. *Quantum Sci. Technol.* 8:61. doi: 10.1088/2058-9565/acfc61
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., and Lloyd, S. (2017). Quantum machine learning. *Nature* 549, 195–202. doi: 10.1038/nature23474
- Blanchard, P., El Mhamdi, E. M., Guerraoui, R., and Stainer, J. (2017). Machine learning with adversaries: Byzantine tolerant gradient descent. *Adv. Neural Inform. Process. Syst.* 30, 118–128.
- Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., et al. (2019). Towards federated learning at scale: system design. *Proceedings of the 2nd SysML conference*.
- Brisimi, T. S., Chen, R., Mela, T., Olshevsky, A., Paschalidis, I. C., and Shi, W. (2018). Federated learning of predictive models from federated electronic health records. *Int. J. Med. Inform.* 112, 59–67. doi: 10.1016/j.ijmedinf.2018.01.007
- Burke, P. J., and Stets, J. E. (2023). *Future Research* (pp. 236–C12P102). Oxford: Oxford University Press eBooks.
- Chen, M., Zhang, Y., Gündüz, D., Saad, W., Yin, C., and Debbah, M. (2021). Robust federated learning against adversarial attacks. *IEEE Trans. Commun.* 69, 6116–6130. doi: 10.1109/TCOMM.2021.3078067
- Dehlaghi-Ghadim, A., Marković, T., Leon, M., Söderman, D., and Strandberg, P. E. (2023). Federated learning for network anomaly detection in a distributed industrial environment. 2023 International Conference on Machine Learning and Applications (ICMLA) 218–225.
- Fan, Z., Wang, Y., Liu, J., Yu, S., and Li, L. (2023). Efficient federated learning for anomaly detection in edge computing systems. *J. Syst. Archit.* 135:102693. doi: 10.1016/j.sysarc.2023.102693

Author contributions

DG: Supervision, Writing – review & editing. VR: Writing – original draft, Resources, Methodology. SM: Formal analysis, Software, Writing – review & editing. JJ: Writing – original draft. DP: Writing – original draft. BD: Conceptualization, Writing – review & editing. FM: Data Curation, Writing – review & editing.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This work is funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under the support UID/05105: REMIT – Investigação em Economia, Gestão e Tecnologias da Informação.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Correction note

This article has been corrected with minor changes. These changes do not impact the scientific content of the article.

Generative AI statement

The authors declare that no Gen AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Geyer, R. C., Klein, T., and Nabi, M. (2018). Differentially private federated learning: a client level perspective. *Arxiv [Preprint]*. doi: 10.48550/arXiv.1712.07557
- Ghosh, A., Chung, J., Yin, D., and Ramage, D. (2020). Efficient federated learning for mobile edge computing. *Arxiv [Preprint]*. doi: 10.48550/arXiv.2409.02548
- Gosselin, R., Vieu, L., Loukil, F., and Benoit, A. (2022). Privacy and security in federated learning: a survey. *Appl. Sci.* 12:9901. doi: 10.3390/app12199901
- Hardy, S., Henecka, W., Ivey-Law, H., Nock, R., Patrini, G., Smith, G., et al. (2017). Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption. *Arxiv [Preprint]*. doi: 10.48550/arXiv.1711.10677
- Hasan, J. H. (2023). Security and privacy issues of federated learning. *Arxiv [Preprint]*. doi: 10.48550/arxiv.2307.12181
- Havlek, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., et al. (2019). Supervised learning with quantum enhanced feature spaces. *Nature* 567, 209–212. doi: 10.1038/s41586-019-0980-2
- Imteaj, A., Thakker, U., Wang, S., Li, J., and Amini, M. H. (2022). A survey on federated learning for resource-constrained IoT devices. *IEEE Internet Things J.* 9, 1–24. doi: 10.1109/jiot.2021.3095077
- Kairouz, P., and McMahan, H. B. (2021). The challenges and opportunities of federated learning. *Nat. Mach. Intell.* 3, 486–497. doi: 10.1038/s42256-021-00303-7
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., et al. (2021). Advances and open problems in federated learning. *Foundat. Trends Mach. Learn.* 14, 1–210. doi: 10.1561/22000000083
- Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., and Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. *Arxiv [Preprint]*. doi: 10.48550/arXiv.1610.05492
- Li, Q., He, B., and Song, D. (2021). Practical one-shot federated learning for cross-silo settings. *IEEE Trans Big Data* 7, 858–870. doi: 10.1109/TBDATA.2020.2996601
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020). Federated learning: challenges, methods, and future directions. *IEEE Signal Process. Mag.* 37, 50–60. doi: 10.1109/MSP.2020.2975749
- Li, S., Wang, M., Zhang, J., and Tian, Y. (2023). Quantum-enhanced federated learning: a review and perspectives. *IEEE Trans. Quantum Eng.* 4, 1–13. doi: 10.1109/TQE.2023.324874
- Liberti, F., Berardi, D., and Martini, B. (2024). Federated learning in dynamic and heterogeneous environments: advantages, performances, and privacy problems. *Appl. Sci.* 14:8490. doi: 10.3390/app14188490
- Liu, Y., Kang, Y., Zhang, Y., and Yang, Q. (2020). A communication efficient vertical federated learning framework. *Arxiv [Preprint]*. doi: 10.48550/arXiv.2004.11013
- Lu, Y., Huang, Q., Wang, Y., and Zhang, L. (2022). Blockchain and federated learning for privacy-preserving data analysis: a survey. *Futur. Gener. Comput. Syst.* 130, 44–59. doi: 10.1016/j.future.2022.01.013
- Lundberg, S. M., and Lee, S. I. (2017). A unified approach to interpreting model predictions. *Adv. Neural Inform. Process. Syst.* 30, 4768–4777.
- Ma, C., Yu, F., Chen, T., and Liu, Y. (2022). On the generalization of federated learning under client drift. *IEEE Trans. Neural Netw. Learn. Syst.* 33, 2479–2491. doi: 10.1109/TNNLS.2021.3088254
- Marfo, W., Tosh, D. K., and Moore, S. V. (2025). Federated learning for efficient condition monitoring and anomaly detection in industrial cyber-physical systems. *Arxiv [Preprint]*. doi: 10.48550/arXiv.2501.16666
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th international conference on artificial intelligence and statistics (AISTATS)*, PMLR, 54, 1273–1282.
- Mohassel, P., and Zhang, Y. (2017). SecureML: A system for scalable privacy-preserving machine learning. *2017 IEEE Symposium on Security and Privacy (SP)*, 19–38.
- Mohri, M., Sivek, G., and Suresh, A. T. (2019). Agnostic federated learning. *Proceedings of the 36th International Conference on Machine Learning (ICML)*, PMLR 97, 4615–4625
- Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., and Niyato, D. (2022). Federated learning for smart healthcare: a survey. *ACM Comput. Surv.* 55, 1–37. doi: 10.1145/3474124
- Park, D., Kim, D., and Rhee, W. (2022). Quantum federated learning for healthcare data. *IEEE Access* 10, 40195–40207. doi: 10.1109/ACCESS.2022.3167101
- Purohit, S., and Manimaran, G. (2024). Federated learning based anomaly detection for distributed energy resource communication. *Seattle, WA, USA. 1–5*. doi: 10.1109/PESGM51994.2024.10688634
- Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). "Why should I trust you?": explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144.
- Santin, G., Skarbovsky, I., Fournier, F., and Lepri, B. (2022). A framework for verifiable and auditable collaborative anomaly detection. *IEEE Access* 10, 82896–82909. doi: 10.1109/access.2022.3196391
- Sattler, F., Wiedemann, S., Müller, K.-R., and Samek, W. (2019). Robust and communication-efficient federated learning from non-IID data. *IEEE Trans. Neural Netw. Learn. Syst.* 31, 3400–3413. doi: 10.1109/TNNLS.2019.2944486
- Schuld, M., and Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. *Phys. Rev. Lett.* 122:504. doi: 10.1103/PhysRevLett.122.040504
- Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., and Bakas, S. (2019). Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation. *Granada, Spain: BrainLes Workshop – MICCAI*, 92–104.
- Wang, S., Tuor, T., Salonidis, T., Leung, K. K., Makaya, C., He, T., et al. (2020). Adaptive federated learning in resource-constrained edge computing systems. *IEEE J. Sel. Areas Commun.* 38, 1205–1221. doi: 10.1109/JSAAC.2020.2993425
- Xu, J., Wang, B., Li, Y., and Yang, Q. (2021). FedHarmony: communication-efficient federated learning via local update skipping and global parameter interpolation. *Proceedings of the 30th international joint conference on artificial intelligence (IJCAI)*, 3933–3939
- Yang, Q., Liu, Y., Chen, T., and Tong, Y. (2019). Federated machine learning: concept and applications. *ACM Trans. Intellig. Syst. Technol.* 10, 1–19. doi: 10.1145/3298981
- Yang, J., Zhang, Y., Chen, T., and Yang, Q. (2020). Federated edge learning for smart cities. *IEEE Netw.* 34, 34–39. doi: 10.1109/MNET.011.2000102
- Yin, D., Chen, Y., Ramage, D., and Cheng, Y. (2021). Gradient compression and differential privacy in federated learning. *Proc. AAAI Conf. Art. Intellig.* 35, 9469–9477.
- Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. (2018). Federated learning with non-IID data. *Arxiv [Preprint]*. doi: 10.48550/arXiv.1806.00582
- Zhao, T., Liu, Y., Zhu, Y., Xu, Y., and Zhang, X. (2021). Federated semi-supervised anomaly detection with communication-efficient training. *Pattern Recogn. Lett.* 145, 68–75. doi: 10.1016/j.patrec.2020.12.024
- Zhuang, Y., Wu, Y., Zhou, W., Zhang, Z., and Li, Y. (2020). A comprehensive survey on federated learning: concepts, applications, challenges, and future directions. *IEEE Access* 8, 53875–53995. doi: 10.1109/ACCESS.2020.2987349