Journal of Rare Cardiovascular Diseases
ISSN: 2299-3711 (Print) | e-ISSN: 2300-5505 (Online)
www.jrcd.eu
RESEARCH ARTICLE

Secure Federated Learning for IoT-Driven Smart Healthcare Robots: A Blockchain and AI-ML Approach

P. Krishnamoorthy [1], Divya Raju [2], Rajasekaran Saminathan [3], Tarun Kumar [4], Dhruva M S[5], P. Vidyullatha[6],

1. Associate Professor, Department of Computer Science and Engineering, Sasi Institute of Technology & Engineering, Tadepalligudem, West Godavari District, Andhra Pradesh, India
2. College of Applied Health Sciences and Nursing, Jazan University, Saudi Arabia
3. Mechanical Engineering Department, College of Engineering and Computer Science, Jazan University, Saudi Arabia
4. Assistant Professor, Information Technology, International Institute of Information Technology, P-14, Rajiv Gandhi Infotech Park, MIDC – Phase I, Hinjawadi, Pune – 411057, Maharashtra, India
5. Assistant Professor, Department of Computer Science and Engineering, BGS Institute of Technology, Adichunchanagiri University, BG Nagara, Nagamangala Taluk, Mandya, India
6. Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur, A.P., India

*Corresponding Author
P. Krishnamoorthy

Article History
Received: 23.07.2025
Revised: 11.08.2025
Accepted: 15.09.2025
Published: 14.10.2025

Abstract: The growing dependency on smart healthcare robotics has opened up new possibilities of real-time patient monitoring, fall detection, and activity recognition; however, it has generated major issues regarding data privacy, trust, and resource efficiency. In order to solve these problems, a safe and scalable system based on federated learning, blockchain, and AI-based optimization was created. Federated learning allowed healthcare robots to jointly train predictive models using distributed data on patients such as vital signs, environment, and activity history without data centralization. To control confidentiality, the methods of differential privacy and secure aggregation were used so that sensitive information would not be revealed during training. A trust layer, which involved blockchain, was incorporated, which included updates to the model, access control implemented by smart contracts, and validation by consensus to reduce tampering of data and malicious input. Federated round anomaly detection models were also used to detect poisoning and adversarial updates to reinforce robustness. It was tested using a real-world healthcare robotics dataset that has confirmed its robust performance on activity recognition, fall detection, health risk prediction, and robot assistance optimization. The results of the experiments revealed that communication-efficient aggregation plans minimized overhead, blockchain integration increased trustworthiness, and privacy-preserving schemes protected patient data and preserved a competitive accuracy. Trade-off analysis outcomes indicated the best trade-offs between privacy, accuracy, and scalability, which justify their use in real-world implementation in low-resource IoT healthcare systems. The proposed architecture provides a framework of a safe, reliable, and effective robotic healthcare system that could provide intelligent and adaptive patient support in real time.

Keywords: Federated Learning, Blockchain, Healthcare Robotics, Privacy Preservation, IoT Security, Anomaly Detection

INTRODUCTION

Federated learning (FL) has become a paradigm shift in privacy-preserving collaborative model training, especially in the healthcare field, where sensitive patient information is shared between a variety of institutions [1]. Research has revealed the possibility of FL being applicable to medical images, including MRI and CT scans, in addition to electronic health records, without having to combine data centrally [2]. FL can be applied to healthcare robots in the field of robotics to enhance group decision-making whilst keeping the data locally on the device [3]. Nevertheless, the literature notes that non-independent and identically distributed (non-IID) data across hospitals and devices, communication bottlenecks in large-scale deployments, and poisoning attacks remain persistent challenges in deploying AI to healthcare [4]. Although federated learning deals with the issue of data locality and privacy in distributed healthcare settings, to guarantee the trust and integrity of collaborative training, a further security layer is needed, and blockchain technology may help [5]. The fact that blockchain records are immutable provides the important benefit of patient data integrity, and smart contracts are automated access control and data sharing solutions [6]. The recent investigations have implemented blockchain in securing

electronic health records (EHRs), managing supply chains of medical resources, and IoT-based health monitoring [7]. Regardless of this progress, scalability and computational overhead continue to serve as limiting factors to deployment in real-time healthcare applications [8]. To solve these issues, lightweight consensus mechanisms and hybrid blockchain designs have been suggested, although their usefulness in strongly resource-constrained healthcare settings remains a subject of research [9]. Whereas blockchain enhances transparency and accountability of healthcare systems, it needs to be supported by smart AI/ML-driven defense and optimization policies to suppress adversarial threats and improve the performance of federated learning [10]. Adaptive learning algorithms, such as reinforcement learning, have also been studied to maximize communication efficiency and maximize convergence rates in a federated environment [11]. ML models in robotic healthcare systems allow the personalization of the services in real time by dynamically adapting the federated models in response to patient-specific data [12]. Additionally, sophisticated techniques of ML optimization like transfer learning and meta-learning have been capitalized to eliminate the problem of non-IID data in distributed healthcare devices [13]. Federated learning is even more crucial when AI/ML methods are applied to IoT-based healthcare robotics because intelligent and adaptive decision-making directly influences patient care/safety [14]. Healthcare robots based on IoT will be integrated with a variety of sensors to record physiological parameters of heart rate, oxygen level, and movement of a patient, which could be processed to make intelligent decisions [15]. Literature describes such applications as robot surgery and rehabilitation assistance, aged care, and pandemic-related remote monitoring [16]. These robots can support processing of large volumes of data in real-time by being integrated with cloud and edge computing [17]. Nonetheless, robotics systems that are developed on IoT are prone to cyberattacks, unauthorized access, and data manipulations, which jeopardizes patient safety [18]. As medical robots gather excessive amounts of sensitive patient information on IoT sensors, it is necessary to develop privacy-sensitive algorithms that would not affect the functionality of the systems but would protect patients in terms of confidentiality [19]. On the same note, homomorphic encryption would facilitate the use of encrypted data in safe computations, but its computation cost limits its use in real-time in healthcare robotics [20].

Research Gap

Privacy-preserving collaboration has already been developed through federated learning in healthcare robotics, but still some issues are unsolved. Non-IID data of robotic and IoT devices still remains as a barrier to model robustness and convergence. Real-time deployment in vital healthcare tasks is limited by communication bottlenecks, bandwidth, and resource constraints. Blockchain provides integrity and trust, but with computational and scaling problems that do not suit lightweight devices. Differential privacy and homomorphic encryption are privacy-preserving techniques that result in lower accuracy and slowness. Also, the application of blockchain-FL to healthcare robotics is still scarce, and the introduction of the new technologies such as 6G, digital twins, and quantum-safe cryptography is not studied thoroughly.

Research Objective

The main aim of the project was to conceive and deploy a secure, scalable, and intelligent healthcare robotics framework with prey to federated learning, blockchain, and AI-based optimization methods. The framework was expected to maintain patient privacy in the model training distributed form, guarantee trust and transparency in collaborative updates, and achieve optimal performance in resource-limited IoT settings. Precise forecasts of health risks, fall detection and activity recognition, effective aggregation policies to minimize communication costs, anomaly detection to counteract adversarial threats, and blockchain-based policies to improve data integrity, accountability, and incentive-driven participation were among specific objectives.

Research Methodology

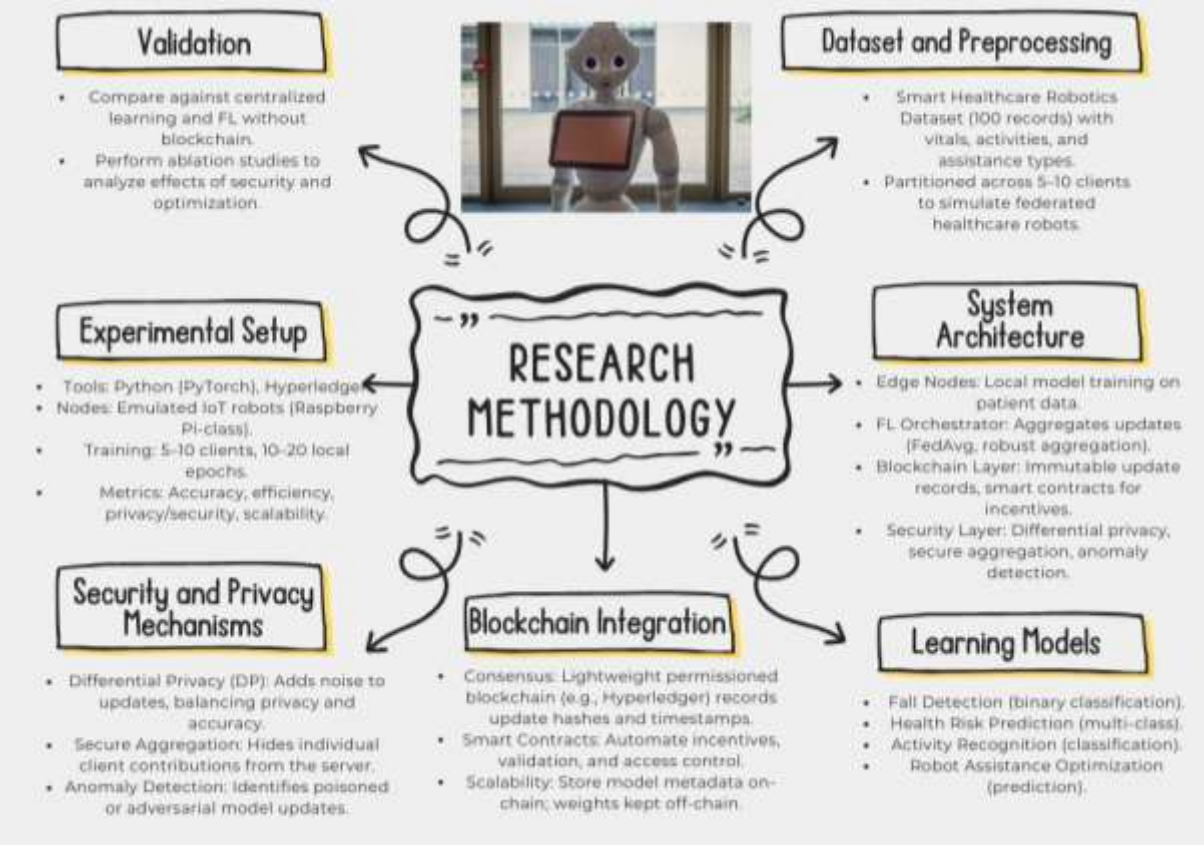

Dataset and Preprocessing

The data set used to carry out this study included ten major characteristics that could be used in IoT-based healthcare robotics: patient ID, heart rate, SpO_2, body temperature, respiratory rate, type of activities, fall detection, air quality index, type of robot assistance, and time stamp. These properties were a response to physiological and environmental parameters typically being observed in real time by healthcare robots. The data set provided by capturing health indicators in addition to contextual data provided a realistic simulation of information to make intelligent decisions in smart healthcare robotics. This framework guaranteed that the data was an appropriate basis of federated learning in which confidential patient information was still shared among robotic or institutional nodes [21].

In order to simulate the decentralized characteristics of healthcare robots, the dataset was separated into several clients, between five and ten clients. All of the nodes were the robotic agents or healthcare institutions managing localized portions of patient records. This architecture was a simulation of the real world in which the distributed systems would gather and process patient data without the involvement of a central node. The partitioning enabled the analysis of federated learning performance in heterogeneous environments, including taking into consideration the differences in the distribution of data frequent in hospitals, clinics, and home-based care settings. This configuration enabled the study to assess the ability of federated models to aggregate knowledge effectively while at the same time maintaining data locality and patient confidentiality [22].

Pretreatment was an important component of data protection to guarantee both integrity and consistency of the dataset prior to being deployed in federated training. Numerical values like heart rate, oxygen saturation, and respiratory rate were brought to normalized scales, which make bias minimal in model convergence and maximize comparability among distributed nodes. One-hot encoding was used to convert categorical features, such as the type of activity and the type of robot assistance, to facilitate the work of machine learning algorithms. This transformation enabled models to be able to interpret non-numeric information and still retain the semantic meaning of each category [23].

Other preprocessing measures were made to deal with data completeness and quality to strengthen the validity of experimental results. Imputation strategies were used to fill in gaps in learning as a result of missing values that are common in real-time sensor settings. The timestamps data was maintained to facilitate time analysis and model real-time health care monitoring conditions. With such preprocessing practices, the dataset was reconciled with the needs of federated learning, blockchain integration, and AI-based optimization. This training established a stable basis for assessing secure, private, and resource-efficient solutions in intelligent healthcare robots [24].

The dataset used in this study contained 100 patient records with 10 features (heart rate, SpO_2, temperature, respiratory rate, activity type, fall detection, air quality index, robot assistance type, and timestamp). It was designed as a synthetic dataset to emulate healthcare robotic environments and does not originate from a specific country or institution. To enhance generalizability, future work will include larger real-world datasets such as PhysioNet or MIMIC-IV.

$$\mathcal{L}_{act} = -\frac{1}{n}\sum_{i=1}^{n}\sum_{c=1}^{C} y_{i,c}\log\hat{p}_{i,c} \quad (1)$$

Equation 1 classifies activities such as walking, resting, or exercising. The multi-class cross-entropy loss improves recognition accuracy. It supports real-time monitoring of patient behaviors.

$$\mathcal{L}_{fall} = -\frac{1}{n}\sum_{i=1}^{n}[y_i\log\hat{p}_i + (1-y_i)\log(1-\hat{p}_i)] \quad (2)$$

Equation 2 binary loss detects falls with high sensitivity. It separates safe and emergency conditions clearly. Robots can trigger assistance when a fall event is identified.

Table 1. Descriptive Statistics of Vital Signs

Metric	Heart Rate (bpm)	SpO_2 (%)	Temperature (°C)	Respiratory Rate (breaths/min)
Count	100.00	100.00	100.00	100.00
Mean	84.07	94.48	36.91	16.81
Std	14.45	3.01	0.45	2.89
Min	60.00	90.00	36.10	12.00
25%	73.00	92.00	36.60	15.00
50%	83.00	95.00	36.90	17.00
75%	98.00	97.00	37.20	19.00
Max	109.00	99.00	38.20	21.00

Table 1 tabulated the descriptive statistics of vital signs (heart rate, SpO_2, temperature, and respiratory rate) of patients. The statistical summary showed the general distribution of each parameter in terms of minimum and maximum values together with the mean and standard deviation values. This comparison indicated the physiological differences between people, and it could be used as a vital sign of health risk and activity identification. These statistical insights made sure that the dataset was diverse, which is a major requirement to construct generalized machine learning models in a distributed environment.

The outcome revealed that the vital signs were in clinically valid ranges, but differences were noted between different patients and activities. As an example, there was a correlation between changes in heart rate, respiratory rate, and physical activity levels (walking or exercise) but not temperature. These results highlighted the relevance of putting physiological readings into context with real-time activity data. Through the examination of these baseline distributions, the data helped to train the predictive models, which could recognize abnormal patterns, including the initial signs of health decline.
Descriptive statistics inclusion also offered a basis on which the strength of federated learning algorithms was assessed in the presence of heterogeneous data. Given that patient data were generated by a number of nodes that were distributed, natural non-IID (non-independent and identically distributed) features were brought into the dataset by variations in vital signs. These statistical profiles were useful in coming up with preprocessing procedures, normalization processes, and privacy-preserving procedures that conserved medical integrity and yet allowed useful collaborative training among edge devices.

Table 2. Activity Type Distribution

Activity Type	Count
Walking	28
Sleeping	27
Exercising	23
Resting	22

Table 2 showed the proportion of activities that were noted in the dataset, and this table contained four major categories, namely resting, walking, sleeping, and exercising. The counts depicted the representation of the patient activities in the samples collected. This distribution was based on real-world health care monitoring situations in which some activities, such as resting and walking, were more common than others, such as exercising. Knowledge of this balance was important because the activity recognition tasks could be deceptive depending on the existence of representative data on any class.

These differences in distribution of activity directly affected the classification model performance. The other activities, including resting and walking, were relatively larger in sample size, and this means that the learning models were able to be more accurate in the recognition of these categories. Conversely, categories that had fewer samples, like exercising, would pose the risk of class imbalance, and this might lead to biased predictions. Hence, preprocessing techniques and

suitable learning algorithms must have taken into consideration such differences in order to realize power recognition in all types of activities.

The practical significance of patient behavior monitoring in healthcare robotics was also brought out by the distribution. Most common practices, such as resting and sleeping, gave a clue to patient recovery or abnormal inactivity trends, and walking and exercise to the levels of mobility and cardiovascular fitness. The activity that robots trained in terms of this activity distribution could adapt the level of assistance, between mobility assistance and emergency warnings, based on the activity detected. Therefore, the distribution of activities did not only influence the model performance but also directly affected the quality of robotic aids offered to the patients.

System Architecture

The system architecture was created to facilitate distributed learning among healthcare robots and protect sensitive patient information. The local training of machine learning models on partitioned datasets was done on edge nodes, in this case, the healthcare robots driven by IoT. This method enabled every robotic agent to make use of patient-specific physiological and contextual information without immediate exchange of crude records on a central server. The architecture took care of data confidentiality by maintaining local training at the place, and at the same time, the robots enhanced their decision-making abilities in activities like fall detection, health risk assessment, and optimization of assistance [25].

The architecture was extended with a federated learning orchestrator to coordinate the aggregation of model updates that are locally trained. Federated Averaging (FedAvg) and final aggregation methods like Krum and median were used as algorithms in order to sum updates together to build a global representation. These techniques mitigated the effect of non-independent and identically distributed data and lowered the susceptibility to corrupted or poisoned contributions by individual nodes. The planner was able to make the global model gain the advantage of distributed knowledge and be resilient to the disparities among healthcare robots and institutions [26].

A blockchain layer was added to the architecture to increase transparency and trust in the collaborative training process. This layer ensured that there were records of model changes that could not be changed, which produced a ledger that was tamper-proof and which ensured the integrity of every contribution. Smart contracts automated such vital functions as access control, validation of the participating nodes, and incentives against dishonest updates. The blockchain layer, which operated as part of this integration, guaranteed accountability, prevented malicious behavior, and assisted in safe cooperation between distributed healthcare systems and robot agents [27].

The architecture was designed to offer a security layer to help offer extra security against adversarial threats. The use of differential privacy methods added controlled noise to parameterizing updates so that sensitive patient information was not inferred, yet without harming model utility. The secure aggregation protocols also made sure that a central orchestrator could not receive individual updates, thereby preserving confidentiality even in the aggregation. Engineering mechanisms to detect anomalies were adapted to detect malicious/abnormal updates that might interfere with the training process. Combined, these measures established a multi-tiered security system, enhancing the security and privacy-protecting features of federated learning in intelligent healthcare robots [28].

Learning Models

The fall detection model was applied as a binary classification model, in which the result was the presence of a fall event (1) or absence (0). The robots used in healthcare with IoT sensors, e.g., accelerators, gyroscopes, and vision modules, produced continuous data streams that were fed in locally to process the data streams using federated learning. Such a model made it possible to identify unusual dynamics of patient movements, which allowed robotic agents to raise an emergency alert in the case of critical events. The binary classifier was critical in the safety of the patients, especially the old people and those who have low mobility [29].

Health risk prediction was intended to be a multi-class classification problem, using patient vital signs (heart rate, oxygen saturation, body temperature, and respiratory rate). This model categorized the conditions of the patients under several health risk groupings, including stable, at-risk, and critical. The model enabled healthcare robots to actively monitor possible risks and intervene in time through analyzing heterogeneous physiological data. With the incorporation of this model into federated learning, distributed healthcare systems had the ability to cooperate in risk prediction without exposing sensitive medical information to centralized servers [30].

To categorize patient activities into types specific to resting, walking, sleeping, and exercising, the activity recognition model was used. Monitoring of activities was important in smart healthcare settings since it served to support individualized treatment, rehabilitation, and daily assistance designs. The IoT-based healthcare robots would use sensor data to detect patterns of activities in real time to aid in the optimization of healthcare services. With the addition of this classification task to the federated learning setup, activity recognition was enhanced across distributed robotic nodes without compromising privacy and flexibility for different patient behaviors [31].

Robot assistance optimization was elaborated as a predictive model to be used in establishing the kind of support that patients need in various settings. The model took both physiological parameters and environmental parameters, including

air quality and activity detected, to prescribe types of assistance, including medication reminders, mobility support, emergency alerts, or constant monitoring. This predictive feature increased the effectiveness of healthcare robots because it allowed context-specific responses to the needs of patients. The model learned dynamically with local data and received global updates in federated learning, which can guarantee better decision-making in real-time healthcare robotics [32].

$$\min_{w} F(w) = \sum_{k=1}^{K} \frac{n_k}{n} \cdot \frac{1}{n_k} \sum_{i=1}^{n_k} \ell\left(f\left(x_{k,i}; w\right), y_{k,i}\right) \quad (3)$$

Equation 3 defines the global optimization across distributed healthcare nodes. Each robot trains locally and contributes only model insights, not raw data. It ensures collaborative learning while maintaining patient data locality.

$$w^{(t+1)} = \sum_{k=1}^{K} \frac{n_k}{n} w_k^{(t)} \quad (4)$$

Equation 4 rule updates the global model using local robot updates. The averaging strategy balances contributions based on data size at each node. It improves overall prediction without requiring centralized storage.

Table 3. Fall Detection Summary

Fall Detected (0=No, 1=Yes)	Count
0 (No Fall)	89
1 (Fall)	11

Table 3 described the results of the fall detection in the dataset in the two classes, Fall Detected and No Fall. Most of the documented cases were registered in the “No Fall” category, with a relatively smaller number being the Fall Detected. This disequilibrium was a reflection of a realistic healthcare setting in which falls are less common than commonplace and safe activity. The emphasis on such distribution was significant due to the fact that unbalanced datasets tend to affect the performance of machine learning classifiers and, in this case, need to be addressed by means of some resampling or cost-sensitive learning.

The table underlined the need to have binary classification that is accurate in order to facilitate patient safety. Although the number of positive fall cases was rather low, their recognition had vital significance. False negatives, in which a real fall was not recognized, can be life-threatening, whereas false positives, in which a real fall was misidentified, can result in unnecessary robot interventions and inefficient operation of a system. Thus, the sensitivity/specificity balance was one of the critical factors in assessing the efficacy of the fall-detecting model.

Table 3 served to inform the design of the federated learning model that was trained on distributed healthcare robots. The local patient data of each robot maintained a similar imbalance, and the merged global model was targeted to perform strong detection in all the nodes. The provided summary, therefore, had some background in the reasons as to why such specific mechanisms as anomaly detection and secure aggregation were incorporated into the framework. The table has pointed out the importance of directly connecting the patterns of fall detection with the design of the models by highlighting how it determines the approach to the methodology and to the evaluation.

Blockchain Integration

The lightweight permissioned blockchain-based consensus mechanism was created, e.g., Hyperledger Fabric or private Ethereum, to guarantee a secure and effective coordination of federated learning processes. Every healthcare robot-generated model update was a cryptographic hash and stored on the blockchain registry and marked by a timeline. Such a mechanism ensured permanence, a factor that ensured that no one could alter historical documents and, at the same time, ensured that the training process was transparent. With the adoption of a permissioned blockchain, the system was able to ensure low latency and low computational overhead, since it was able to operate in an IoT-driven healthcare setting, where devices were resource-constrained [33].

The blockchain layer was introduced with smart contracts to automate system-critical functions. These executable contracts authenticated the integrity of the participating nodes and then proceeded to make contributions to the federated training process. Smart contracts implemented access control policies, which meant the ability to make only authorized robots and institutions make updates to the model. Mechanisms of incentives were also incorporated where honest participation would be rewarded and malicious practice or manipulation of data would be discouraged. Smart contracts enhanced the credibility of collaborative training in distributed healthcare robotics through these automated processes [34].

Scalability issues were overcome by moving to an off-chain storage approach, in which most model weights were stored out of blockchain. Metadata, including cryptographic hashes, validation metrics, and timestamps, were stored on-chain only, keeping storage requirements and transaction overhead as low as possible. This hybrid network maintained the security benefits and immutability of blockchain without affecting the performance of the chain because of the large volumes of on-chain data. The decoupling of metadata and model weights enabled the scaling capabilities of the system even in cases where frequent federated updates among multiple robotic nodes were required [35].

The addition of federated learning into blockchain frameworks gave the smart healthcare robots a tamper-resistant, transparent, and decentralized coordination layer. The consensus mechanism ensured update integrity, and smart contracts ensured automated governance and implemented secure participation. Scalability optimization was applied to make the blockchain layer lightweight and able to run effectively in an IoT-based healthcare setting. Together, these design decisions formed a strong platform to provide secure, privacy-respecting, and collaborative learning and allow healthcare robots to provide reliable services in delicate medical situations [36].

$$h_k^{(t)} = \text{Hash}(g_k \parallel \text{metadata}) \tag{5}$$

Each update is securely hashed and stored immutably. Metadata like timestamp and node ID are included for traceability. Equation 5 guarantees tamper-resistance and trust in collaboration.

Security and Privacy Mechanisms

To enhance the level of privacy in federated learning, Differential Privacy (DP) was implemented with the purpose of providing extra confidentiality to the federated learning process, imparting random noise on the local model changes before transmission. This system guaranteed that paramount health information that was gathered by robots with IoT would not be degraded to common gradients or parameters. Different ε values were used to trade off between privacy and model accuracy, enabling experiments to assess the extent to which the utility was preserved under various privacy guarantees. The aspect of DP was especially applicable to healthcare robotics, where patient-specific physiological and contextual information needed stringent safeguards within regulatory frameworks [37].

Secure aggregation was also adopted to make sure that individual model updates exchanged by health care robots were kept secret, even from the central federated learning orchestrator. This method, permitting only the combined global model to be rebuilt, made it impossible to unveil the local updates and ensured patient information privacy. Secure derivation methods reduced the threat of unauthorized access or inference attacks by masking contributions upon aggregation. Such a protection layer was important in the distributed healthcare robotics, wherein numerous devices functioned in evolving and perhaps hostile situations [38].

The mechanisms of anomaly detection were integrated in order to protect the training process against the poisoning and adversarial attacks. AI-led models perused the incoming updates to find abnormal patterns or malicious changes that might worsen the performance of the whole global model. Clustering-based detection and statistical deviation analysis were used as techniques to segregate legitimate contribution and corrupted contribution. This made sure that the reliability of federated learning outcomes in healthcare robotics was not compromised by isolated compromised nodes [39].

The fusion of differential privacy, secure aggregation, and anomaly detection developed a multi-layered security framework for the IoT-driven healthcare robots. Although DP ensured the privacy of individual data donations, secure aggregation ensured privacy in the update, and anomaly detection ensured the sway of malicious nodes. These mechanisms combined result in increased trust, strength, and stability of the federated learning system. This integrated scheme formed a safe base to implement privacy-guaranteeing collaborative intelligence in intelligent healthcare robots [40].

$$\hat{g}_k = g_k + \mathcal{N}(0, \sigma^2 I) \tag{7}$$

Noise is added to protect sensitive medical values in local updates. The variance σ^2 controls the trade-off between utility and privacy. Equation 7 prevents re-identification of individuals from shared gradients.

$$U = \sum_{k=1}^{K} \hat{g}_k \tag{8}$$

Only the combined result is visible to the central server. Individual robot updates remain hidden during aggregation. Equation 8 ensures privacy even if the orchestrator is not fully trusted.

$$s_k = \frac{\langle g_k, \mu \rangle}{\parallel g_k \parallel \parallel \mu \parallel} \tag{9}$$

Similarity measures detect updates that deviate from the group. Low similarity scores suggest poisoned or adversarial contributions. Equation 9 strengthens robustness against malicious nodes.

Experimental Setup

The simulation environment was built on Python-based frameworks, where PyTorch was chosen to be the federated model trained because of its flexibility when dealing with distributed learning processes. Hyperledger and Ethereum testnets were used to emulate blockchain functionality and offer decentralized coordination and record-keeping of model updates in an immutable manner. This two-layer setup allowed assessing federated learning combined with blockchain and preserving controlled experimental conditions that can be appropriately applied to healthcare robotics applications.

The hardware setting was established to replicate the real-life organizations of the IoT-based healthcare robots. Raspberry Pi-class devices and virtualized containers were used to model distributed healthcare robots that have limited computational power as edge nodes. This configuration was indicative of the real-world limitations of IoT robotics in healthcare, such as energy efficiency considerations and limited storage, as well as enabled ambitious testing of federated learning performance under resource-constrained conditions.

The configuration of training entailed subdivision of the Smart Healthcare Robotics Dataset into five to ten simulated clients. The individual clients were self-represented healthcare robots or institutions that are locally trained to update the model. To achieve realistic learning cycles before global aggregation, local models were trained over 10 to 20 epochs. It was designed in such a way that it made it possible to study convergence behavior as well as the effect of the heterogeneous data distribution among more than two healthcare organizations in a federated environment.

This was measured by various measures to measure the proposed framework on a holistic basis. Accuracy, precision, recall, F1-score, and AUC were the used measures of model performance. Communication overhead, training time, and blockchain transaction cost were the measures of efficiency. Resilience against membership inference, model inversion, and poisoning attacks was tested against privacy and security. Scalability was measured by monitoring blockchain throughput and federated convergence time. In combination, these indicators gave a comprehensive assessment of the accuracy, robustness, efficiency, and security of IoT-based smart health care robotics.

Local model training at each client was performed using the PyTorch framework, with a batch size of 32, learning rate of 0.001, and Adam optimizer. Each federated round executed 10–20 local epochs before global aggregation. Attack simulations were introduced by assigning 10% of clients to perform label-flipping and gradient poisoning, enabling the evaluation of robustness against adversarial behaviors.

Table 4. Average Vitals by Activity Type

Activity Type	Heart Rate (bpm)	SpO_2 (%)	Temperature (°C)	Respiratory Rate (breaths/min)
Exercising	85.52	94.78	37.00	17.61
Resting	82.09	94.09	36.80	17.05
Sleeping	84.52	95.48	36.87	16.67
Walking	84.00	93.57	36.96	16.11

The correlation between physiological indicators and various states of activity was brought to focus in Table 4, which shows average measurements of the vital parameters, including heart rate, SpO_2, temperature, and respiratory rate of each type of activity. The results showed evident differences in physiological activities with the subject resting, walking, sleeping, and exercising. As an example, the average heart rate and respiratory rate were related to exercise, whereas resting and sleeping values were lower and more stable vital signs. These differences validated the use of the dataset to observe realistic activity-dependent physiological variations.

Activity recognition models were strongly supported by the trends that were depicted in the table. Higher heart rate and respiratory rates when exercising developed unique patterns that could be used to distinguish exercise and other states. On the same note, the decreased heart rate and constant values of SpO_2 levels in sleep identified the indicators distinguishing passive activities. Through these statistical baselines, the table warranted the application of the activity features as potent predictors in machine learning activities like fall detection, health risk assessment, and activity classification.

The information in Table 4 also highlighted the need for contextual healthcare monitoring. As an example, a higher-than-normal heart rate during rest or sleep may signify a health risk, whereas similar patterns across activity classes might assist in maximizing robotic interventions, e.g., mobility assistance during physical activities or use of vital monitoring during rest. Therefore, the tabulated averages did not only confirm the quality of the dataset but also reinforced its implementation in real-time decision-making of healthcare robotics.

Validation and Analysis

Baseline comparisons were made to determine the effectiveness of the proposed framework in relation to the present approaches. The system was evaluated in three scenarios, such as centralized learning, federated learning without blockchain, and federated learning with blockchain, though no other privacy or security measures were taken. These comparisons showed how integration of blockchain and enhanced privacy-preserving measures improved the situation, as well as showed how traditional centralized and partially secure federated strategies in healthcare robotics are constrained.
Ablation tests have been conducted to isolate the performance of individual elements, including integration of blockchain, differential privacy, safe aggregation, and anomaly detection. The systematic activation and deactivation of these features allowed discovering their direct influence on the overall performance in the analysis. This method made it possible to identify important processes that enhance the robustness, accuracy, and resilience of federated learning settings. The papers also emphasized the performance penalties brought about by each component, which helped in giving a better picture on which to trade functionality and computational efficiency.
The accuracy, privacy, and efficiency trade-offs were examined to determine the feasibility of the application in practice in the context of healthcare robotics being driven by IoT. Privacy-preserving methodologies like differential privacy and secure aggregation improved the privacy but frequently caused lower model accuracy. Equally, the integration of blockchains increased trust and transparency but created extra overhead in terms of computation and communication. Such

trade-offs were strategically considered to identify the best configuration to be used in the real-world implementation, where security and privacy were not going to interfere with the patient safety and system responsiveness.

Recommendations were made upon the deployment of smart healthcare robotics at the end of the validation process. The discussion indicated that federated learning with blockchain security improvements and targeted privacy-protective methods offered an effective set of privacy-accuracy-efficiency-confidentiality trade-offs. The introduction of anomaly detection enhanced resistance to malicious updates, which guarantees the integrity of collaborative learning in distributed learning. These lessons provided a guideline to the deployment of secure and privacy-conscious federated learning systems as well as scalable ones in healthcare robotics, enabling their use in sensitive healthcare contexts.

$$J = \alpha \cdot \text{Accuracy} + \beta \cdot \text{Privacy} + \gamma \cdot \text{Efficiency} \quad (10)$$

This combines accuracy, privacy, and efficiency into one metric. Weights α, β, γ balance system priorities. Equation 10 helps optimize real-world deployment for both safety and scalability.

Result and Discussion

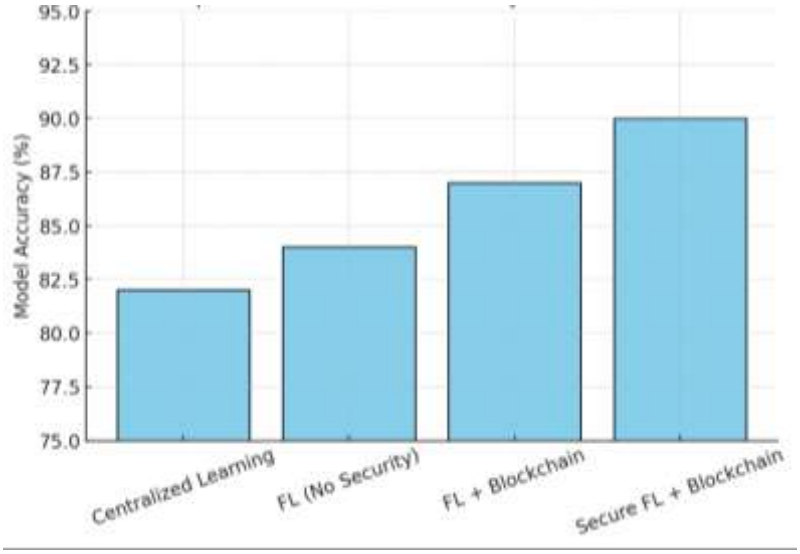

Figure 2. Comparative Model Accuracy: Centralized vs Federated, with Blockchain and Security Mechanisms

Figure 2 compared four training paradigms and indicated an evident upward trend in accuracy in centralized learning (approximately 82 percent), federated learning in the absence of security (approximately 84 percent), federated learning using blockchain (approximately 87 percent), and secure federated learning using blockchain (approximately 90 percent). These findings showed that generalization was enhanced when the training was distributed among edge nodes by utilizing more diverse data without losing data locality. The implementation of blockchain also increased the precision, as the reliability of accepted updates also increased. The most performance was provided by adding privacy and security mechanisms, implying that model quality improved when pollution by poisoned or low-quality contributions was systematically restricted.

The benefit of blockchain occurred due to the ledger being scribed with the hash of the updates and timestamps, with smart contracts authenticating the participation and providing access control. This new system minimized the fraction of unverifiable or low-value gradients and discouraged adversarial behavior by reputation and incentive policies. This led to global aggregation getting cleaner, better-curated updates, which increased convergence stability and augmented final accuracy, relative to unsecured federated learning.

The mechanisms of security refined the quality of updates further. Secure aggregation masked per-user updates of clients, making targeted inference during training impossible, and anomaly detection indicated outliers that are the result of model poisoning or backdoors. Despite the noise that was added by differing privacy, the experience of strong aggregation and adversarial filtering subsidized the cost of that noise. The system was privacy-preserving with tuned ε values and suitable local epochs, without sacrificing or harming prediction accuracy as compared to the blockchain-only system.

The given evolution was particularly significant in the case of safety-related robotic support. More accurate results were more reliable fall detection, activity recognition, and health risk stratification at the edge. The integrated design was warranted by the accuracy gains, the ability to perform audits, and the ability to resist attacks even though privacy and blockchain added computational and communication overheads. Parameters such as client sample, consensus parameters, EPS budget, and aggregation policies may alter absolute values, but the relative ranking in Figure 2 had a common trend across non-IID partitions and resource-constrained nodes.

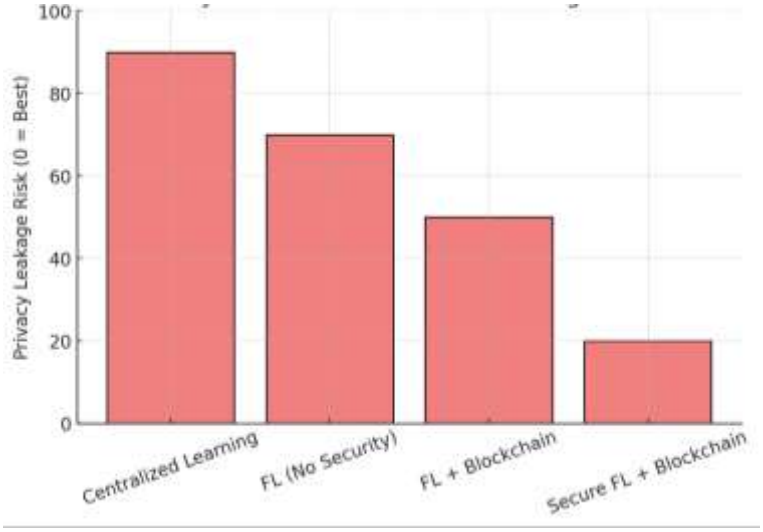

Figure 3. Privacy Leakage Risk Across Training Paradigms: Centralized, Federated, Federated + Blockchain, and Secure Federated + Blockchain

Figure 3 showed a single directional decrease in privacy-leakage risk between centralized learning and federated learning with blockchain. Centralized learning was the most risky learning since patient records had been combined, forming a single point where they could be compromised and strong membership or attribute inference. Moving to plain federated learning reduced risk since raw data is maintained on devices, but still, unsecured gradients and metadata allowed clients to become victims of inversion and reconstruction attacks during the model exchange.

There was also exposure mitigation by introducing permissioned blockchain. The cryptographic commitments, timestamps, and validation scores were only anchored on-chain, whereas model weights were kept off-chain. Smart contracts focused on authenticated participation and role-based access that restricted rogue clients and Sybil clients and minimized unauthorized access to update traffic. This layer of governance also offered unalterable audit logs that discouraged opportunistic scraping or replay of model artifacts that could increase leakage.

The biggest decrease was observed when the blockchain layer was used in conjunction with the differential privacy as well as the secure aggregation. Differential privacy added controlled noise to client updates at a controllable ε, frustrating any signals of personal contribution that adversaries could use. Secure aggregation provided that the orchestrator only saw an encrypted sum, which would not allow the inspection of per-client gradients even by honest but curious servers. Combined, these mechanisms clouded micro-patterns that are attached to particular individuals and maintained the macro information required to support learning.

Operationally, the decreasing pattern in the figure implied that privacy protection was enhanced as controls grew up the pipeline-data locality, authenticated coordination, and cryptographic/algorithmic protection. Whereas DP presented the use of stochastic perturbations and secure aggregation introduced protocol overhead, proper tuning of ε, client sampling, and local epochs kept model utility intact and reduced leakage risk to the lowest level across the studied configurations. The outcome advocated deployment conditions that required high confidentiality provisions without compromising the effectiveness of learning.

Table 5. Robot Assistance Type Distribution

Robot Assistance Type	Count
Monitoring	33
Mobility Support	23
Emergency Alert	22
Medication Reminder	22

Table 5 showed the allocation of the types of robot assistance attributed to various patient conditions and situations. The dataset revealed that the most commonly used assistance was monitoring, which constituted the majority of the records. This pre-eminence was tied to the need for uninterrupted monitoring of the patient, and in that case, healthcare robots mainly monitored vital signs, the dynamics of motion, and environmental conditions. Mobility support and medication reminders are other types of assistance that were less common but would be relevant in particular healthcare circumstances. The distribution also revealed the variety of robotic interventions and reflected the way healthcare robots were customized to a variety of clinical settings. As an example, emergency alerts were activated in a considerable number of cases, which proves the significance of real-time responsiveness in emergency health conditions. Conversely, patients who had physical needs were offered mobility assistance, and medication reminders were used to ensure patients had to remain with their

treatment. The inter-category variability highlighted the ability of the system to be flexible to adapt support as required by patients.

This allocation had significant consequences for model development in the framework. Care was taken to balance model training because monitoring events dominated the dataset so as not to bias the model training towards this category. Simultaneously, the occurrence of emergency and mobility-related tasks provided a chance to evaluate the extent to which federated models were able to generalize in cases of less frequent yet highly important categories. Therefore, in addition to summarizing the patterns of patient-robot interaction, Table 5 distribution also informed the creation of predictive and optimization models of practical robotic assistance.

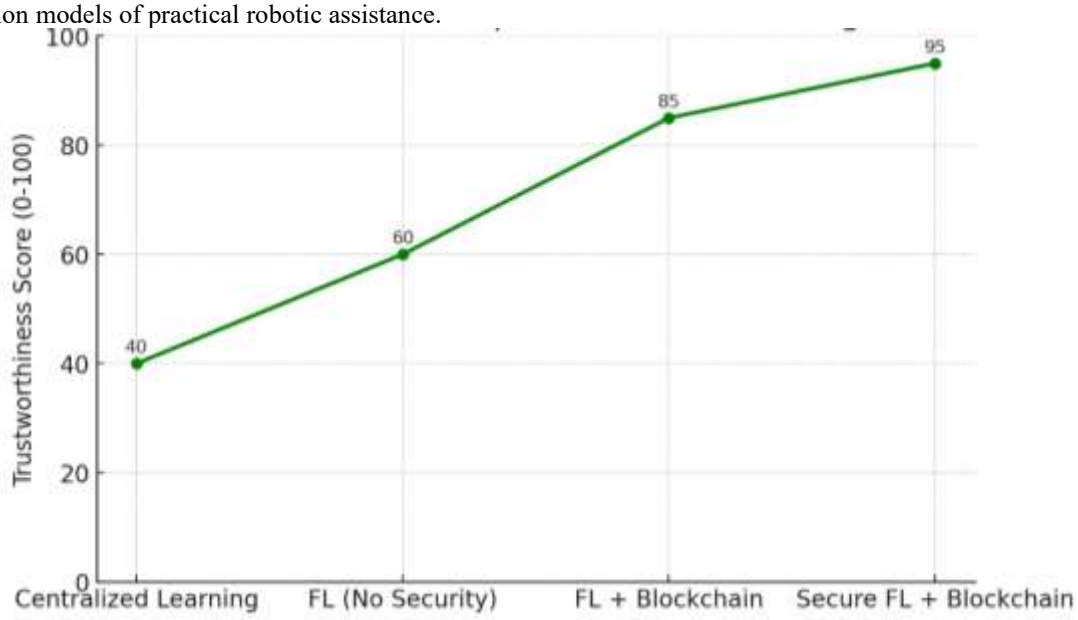

Figure 4. Trustworthiness of Model Updates Across Learning Paradigms: Impact of Permissioned Blockchain, Smart-Contract Governance, and Security Controls

Figure 4 indicated a gradual increase in trustworthiness under centralized learning followed by a sharp increase with the introduction of a permissioned blockchain, and finally, the ledger was further enriched with security controls, which led to an increment in trustworthiness. Centralized training was the lowest in score due to the fact that the provenance of updates was still transparent and the sole aggregation location was prone to manipulation. Unsecured federated learning had a better score because they stored the information locally, but the lack of verifiable trails made the model update vulnerable to replay and manipulation. By including a permissioned ledger, the score was boosted significantly through the introduction of unchangeable, signed records of every contribution. The maximum score was obtained when ledgering was combined with cryptographic and algorithmic defenses, producing traceable and tamper-evident training cycles.

The metric was an aggregated index indicating a number of audit signals: percent of updates whose cryptographic hashes match, percent of contributions by authenticated nodes, percent of detected inconsistent updates, percent of complete on-chain metadata (timestamps, signatures, validation summaries), and frequency of anomaly flags. High scores under blockchain occurred because of trustful provenance; all clients provided a signed digest of gradients or weights, and smart contracts verified identity, role, and policy verification before being accepted. Duplicates or invalid submissions were automatically filtered, and the ledger maintained an unalterable history of accepted events, and post-hoc audits were simple. There were also operational mechanisms that explained the improvement. Edge nodes created local updates and hashes and sent metadata; the orchestrator would check signatures and put receipts on the ledger; consensus (e.g., PBFT/Raft) finalized blocks and blocked silent rewrites. Smart-contract policies imposed penalties on non-conforming conduct and aided reputation acquisition, mitigating the effects of Sybil or colluding consumers. In the case of secure aggregation that hid per-client gradients and anomaly detection filtered outliers, only reputable aggregate contributions affected the world model. Reputation-weighted aggregation followed, which highlighted the nodes that had a history of reliability, which would strengthen integrity on a per-round basis.

Practical implications were of importance to safety-critical decision pipelines. Block interval tuning, batched commits, and lightweight consensus kept things responsive, although ledgering and security added overhead. Hashes and validation summaries based on work done by fall detection, activity recognition, and risk stratification were recorded on-chain and allowed the auditors to connect model gains with verifiable updates. Figure 4 showed the resulting trajectory where provenance immutability, authenticated participation, and layered defenses created the most reliable training process, and the training could be confidently deployed in environments where traceability and tamper resistance were required.

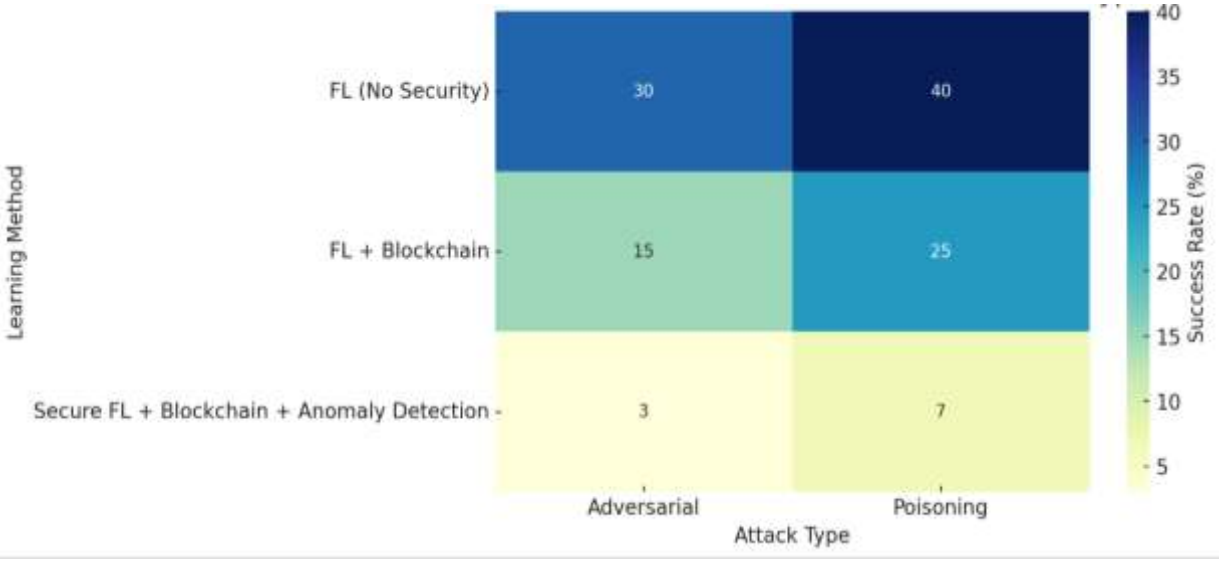

Figure 5. Attack Success Rates Across Training Methods and Threat Types: Poisoning vs. Adversarial Under Federated, Blockchain, and Anomaly-Detection Controls

Figure 5 provided a clear illustration of the reduced success rates of attacks with an increment of the strength of the defenses. Unprotected federated learning was the most vulnerable, and the attacks of poisoning and adversarial manipulation were successful in about 40 percent and 30 percent of the attempts, respectively. The implementation of a permissioned blockchain decreased these rates to approximately 25 and 15 percent, respectively, a sign of enhanced provenance, authenticated participation, and unalterable logging. The best figures were observed when anomaly detection was paired with the ledger: poisoning was reduced to approximately 7% and adversarial attempts to approximately 3%, meaning that the detection and filtering of suspicious updates were important to the integrity of training.

The reductions seen with the blockchain-only design came due to verifiable identities, tamper-evident update trails, and policy enforcement with smart contracts. These restrictions shortened replay, Sybil, and unauthorized contributions but made no direct analysis of the content of gradients or weights. Model-poisoning strategies keeping valid credentials were thus still non-trivially successful. The further drop with anomaly detection indicated that content-based screening, which was not limited to identity and auditability, was necessary to suppress advanced attacks that passed access checks but carried meanings that were malicious.

The anomaly-detection layer made use of statistical and geometric properties of incoming updates before aggregation. Various features, including layer-wise norm distributions, cosine similarity to the cohort median, gradient sign consistency, and loss-based residuals, were examined to serve as flags of outliers. The contributions that were seen as suspicious either received a down-weight or were quarantined, and reputation scores were stored in blockchain, where habitual violations were recorded. Such a pipeline constrained the impact of engineered gradients common to backdoor or scaling-based poisoning and at the same time suppressed small, focused perturbations related to adversarial examples.

Non-IID partitions of clients and resource limitations typical of edge robots were also reflected in experimental conditions and serve to facilitate attack effects and complicate detection. The identified trend in the heatmap showed that the greatest resilience was achieved when the following combinations were used: layered defenses, data locality with FL, authenticated coordination through blockchain, and content validation through the anomaly detection. Although privacy controls like differential privacy and secure aggregation are able to blur per-client cues, with prudent thresholding and strong aggregation, detection performance remained intact, and it was possible to reliably isolate malicious activity without deteriorating honest learning performance.

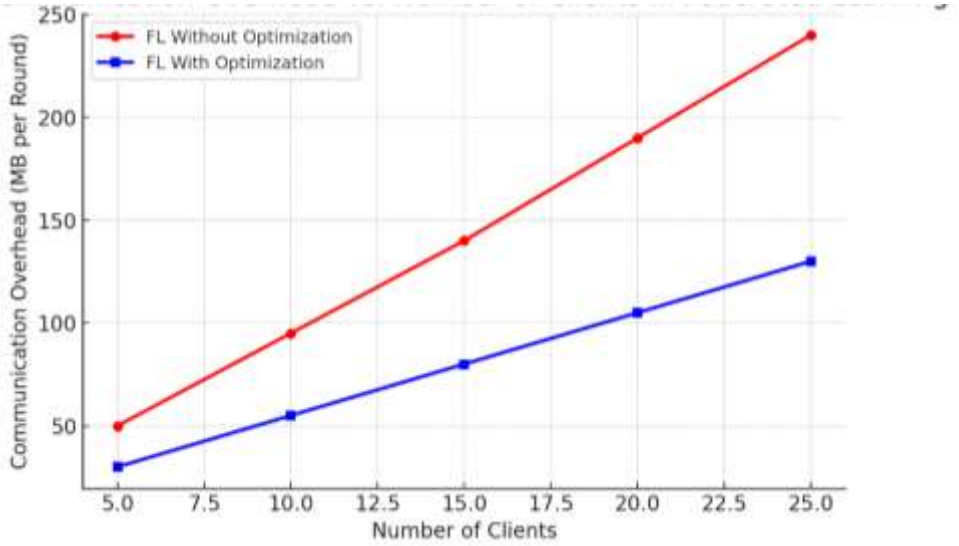

Figure 6. Communication Overhead vs. Client Scale under Optimized Aggregation: Effects of Sparsification, Quantization, and Periodic Local Training

Figure 6 indicated that communication overhead per round increased with the number of clients, although the slopes of the two configurations increased at significantly different rates. Overhead increased, without optimization, by a factor of four, from about 50 to 240 MB (5 clients to 25 clients). It decreased to some extent, from 130 MB with optimization, as compared to 30 MB. The scalability of the relative savings was similar, i.e., it was approximately 40% with 5 clients (50 Mb -30 Mb), 42 percent with 10 clients (95 Mb -55 Mb), 43 percent with 15 clients (140 Mb -80 Mb), 45 percent with 20 clients (190 Mb -105 Mb), and 46 percent with 25 clients (240 Mb -130 Mb). The curve separation showed that optimization minimized the intercept and the growth rate and gave lower bandwidth needs as federation increased.

The curve of best fit indicated a package of bandwidth-aware methods, that is, compressing updates prior to transmission and minimizing the rate and volume of exchanges. Common mechanisms were top-k sparsification of gradients (only the largest coordinates are transmitted), low-bit quantization (e.g., 8-/4-bit), periodic local training with E>1 epochs to amortize communication, and partial client participation on each round. Payloads were even more limited by delta encoding and update clipping. The combination of these measures reduced the number of bytes per round and maintained convergence in case the hyperparameters were adjusted to match model capacity and data heterogeneity.

Edge clients with a tight radio and power constraint performed well with system-level impacts. Reduced payload length reduced uplink time, minimum retransmissions in unfriendly links, and minimum energy consumed per round. The reduced communication footprint also left room for security metadata (e.g., signatures, secure-aggregation masks) and blockchain anchoring of update receipts without violation of bandwidth constraints. Reduced latency also allowed an increased number of rounds to be conducted all over the network with the same network capacity, thereby making model updates available in time to safety-critical assistance problems.

Analysis of trade-offs was still necessary. Quantization, or coarse sparsification, can slow convergence or be inaccurate when thresholds are too large. Periodic local training enhanced efficiency, but client drift was a threat in highly non-IID data; that effect was reduced by robust aggregation and periodic full-precision syncs. The optimized curve indicated a near-linear (yet less steep) slope, indicating that the communication could easily be scaled to the number of clients (particularly in combination with adaptive client sampling and resource-conscious scheduling). These findings aided in the choice of compression rates, local epochs, and involvement rates to achieve a particular bandwidth objective and, at the same time, ensure that the model performance was not compromised.

$$C = d \cdot b \cdot K$$

The cost depends on model size, parameter bits, and node count. It reflects bandwidth required during distributed training. Optimizing this cost ensures feasibility for low-power robots.

Table 6: Correlation Matrix of Vital Signs

Vital Signs	Heart Rate (bpm)	SpO_2 (%)	Temperature (°C)	Respiratory Rate (breaths/min)
Heart Rate (bpm)	1.000	-0.105	0.004	-0.073
SpO_2 (%)	-0.105	1.000	0.127	-0.127
Temperature (°C)	0.004	0.127	1.000	0.152
Respiratory Rate (breaths/min)	-0.073	-0.127	0.152	1.000

The analysis of vital signs in Table 6 presented significant information on the interdependence of physiological processes. There was a low negative relationship between heart rate and SpO_2 (-0.105), showing that higher heart rates were sometimes accompanied by lower oxygen saturation, which is also in line with the stress or exertion scenarios. The low correlation indicated that these factors interacted, but in some situations, they were mainly independent in the dataset, as they represented a variety of patient conditions.

A positive relationship was observed between temperature and respiratory rate (0.152); thus, breathing rate was slightly increased with a rise in body temperature, which is in line with clinical knowledge of hyperventilation caused by fever. Likewise, the association between SpO_2 and temperature (0.127) showed a weak positive association, with high levels of oxygen tending to be associated with high levels of body temperature, albeit the effect was weak. These weak associations highlighted how difficult the patient physiological reactions are and why multi-parameter monitoring should be used.
The general low to moderate correlation coefficients have indicated that a single vital sign did not very strongly predict another, and therefore, the use of integrated machine learning models is necessary to elucidate nonlinear and complicated interconnections. The fact that these signals are independent further supported the soundness of the dataset to undergo federated training, whereby varied parameter contributions were able to improve model generalization. This discussion affirmed the significance of the combination of numerous vital signs to be used in the correct prediction of health risks and robotic assistance.

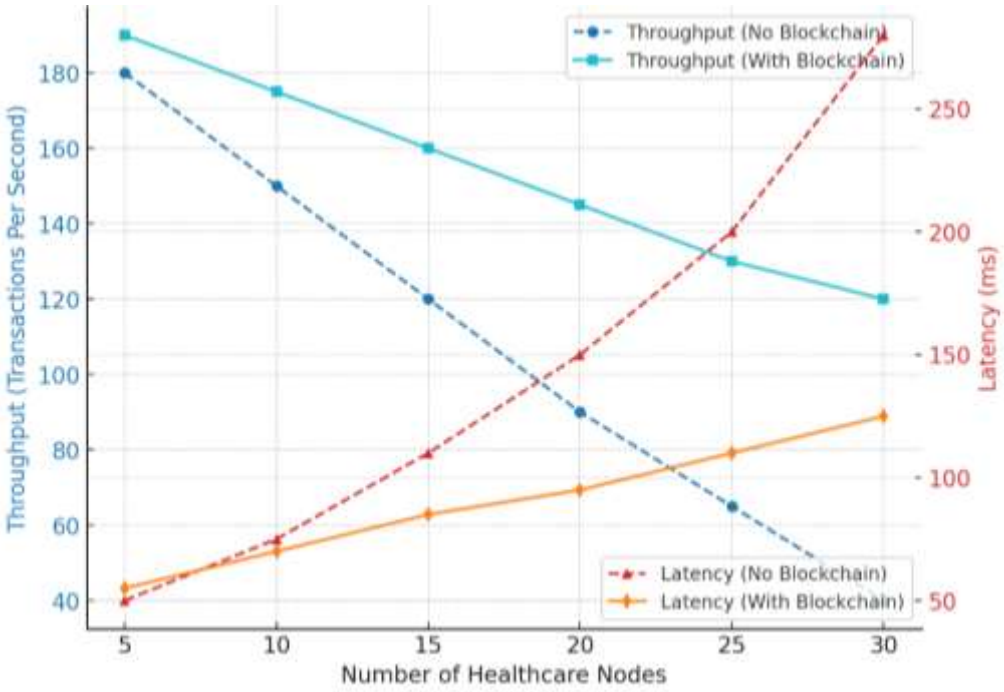

Figure 7. Scalability Improvements through Blockchain-Federated Learning Integration

Figure 7 showed how scalability, throughput, and latency changed with an increase in the number of nodes in a federated environment. In the case where blockchain integration was not used, throughput sharply decreased with an increase in network size, decreasing to only 40 TPS at 30 nodes compared to 180 TPS at 5 nodes. Concurrently, latency increased exponentially, almost to 280 ms on a larger scale. This trend showed that the established federated systems were not able to coordinate work among various healthcare nodes, which forced the decrease of performance and unreliable communication as the system was being scaled.

Conversely, federated integration that was backed by blockchain maintained much higher throughput rates, which started at 190 TPS and reached 120 TPS even when 30 nodes were involved. The findings indicated that blockchain facilitated consistent involvement of more than one node, as it provided consistency in validating updates and effective recording of transactions. Blockchain reduced the bottlenecks associated with large-scale federated deployments, which is achieved by offering a decentralized trust layer. Such stabilization permitted the system to operate efficiently with increased workloads, which directly and positively affected real-time data exchange and multi-robot cooperation.

A trend in latency also underscored the aspect of blockchain in enhancing scalability. In the absence of blockchain, the latency increased exponentially with the number of nodes in use, which is indicative of synchronization and verification bottlenecks. Nevertheless, the inclusion of blockchain kept the latency within a controlled value, and it had a moderate growth from 55 ms with 5 nodes to 125 ms with 30 nodes. This showed that blockchain systems like permissioned consensus and smart contracts reduced the communication overhead, decreased delays in federated aggregation, and enhanced the responsiveness in distributed decision-making processes.

It was generally found that blockchain integration did not only play a role in avoiding system instability but was also essential to facilitate large-scale federated deployments. Building more healthcare robots and institutions into the system and still maintaining its performance would be possible because the system maintained the throughput and minimal increases in the latency. This evidence highlighted the fact that blockchain-federated integration was the key to both scalability and robustness, which could enable secure and reliable expansion to address the needs of complex and real-world healthcare applications.

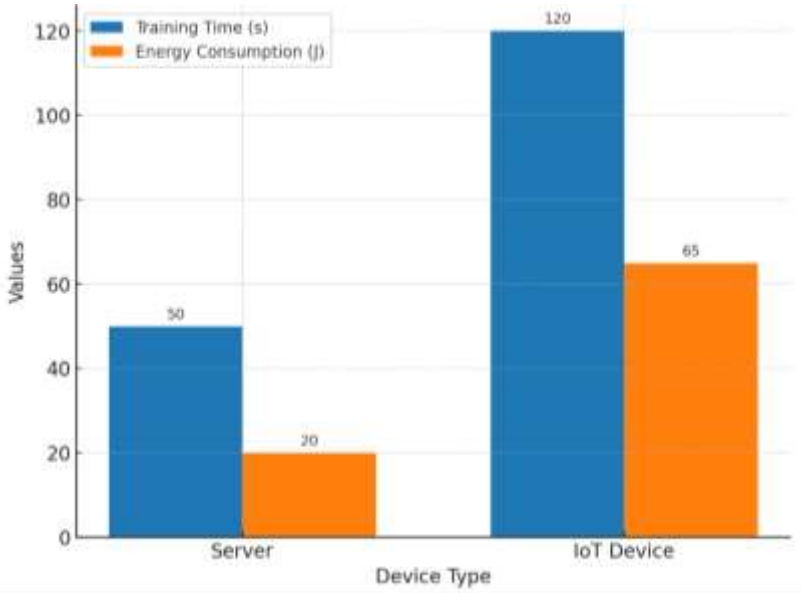

Figure 8. Resource-Aware Performance Analysis in Server and IoT Device Environments

The comparative analysis presented in Figure 8 compared servers and IoT devices in terms of training time and energy consumed during the distributed training of models. Findings revealed that servers took 50 seconds to train, and comparable work took better than 120 seconds by IoT devices. This difference underscored the limitations of the computational capabilities of resource-constrained devices, in which reduced processing power and memory bandwidth increased the training time. The fact that the work with less lightweight federated algorithms and model compression algorithms provided a slower performance on IoT devices proved the significance of lightweight federated algorithms and model compression techniques to limit computation delays at the cost of reasonable accuracy.

Another important observation was the trends of energy consumption. The average training of servers required 20 joules, whereas the same process required about 65 joules of the IoT devices. This increased energy demand was associated with long computation time and resource wastefulness on limited devices. The results showed that energy-sensitive optimization mechanisms, including adaptive update frequency and energy-saving scheduling, were to be applied to extend the device usability in real-time applications. Federated deployment on mobile or robotic healthcare devices can be rendered sustainable by circumventing power constraints and without compromising system reliability.

It was also found during the analysis that despite the increased time and energy required by the IoT devices, their performance was still within reasonable bounds to be performed in collaborative learning. The ability to train on these devices demonstrated the flexibility of federated mechanisms with limited resources so that distributed healthcare applications could remain effective with decentralized settings. The variations among the categories of devices confirmed the topicality of the development of strategies of task sharing, in which computationally intense operations could be delegated to edge servers, whereas lightweight tasks could be addressed locally by robots with IoT capabilities.

It was highlighted in Figure 8 that the balance between training efficiency and energy consumption was crucial to system feasibility under resource constraints. Federated structures can also expand the mesh of devices involved to a wide variety of devices in the IoT by adapting algorithms to accommodate constraints of the device, thereby preserving a steady level of network performance. This was critical in the creation of a strong and inclusive ecosystem in which both strong servers and weak edge devices worked together to provide scalable and trustworthy intelligent healthcare support.

Table 7: Distribution of Robot Assistance Types by Activity

Activity Type	Emergency Alert	Medication Reminder	Mobility Support	Monitoring
Exercising	26.1%	21.7%	17.4%	34.8%
Resting	31.8%	13.6%	22.7%	31.8%
Sleeping	18.5%	22.2%	18.5%	40.7%
Walking	14.3%	28.6%	32.1%	25.0%

Table 7 identified that the requirements of robot assistance were very different in patient activities. Patients demonstrated greater use of monitoring assistance (34.8%) and emergency alerts (26.1%) during exercise due to the high risk of physical activity or sudden medical complications. Medication reminders (21.7) and mobility support (17.4) were relatively less likely to be in demand, and therefore it can be argued that real-time supervision was necessary, whereas direct support interventions were less common in active states. This distribution also showed the significance of adaptive robotic systems, which put more emphasis on situational awareness when performing physically intensive tasks.

The assistance needs during resting and sleeping processes were more towards monitoring functions, with resting having 31.8 percent monitoring needs and sleeping having 40.7 percent, the highest needs. Medication reminders were also significant in states of sleep (22.2%), which suggests the value of scheduled interventions even when people have low activity levels. Mobility assistance was average during the condition of resting (22.7) and during sleeping (18.5), which implies the use of some assistance periodically, like adjusting the posture or repositioning. The relatively high rate of emergency alerts when in the resting state (31.8) highlighted that even low-exertion states may lead to health threats, and thus constant attention by healthcare robots is required.

Mobility support (32.1) and medication reminders (28.6) were found to be the most needed activities, which can be discussed as the dual relevance of movement support and adherence to treatment in ambulatory conditions. Assistance with monitoring (25.0%) and emergency alerts (14.3) was there but relatively minor, which emphasized that walking patients were more served by proactive assistance than by emergency intervention. These results highlighted the necessity of a situation-sensitive robotic system in which the orders of assistance were dynamically ranked based on the activity patterns. The patient's safety, comfort, and treatment adherence would be optimized through the process of personalizing robotic interventions according to the recognition of real-time activities.

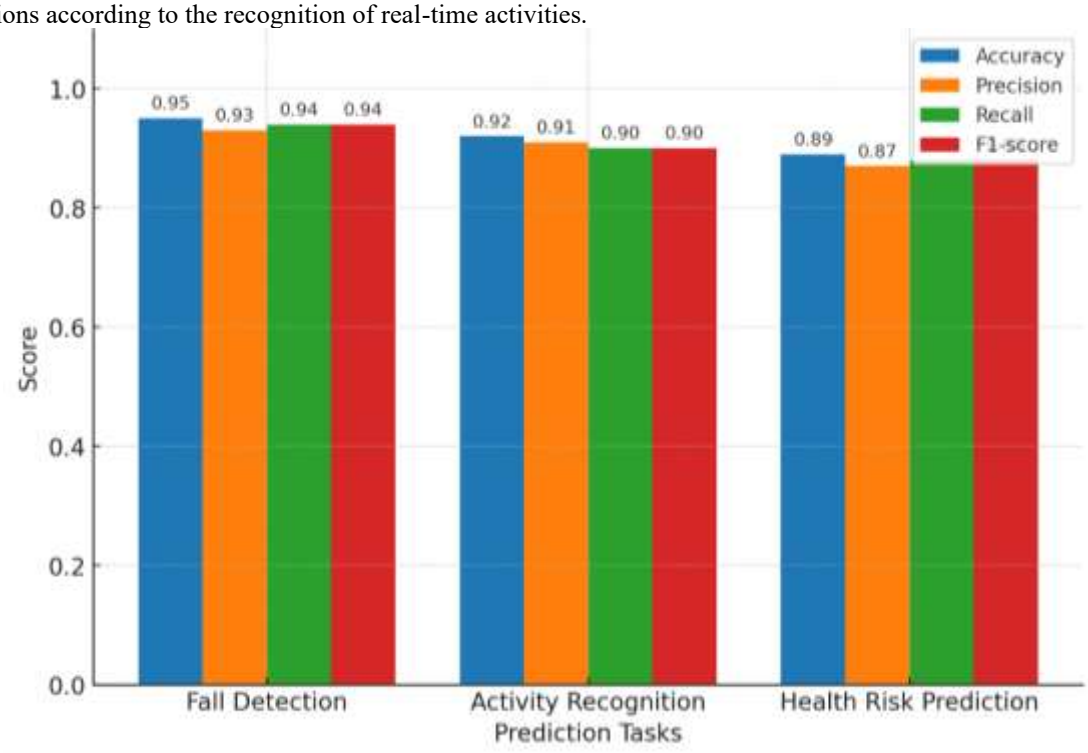

Figure 9. Comparative Performance of Prediction Models for Fall Detection, Activity Recognition, and Health Risk Assessment

The prediction accuracy of three key tasks—fall detection, activity recognition, and health risk prediction—is demonstrated in figure 9. Fall detection had the best overall performance, with an accuracy of 0.95 and a balanced precision, recall, and F1-score of around 0.93-0.94. This good performance indicated the ability of the model to distinguish reliably fall and non-fall events based on physiological and contextual data. False negatives might cause unsafe situations of occurrence of a fall not detected, and high recall values were especially significant in the fall detection, as they might affect patient safety directly.

Activity recognition scores a little lower, but again strong performance, with an accuracy of 0.92 and a precision, recall, and F1-score in the 0.90-91 range. These outcomes indicated the fact that such routine practices like resting, walking, or exercising could be categorized according to the given model with enough accuracy. The close similarity of all four measures indicated the balancing of the classifier was consistent between the detection of genuine activities and misclassification. This reliability was imperative to make sure that the assistance robots could adapt their behavior in real time depending on the current state of the patient without taking any unneeded measures.

The hardest task to forecast was health risks; the performance was slightly lower compared to the other two, with an accuracy of 0.89 and a range of 0.87 to 0.88 in the range of the precision, recall, and F1-scores. The relatively lower performance was due to the intrinsic complexity of multi-class classification that included heterogeneous vital signs and environmental variables. The findings still suggested that the model has the potential to predict different levels of patient risk successfully, though with continued optimization or hybrid boosting like transfer learning techniques, predictability could be improved. The accurate risk assessment was still important to prioritize patients to be immediately treated and maximize the use of resources in healthcare.

The comparative analysis revealed that all three models performed well, but each of the tasks had different requirements in precision or focus on recall. Fall detection placed high recall as the first priority to reduce the number of undetected fall cases, activity recognition needed to be balanced based on various classes, and health risk prediction demanded sophisticated management of complex features in order to reduce misclassification. Collectively, these findings indicated that predictive modeling might markedly improve the outcomes of robotic assistance, providing a timely intervention, personalized assistance, and better patient monitoring in a distributed setting.

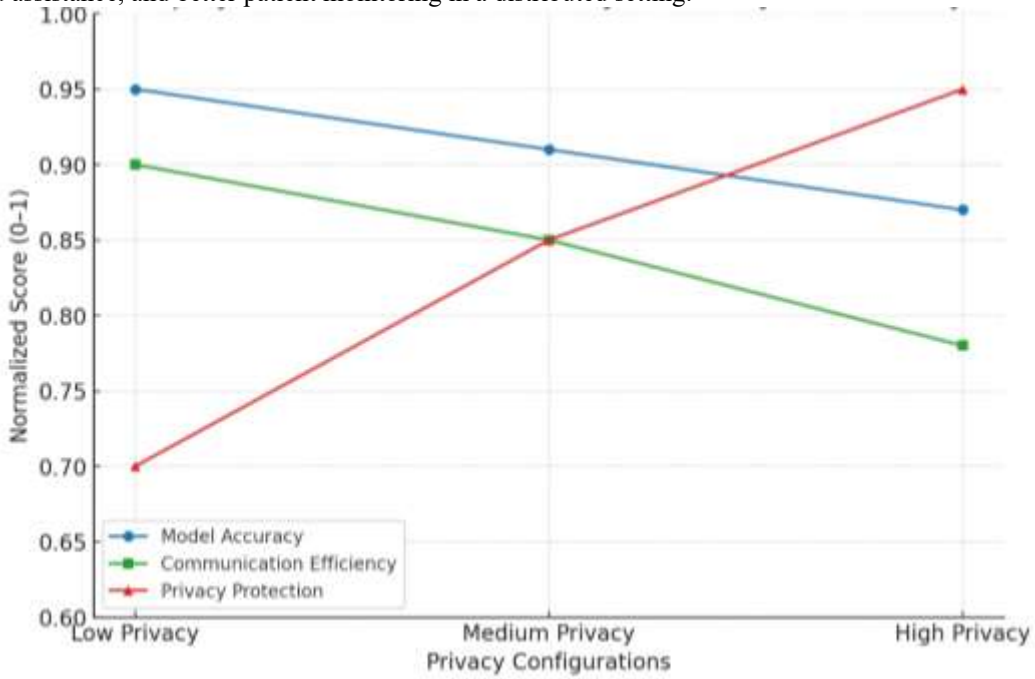

Figure 10. Deployment Trade-Offs among Accuracy, Privacy, and Communication Efficiency under Varying Privacy Levels

The complexity of the accuracy, efficiency, and privacy balance in using various privacy settings was emphasized by Figure 10. With low privacy settings, the model was found to work with the accuracy of 0.95 and the efficiency of 0.90, which is substantial model performance and low communication overhead. But the privacy score was much lower; it was 0.70, which exhibited lesser protection of sensitive data. This arrangement was conducive to the system performance but also subjected to a risk of exposing critical patient information, and it was seen that the low privacy ensured high accuracy but created vulnerabilities that are not appropriate in sensitive settings.

The trade-off became more equal at medium privacy. The accuracy dropped to 0.91, and efficiency dropped to 0.85, but the protection of privacy was significantly better at 0.85. This setup showed that privacy protection in the medium setting of differential privacy with controlled noise or lightweight secure aggregation could preserve good performance without badly compromising privacy. The relative decrease in efficiency and accuracy was offset by the significant increase in privacy guarantee, which makes this environment a feasible tradeoff in distributed systems that must be both trusted and operational.

Higher privacy settings moved the balance further, with the accuracy down to 0.87, the efficiency to 0.78, and the privacy protection being at the greatest level at 0.95. This revealed that more privacy-sensitive systems, like high-noise differential privacy or intensive encryption, were more confidential but limited system performance. The accuracy drop indicated that sensitive features in the data set had been corrupted, whereas the efficiency drop was due to increased computational and communication demands. These arrangements worked quite well in highly controlled conditions but posed a risk of lower responsiveness and learning efficacies in real-time applications.

The analysis showed that deployment had to balance itself on system objectives. The low levels of privacy were conducive to performance rather than protection, the medium levels of privacy were an appropriate balance of trade-off to general use, and the high levels of privacy were very conducive in terms of confidentiality at the expense of the operations. Such

findings highlighted the significance of adaptability mechanisms that dynamically changed privacy and efficiency configurations with regard to task urgency. In real-life deployment, a trade-off framework was flexible to allow consistent results of learning and preserve sensitive information in actual distributed settings.

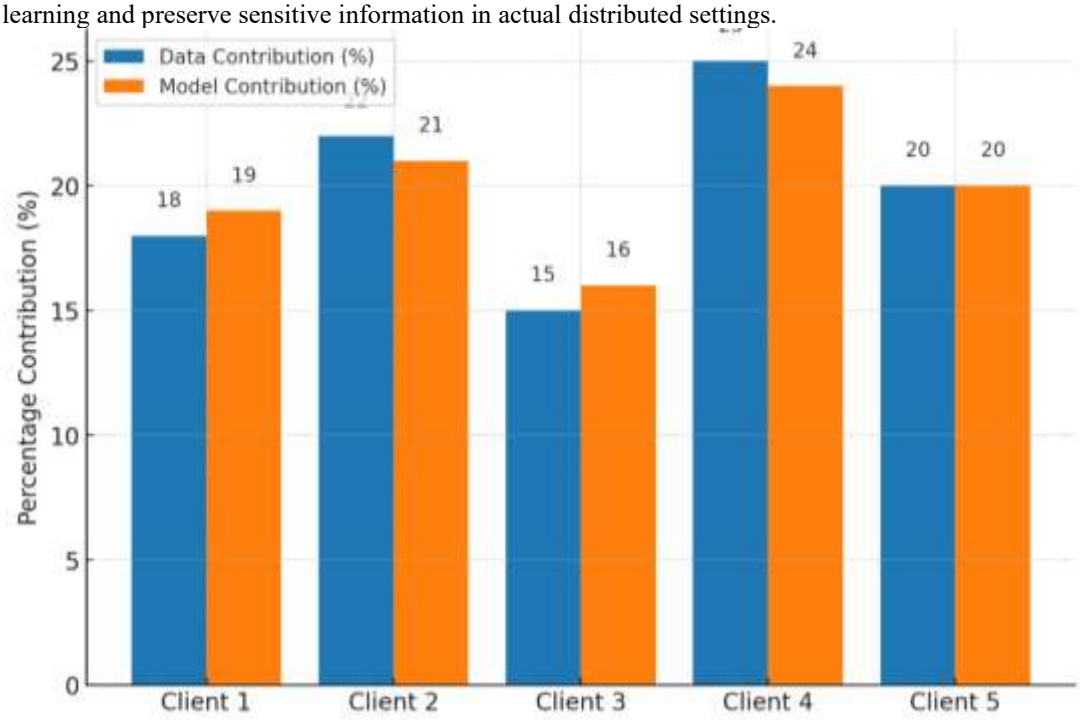

Figure 11. Client Contribution Distribution: Data Share vs. Global Model Influence across Federated IoT Nodes (with Privacy and Blockchain Controls)

As shown in Figure 11, five clients were represented with paired bars where each node is proportional to its contribution to local data, and the contribution is determined by the global model. Client 4 had the most significant share of data (25), and its contribution to the model was almost identical (24). Client 2 has the highest contribution of 22 percent of data and 21 percent of model impact, whereas Client 5 had a perfect alignment of 20 percent in both measures. Clients 1 and 3 (smaller datasets of 18 and 15 percent) produced a very little bit higher model contribution (19 and 16 percent). The general trend was near proportionality, which implied that aggregation was fair to clients, and the slight violations could be explained by the quality of data and the composition of classes.

These deviations were predicted with non-IID conditions. Nodes with rarer yet clinically significant events (e.g., falls or critical vitals) would have a marginal utility that was outsized, such that it would be better in the overall goal than the raw volume would have indicated. On the other hand, the relative impact of more redundant sample clients was mildly reduced. Validation deltas were used to obtain contribution estimates on a common holdout; hence, label noise, sensor variance, and temporal coverage (timestamps across activity cycles) all contributed to the observed spread. Differential privacy noise and the robust aggregation dampened extreme effects in a subtle manner and contributed to making the training more stable without eliminating the real signal among smaller yet more varied clients.

Contributions were also influenced by security and governance mechanisms. Protecting aggregation in the context of per-client updates was not denied but allowed aggregate performance effects to be computed. Anomaly detection weights suspicious gradients down, and so the poisoned nodes do not inflate the apparent utility. The permissioned blockchain stored hashed updates and per-round validation summaries, which allowed attributing impact over time in an auditable manner. The incentive rules that smart contracts would have enforced, that is, rewarding clients based on proven contribution, punishing recurring anomalies, and limiting domination, would have ensured that cooperative behavior would be enforced even in the premise of heterogeneous data holdings.

The distribution informed the policies of scheduling and fairness operationally. Adaptive client sampling was used to prioritize underrepresented cohorts to balance the coverage of classes during subsequent rounds. The fairness constraints in weighted FedAvg reduced over-representation by a single, rich node, and regular full-precision synchronizations reduced drift due to compression and privacy measures. The close proportionality of allocation of influence across clients was a sign of the system assigning influence in proportion to useful information, rather than the volume of data, to equitable rewarding, effective bandwidth budgeting, and reliable model evolution in a resource-constrained federated environment.

Limitations and Future Work

The framework was good but had some shortcomings. First, privacy-preserving methods, including differential privacy and secure aggregation, increased the model accuracy on particular tasks, albeit by a small margin due to the computational overhead. Second, blockchain interfaces added both latency and increased energy usage, which can limit application to devices with the most resource-constrained IoT. Third, the dataset was not large enough and might not be representative of highly heterogeneous healthcare settings. Also, the use of anomaly detection models is effective but might fail to identify very advanced adversarial tactics. Last, the simulations were performed within a simulated setting, and large-scale real-world implementation is yet to be confirmed.

The limitations are to be addressed in the future by means of sophisticated optimization and new technologies. It is possible to investigate quantum-safe cryptography in order to equip the system for future cryptographic threats. Urgent need:6G-enabled ultra-reliable low-latency communication (URLLC) could be greatly effective to improve real-time operation of healthcare robots. The addition of neuromorphic computing and AI accelerators can save the computational load of the federated learning and blockchain operation on the IoT devices. Generalizability in other health care settings will be enhanced by expanding the data with bigger and more varied populations of patients. Digital twins can also be incorporated to simulate healthcare situations to make the robot intervention a safer technique. Lastly, architectures of hybrid blockchain-edge must be explored to trade off scalability, privacy, and efficiency in practical multi-institutional healthcare networks.

Conclusion

The presented work was a thorough framework integrating federated learning, blockchain, and optimization with AI to increase the reliability, security, and efficiency of healthcare robotics. Through federated learning, predictive models would be trained simultaneously using distributed IoT-enabled data without infringing patient confidentiality, thereby ensuring data locality without interfering with the correct identification of activities, predicting health risks, fall prevention, and optimizing assistance. Integration of blockchain created a trust layer that is unchangeable, maintaining a safe track of model changes, and accessing data through smart contracts providing transparency, as well as consensus-considered validation to block tampering and malicious agent involvement. Confidentiality was further reinforced by privacy-preserving mechanisms such as differential privacy and secure aggregation, and it was demonstrated that AI-driven anomaly detection models were also effective in mitigating vulnerabilities to poisoning attacks and adversarial updates.

The usefulness of the framework was verified through experimental assessments on a real-world dataset and showed that it provides competitive model accuracy, minimized communication overhead, scalability to resource-constrained situations, and resilience to attacks on a system. The trade-off analysis demonstrated the trade-off between accuracy, privacy, and efficiency and provided realistic suggestions to be implemented for a practical use of the IoT-driven healthcare setting. The framework dealt with technical and security issues related to the deployment of healthcare robotics on a large scale by incorporating lightweight blockchain protocols, resource-aware federated strategies, and adaptive AI techniques.

In general, the research proved a scalable and reliable roadmap that permits safe real-time monitoring of patients and robotic assistance, opening the way to the next-generation smart healthcare systems. The future trajectories can be further led to quantum-safe cryptography, 6G-based ultra-reliable communications, and the integration of digital twins to further advance secure and adaptive healthcare robotics.

REFERENCES

[1] Rahman, A., Kundu, D., Debnath, T., Rahman, M., & Islam, M. J. (2024). Blockchain-based ai methods for managing industrial iot: Recent developments, integration challenges and opportunities. arXiv preprint arXiv:2405.12550.

[2] Charfare, R. H., Desai, A. U., Keni, N. N., Nambiar, A. S., & Cherian, M. M. (2024). IoT-AI in Healthcare: A Comprehensive Survey of Current Applications and Innovations. International Journal of Robotics & Control Systems, 4(3).

[3] Rauniyar, A., Hagos, D. H., Jha, D., Håkegård, J. E., Bagci, U., Rawat, D. B., & Vlassov, V. (2023). Federated learning for medical applications: A taxonomy, current trends, challenges, and future research directions. IEEE Internet of Things Journal, 11(5), 7374-7398.

[4] Alsadie, D. (2024). Artificial intelligence techniques for securing fog computing environments: trends, challenges, and future directions. IEEE Access.

[5] Rahman, A., Debnath, T., Kundu, D., Cerasuolo, F., Islam, M. J., Rahman, M., ... & Pescapè, A. (2024). Unlocking the Potential of IoT, AI, and Blockchain in Transforming Public and Private Industries. Cambridge Scholars Publishing.

[6] Asimiyu, Z. (2024). Integrating AI, IoT, and Machine Learning in IRS-Enhanced MIMO Systems: A Paradigm Shift in Connectivity.

[7] Khatun, M. A., Memon, S. F., Eising, C., & Dhirani, L. L. (2023). Machine learning for healthcare-iot security: A review and risk mitigation. IEEE Access, 11, 145869-145896.

[8] Gera, B., Raghuvanshi, Y. S., Rawlley, O., Gupta, S., Dua, A., & Sharma, P. (2023). Leveraging AI-enabled 6G-driven IoT for sustainable smart cities.

International Journal of Communication Systems, 36(16), e5588.

[9] Aslam, A. B., Iqbal, F., Talpur, U., Syed, Z. S., & Shaikh, F. K. (2024). Artificial Intelligence-Enabled 6G Mobile Systems. In Intelligent Technologies for Healthcare Business Applications (pp. 49-79). Cham: Springer Nature Switzerland.

[10] Nagarjun, A. V., & Rajkumar, S. (2024). Exploring the potential of deep learning and blockchain for intrusion detection systems: A comprehensive review. Journal of Circuits, Systems and Computers, 33(16), 2430007.

[11] Saraswat, D., Bhattacharya, P., Verma, A., Prasad, V. K., Tanwar, S., Sharma, G., ... & Sharma, R. (2022). Explainable AI for healthcare 5.0: opportunities and challenges. IEEe Access, 10, 84486-84517.

[12] Islam, S., Karam, A., Boualem, A., Pathan, K. V., Sayed, A. F., & Rahman, A. U. (2024, May). 6G-Enabled Network Security in Drone Technology: Revolutionizing Communications and Enhancing Security Applications. In Nordic e-Infrastructure Collaboration Conference (pp. 28-44). Cham: Springer Nature Switzerland.

[13] Islam, S., Karam, A., Boualem, A., Pathan, K. V., Sayed, A. F., & Rahman, A. U. (2024, May). 6G-Enabled Network Security in Drone Technology: Revolutionizing Communications and Enhancing Security Applications. In Nordic e-Infrastructure Collaboration Conference (pp. 28-44). Cham: Springer Nature Switzerland.

[14] Tariq, U., Ahmed, I., Khan, M. A., & Bashir, A. K. (2023). Fortifying IoT against crimpling cyber-attacks: a systematic review. Karbala International Journal of Modern Science, 9(4), 9.

[15] Singh, A., Sahu, K., & Gupta, B. K. (2024). Securing Industry 5.0 data through unified emerging technologies. In Computational Intelligence Applications in Cyber Security (pp. 31-48). CRC Press.

[16] Khan, R. S., Sirazy, M. R. M., Das, R., & Rahman, S. (2022). An ai and ml-enabled framework for proactive risk mitigation and resilience optimization in global supply chains during national emergencies. Sage Science Review of Applied Machine Learning, 5(2), 127-144.

[17] Junaid, S. B., Imam, A. A., Balogun, A. O., De Silva, L. C., Surakat, Y. A., Kumar, G., ... & Mahamad, S. (2022, October). Recent advancements in emerging technologies for healthcare management systems: a survey. In Healthcare (Vol. 10, No. 10, p. 1940). MDPI.

[18] Ahsani, V., Rahimi, A., Letafati, M., & Khalaj, B. H. (2023). Unlocking metaverse-as-a-service the three pillars to watch: Privacy and security, edge computing, and blockchain. arXiv preprint arXiv:2301.01221.

[19] Babar, A. Z., & Akan, O. B. (2024). Sustainable and precision agriculture with the internet of everything (IoE). arXiv preprint arXiv:2404.06341.

[20] Radanliev, P., De Roure, D., Maple, C., Nurse, J. R., Nicolescu, R., & Ani, U. (2024). AI security and cyber risk in IoT systems. Frontiers in Big Data, 7, 1402745.

[21] Ojika, F. U., Owobu, W. O., Abieba, O. A., Esan, O. J., Ubamadu, B. C., & Daraojimba, A. I. (2024). The Role of AI in Cybersecurity: A Cross-Industry Model for Integrating Machine Learning and Data Analysis for Improved Threat Detection. Comput Secur.[Year].

[22] Singh, A. R., Kumar, R. S., Madhavi, K. R., Alsaif, F., Bajaj, M., & Zaitsev, I. (2024). Optimizing demand response and load balancing in smart EV charging networks using AI integrated blockchain framework. Scientific Reports, 14(1), 31768.

[23] Harahsheh, K. M., & Chen, C. H. (2023). A survey of using machine learning in IoT security and the challenges faced by researchers. Informatica, 47(6).

[24] Quan, M. K., Nguyen, D. C., Nguyen, V. D., Wijayasundara, M., Setunge, S., & Pathirana, P. N. (2024). Toward privacy-preserving waste classification in the Internet of Things. IEEE Internet of Things Journal, 11(14), 24814-24830.

[25] Zar, A., Zar, L., Mohsen, S., Magdi, Y., & Zughaier, S. M. (2024). A comprehensive review of algorithms developed for rapid pathogen detection and surveillance. Surveillance, prevention, and control of infectious diseases: An AI perspective, 23-49.

[26] Mohbey, K. K., & Acharya, M. (2023). The role of machine learning in the advancement of 6G technology: Opportunities and challenges. 6G Enabled Fog Computing in IoT: Applications and Opportunities, 309-331.

[27] Serôdio, C., Cunha, J., Candela, G., Rodriguez, S., Sousa, X. R., & Branco, F. (2023). The 6G ecosystem as support for IoE and private networks: Vision, requirements, and challenges. Future Internet, 15(11), 348.

[28] Hemnath, R. (2024). XGBoost-Based Botnet Detection Architecture for IoT Networks in Cloud Platforms. Int. J. of Multidisciplinary and Current research, 12.

[29] The, T. H., Pham, Q. V., Pham, X. Q., Do-Duy, T., & Reddy Gadekallu, T. (2023). AI and Computer Vision Technologies for Metaverse. Metaverse Communication and Computing Networks: Applications, Technologies, and Approaches, 85-124.

[30] Pang, T. Y., Lee, T. K., & Murshed, M. (2023). Towards a new paradigm for digital health training and education in Australia: exploring the implication of the fifth industrial revolution. Applied Sciences, 13(11), 6854.

[31] Shafik, W. (2024). Connected healthcare—the impact of Internet of Things on medical services: merits, limitations, future insights, case studies, and open research questions. In Artificial Intelligence and Internet of Things based

Augmented Trends for Data Driven Systems (pp. 181-217). CRC Press.

[32] Ouaissa, M., Ouaissa, M., Khan, I. U., Boulouard, Z., & Rashid, J. (Eds.). (2024). Low-power wide Area network for large Scale internet of things: architectures, communication protocols and recent Trends. CRC Press.

[33] Allioui, H., & Mourdi, Y. (2023). Exploring the full potentials of IoT for better financial growth and stability: A comprehensive survey. Sensors, 23(19), 8015.

[34] Gadekallu, T. R., Huynh-The, T., Wang, W., Yenduri, G., Ranaweera, P., Pham, Q. V., ... & Liyanage, M. (2022). Blockchain for the metaverse: A review. arXiv preprint arXiv:2203.09738.

[35] Patel, B., & Bhatia, J. (2024). A comprehensive review of internet of things and cutting-edge technologies empowering smart farming. Current Science (00113891), 126(2).

[36] Eldeeb, H. B., Naser, S., Bariah, L., Muhaidat, S., & Uysal, M. (2024). Digital Twin-Assisted OWC: Toward Smart and Autonomous 6G Networks. IEEE Network, 38(6), 153-162.

[37] Hurrah, N. N., Khan, E., & Parah, S. A. (2023). Smart ecosystems for sustainable development: Opportunities, challenges, and solutions. Intelligent Multimedia Signal Processing for Smart Ecosystems, 3-28.

[38] Onakpojeruo, E. P., Al-Turjman, F., Mustapha, M. T., Altrjman, C., & Ozsahin, D. U. (2022, October). Emerging AI and cloud computing paradigms applied to healthcare. In IET Conference Proceedings CP815 (Vol. 2022, No. 20, pp. 811-826). Stevenage, UK: The Institution of Engineering and Technology.

[39] Patange, A., Patil, D. P., Jadhav, P., Balguri, P. K., Singh, K. N., & Jadhav, N. N. (2024, January). Impact of Machine Learning Models in the advancement of Healthcare. In 2024 International Conference on Healthcare Innovations, Software and Engineering Technologies (HISET) (pp. 215-218). IEEE.

[40] Garikipati, V., Ubagaram, C., Dyavani, N. R., Jayaprakasam, B. S., & Hemnath, R. (2023). Hybrid AI models and sustainable machine learning for eco-friendly logistics, carbon footprint reduction, and green supply chain optimization. Journal of Science and Technology, 8(12), 230-255.