Decentralized AI For Home Health Monitoring

Leeladhar Revankar
leeladharrevankar51@gmail.com

Valencia Veigas
velanciaveigas2@gmail.com

***Abstract*- The growing demand for personalized and privacy preserving healthcare solutions has accelerated interest in federated learning (FL) for in-home health monitoring. This paper proposes a cloud-edge based personalized FL framework that enables intelligent data processing directly at the source—wearable and smart home devices—without transmitting raw data to centralized servers. By leveraging edge computing, the system ensures low-latency model updates while reducing bandwidth consumption and privacy risks. The proposed architecture supports personalized model training to address data heterogeneity across users, enhancing accuracy and adaptability in real-world scenarios. A case study involving elderly patients with wearable sensors demonstrates improved model performance and reduced latency compared to centralized approaches. This research highlights the potential of combining cloud-edge infrastructures with federated learning to deliver scalable, secure, and context-aware health monitoring systems for modern smart homes.**

1 INTRODUCTION

The healthcare industry is in the midst of a major shift, driven by an aging global population and the rapid growth of Internet of Things (IoT) technologies. At the heart of this transformation is in-home health monitoring—a system where wearable devices and smart sensors continuously track patients' vital signs, enabling proactive care instead of reactive treatment. These tools generate a steady stream of health data that, if used correctly, can support early diagnosis and personalized healthcare decisions through artificial intelligence (AI).

However, traditional AI systems typically rely on centralized data processing, where patient data must be transferred to a central server for training and analysis. This approach creates serious concerns around privacy, data security, and latency. It increases the risk of data breaches and often fails to comply with strict privacy laws such as the Health Insurance Portability and Accountability Act (HIPAA). Moreover, central servers can become bottlenecks, slowing down critical insights in real-time healthcare scenarios.

This is where Federated Learning (FL) steps in. Instead of sending sensitive data to a centralized location, FL allows AI models to be trained directly on the users' devices. The data never leaves the local system—instead, only model updates are shared. This decentralized approach inherently supports data privacy and sovereignty, making it especially well-suited for personal and medical applications.
While FL offers a solid foundation, it still faces limitations, particularly when dealing with non-identical data non-from diverse patients. Health data can vary widely from person to person, and treating every user the same leads to less accurate models. To overcome this, we turn to Personalized Federated Learning (PFL), a technique that allows the system to adapt and fine-tune the model for everyone's unique patterns and behaviors.

To further enhance responsiveness and privacy, this research proposes a cloud-edge PFL framework—a structure that brings computation closer to the user by utilizing edge devices. Processing data locally not only reduces communication delays but also limits exposure to external systems, strengthening patient privacy. On top of this, the framework incorporates blockchain technology, using smart contracts to maintain an immutable and transparent record of data usage, access control, and model updates.

This paper introduces a three-layer architecture (client-edgecloud) designed specifically for in-home health monitoring. It brings together the advantages of FL, edge computing, PFL, and blockchain to create a system that is not only intelligent and personalized, but also secure, scalable, and ethically sound. A real-world case study is presented to demonstrate the effectiveness of this framework in improving health predictions while maintaining privacy and performance standards.

2 LITERATURE REVIEW

Over the past few years, the combination of federated learning (FL), edge computing, and personalized healthcare has received growing attention, especially with the rise of smart homes and wearable health devices. Researchers and technologists alike are exploring how to create intelligent, privacy-conscious systems that can operate efficiently in real-world environments like inhome health monitoring.

One of the early contributions in this space comes from researchers who recognized the limitations of centralized healthcare systems—particularly in how they handle sensitive data. These systems often require uploading private patient information to the cloud, which raises both privacy and regulatory concerns. Federated Learning emerged as a solution, allowing machine learning models to be trained locally on users' devices without moving raw data across networks.

In a notable study, Wu et al. introduced a cloud-edge framework for FL that distributes model training between edge devices and cloud servers. This setup helps in reducing latency and makes real-time health monitoring more practical. The study demonstrated how edge nodes could handle local data aggregation and personalization, keeping the cloud responsible for more generalized tasks.

Building on that foundation, several surveys like those by Nguyen et al. and others have offered a broader view of FL in healthcare. These works have emphasized the unique challenges in this domain—such as non-identical data distributions (non-IID), limited computing power on devices, and the critical importance of data privacy. They also shed light on how FL can help

healthcare systems collaborate without sharing raw data across hospitals or regions.

To improve model accuracy across diverse users, researchers have explored Personalized Federated Learning (PFL). Rather than treating all devices the same, PFL adapts the model to individual user behavior. This is especially useful in in-home health monitoring where each patient might have different symptoms, routines, and medical needs. Approaches like FedPer and cluster-based FL have shown promise in tailoring models without sacrificing the benefits of global collaboration.

Privacy and security are recurring themes in this research space. Frameworks like FRESH (Federated learning with Ring Signatures and Homomorphic encryption) aim to secure the model updates being shared between clients and servers. These methods ensure that sensitive health patterns cannot be inferred even from model gradients.

At the practical level, systems like FedStack have been designed to monitor physical activities in smart home environments. They use stacked models at the edge to improve recognition of patient behaviors and detect anomalies early. Similarly, digital twin-based FL systems are being explored for chronic disease management. These systems simulate patient conditions using virtual models and apply FL to personalize treatments over time.

Recent studies are also considering the role of blockchain and smart contracts to enhance transparency and trust. Smart contracts can automate participation rewards, enforce data usage policies, and provide audit trails—making the overall system more secure and reliable. The intersection of federated learning, edge computing, and personalized healthcare has sparked a growing wave of innovation and research. As healthcare systems move toward decentralized, patient-centric models, researchers have explored how these technologies can work together to provide secure, intelligent, and efficient health monitoring—especially in home-based settings.

Federated Learning (FL) has gained attention for its ability to train machine learning models without collecting raw data in a central location. This feature is particularly relevant in healthcare, where data privacy is not just important—it is mandatory. Instead of gathering data in a single repository, FL allows hospitals, wearable devices, and mobile health apps to collaborate on AI model training by only exchanging model updates. This reduces privacy risks and supports compliance with regulations like HIPAA and GDPR.

Several surveys, such as the one by Nguyen et al., have emphasized FL's relevance in smart healthcare. They discuss FL's ability to enable cross-device collaboration while preserving confidentiality. Another review by Kairouz et al. (Google AI) presents a foundational understanding of FL's strengths and limitations, highlighting its scalability, communication bottlenecks, and non-IID data challenges.

Need for Personalization: Beyond One-Size-Fits-All

In real-world applications, especially in healthcare, no two users have the same data. Patients differ in age, health conditions, lifestyle habits, and wearable usage patterns. This makes standard federated models less effective when applied uniformly.

This has led to the emergence of Personalized Federated Learning (PFL) techniques. Unlike global FL models that treat all clients equally, PFL focuses on fine-tuning the model at the client level, adjusting it to reflect individual behavior and physiological signals. Studies like FedPer (Federated Personalization) and PerFedAvg introduce ways to combine global knowledge with userspecific adaptations, resulting in higher accuracy for applications like heart-rate monitoring and fall detection.

The Role of Edge Computing while FL decentralizes model training, it still often depends on the cloud for orchestration. In latency-sensitive healthcare applications, especially where immediate action is needed (e.g., heart attack detection), relying on cloud connectivity is not always ideal. Edge computing solves this by moving computation closer to where data is generated—on devices like smartwatches, smartphones, or local edge servers.

Wu et al. proposed a cloud-edge FL framework where edge devices aggregate updates from wearables and perform intermediate model refinement before syncing with the cloud. This setup reduces delays and improves responsiveness. Other works explore "edge intelligence" strategies, where edge nodes make on-the-fly decisions without needing cloud communication for every update.

Security & Privacy Enhancements even with FL, the risk of indirect data leakage remains. Researchers have shown that it's sometimes possible to infer sensitive information from model updates—a problem especially concerning in health data.

To address this, several frameworks have been developed. For example, FRESH (Federated learning with Ring Signatures and Homomorphic encryption) applies cryptographic techniques to secure both the model and communication pipeline. Other approaches use differential privacy, ensuring that the shared data cannot be traced back to a specific individual.

Smart Healthcare A number of real-world FL applications in healthcare have already been tested or simulated. For instance: FedStack uses personalized FL for activity recognition in smart homes, allowing systems to learn from each user's unique patterns.

Gupta et al. combined FL with digital twin technology to detect anomalies in chronic disease patients, showing how virtual simulations can personalize real-world interventions. Sufian et al. used edge devices to classify ECG data in real-time, achieving both privacy preservation and clinical-grade accuracy. These systems demonstrate the practical feasibility and effectiveness of combining FL with edge-based processing in real healthcare environments.

Emerging Trends: Blockchain and Smart Contracts

More recent studies explore the use of blockchain to add an additional layer of trust and transparency to FL systems. By using smart contracts, researchers have proposed frameworks where access control, model update logging, and participant rewards can be automated and recorded immutably.Blockchain-based FL systems help ensure that all data interactions are traceable, tamper-proof, and auditable.

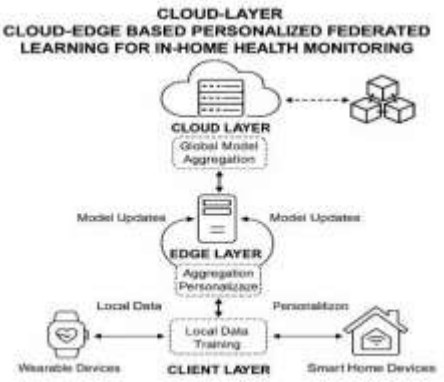

3 METHODOLOGY

The core objective of this study is to develop a privacy-preserving, personalized, and real-time health monitoring system that operates efficiently across smart devices within a patient's home. To achieve this, we propose a three-layer architecture that combines Personalized Federated Learning (PFL), edge computing, and cloud coordination—ensuring both intelligence and privacy at every level of the system.

3.1 System Overview

The system is designed using a three-tier structure:
Client Layer (Wearables & Smart Devices)

This layer includes the devices directly interacting with the user, such as smartwatches, fitness bands, ECG sensors, and other IoMT (Internet of Medical Things) tools. These devices collect raw health data like heart rate, sleep cycles, temperature, or motion activity. Instead of uploading this sensitive data, they perform local training using lightweight machine learning models.

Edge Layer (Local Edge Gateways)

Positioned between the client and the cloud, the edge layer acts as an intelligent gateway. It aggregates local model updates from multiple nearby devices, performs initial encrypted, aggregated updates to the cloud, which uses them to refine the global model.

Model Distribution & Iteration:The improved global model is redistributed to all clients, beginning a new training round. Over time, this iterative process results in highly personalized models across devices.

3.3 Privacy-Preserving Mechanisms

Given the sensitivity of health data, several safeguards are built into the system

Differential Privacy: Ensures that individual user data cannot be reverse-engineered from shared updates.

Secure Aggregation: Only aggregated model updates are visible to the cloud, preventing any single client's data from being isolated.

Blockchain Logging: Each model update and access event is recorded using smart contracts on a blockchain, creating a transparent and tamper-proof audit trail.

3.4 Personalization Strategies

To make the models truly user-specific, we integrate Personalized

Federated Learning (PFL) techniques:

Meta-learning approaches help the model learn how to adapt quickly to new data from each individual.

Client clustering groups similar users (e.g., based on age, activity level, or health conditions) for more relevant model tuning.

Fine-tuning layers are added on-device to further personalize the model output based on the user's latest data.

3.5 Real-World Deployment Scenario

In a typical home setup, an elderly patient might wear a smartwatch that monitors ECG data, step count, and sleep patterns. The device trains a local model using recent data. An edge gateway installed in the home collects updates from multiple wearables (or even multiple family members), refines the model,

and alerts the patient or healthcare provider if anomalies are detected. All of this happens without sending personal data to any external cloud—only secure, encrypted updates are shared for global model improvement.

This methodology ensures that health monitoring becomes more personalization, and can even detect anomalies in real-time personalized, intelligent, secure, and responsive—all while (e.g., sudden heart rate spikes). This level significantly respecting the patient's privacy and data ownership. reduces latency and ensures quicker feedback without

relying on a distant cloud server.

4 TOOLS AND TECHNOLOGY

Cloud Layer (Central Orchestrator)

The cloud layer coordinates the global model aggregation,

To implement and evaluate the proposed cloud-edge-based receives encrypted updates from edge nodes, and personalized federated learning system, a variety of tools and maintains a high-level understanding of overall model technologies are integrated at different levels of the architecture. performance. It handles intensive computations, long-term

These components work together to ensure secure data handling, storage, and periodic model optimization across the entire efficient processing, and accurate personalized predictions. network.

Programming & AI Frameworks

3.2 Personalized Federated Learning Process

Python – The primary programming language used for

developing AI models and integrating system components.

Here's how the learning process works in practical steps:

TensorFlow Federated (TFF) – Enables simulation and

Initial Model Broadcast: The global model is initialized implementation of federated learning algorithms. in the cloud and pushed to all edge nodes and client

devices.

PySyft – A Python library designed for secure and private deep

Local Model Training: Each wearable or smart device learning, supporting federated learning with privacy features like trains the model locally using the user's recent health data. differential privacy and homomorphic encryption. This data never leaves the device.

Scikit-learn & PyTorch – Used for developing and benchmarking

Edge-Level Personalization: Model updates are sent to traditional and personalized models. the edge node, where they are aggregated and fine-tuned

to

reflect local patterns (e.g., older users vs. younger ones,

Edge Computing Platforms resting

heart rate ranges, etc.)

Raspberry Pi 4 / NVIDIA Jetson Nano – Compact, affordable

Cloud Aggregation: The edge nodes periodically send edge devices capable of running lightweight machine learning models and collecting health data from wearable sensors.

EdgeX Foundry – A flexible, open-source edge platform for managing devices, data collection, and processing near the data source.

Cloud Infrastructure

Google Cloud Platform (GCP) / Amazon Web Services (AWS) – Used for global model aggregation, data storage, and orchestration of federated learning cycles. Kubernetes – Container orchestration tool used to deploy and scale federated learning services in the cloud.

Security & Privacy Tools

OpenMined's PyGrid – For secure multiparty computation and privacy-preserving federated learning.

TenSEAL – Enables encrypted computations using homomorphic encryption, ensuring sensitive patient data remains secure even during model updates.

Differential Privacy Libraries – Such as Google's DP Library, for adding statistical noise to model updates and preventing data reconstruction. Health Data Sources & Simulators

MIMIC-III – A publicly available critical care dataset used for training and testing medical AI models.

UCI HAR Dataset – Human Activity Recognition data for monitoring movement patterns in elderly or chronically ill patients.Synthetic IoMT Data Generators – For simulating personalized sensor data when real-time inputs are unavailable.

Ethereum – A decentralized platform for deploying smart contracts to manage access control and participation rewards.
Remix IDE – A browser-based environment to write, test, and deploy Ethereum smart contracts using Solidity.
Hyperledger Fabric – A permissioned blockchain framework for private healthcare networks.

Implementing a cloud-edge-based personalized federated learning system for in-home health monitoring requires an ecosystem of technologies that span from data acquisition at the client end to secure model training and deployment in the cloud. Below is a more detailed view of the tools and technologies involved:

1. **Programming Languages and AI Frameworks**

Python: Widely used for building machine learning models and orchestrating federated learning pipelines. TensorFlow Federated (TFF): An open-source framework that allows simulation and deployment of federated learning models in Python.
PyTorch +PySyft: PyTorch is used for model development, while PySyft adds capabilities for secure and privacy-preserving federated learning using differential privacy and encrypted computation.
Keras: Useful for prototyping deep learning models, especially on edge devices due to its simplicity.

2. **Edge Device Platforms**

Raspberry Pi 4 / Jetson Nano / Coral Dev Board: Low-cost, low-power devices used to simulate real-world wearable or IoT health sensors capable of executing local ML models.
Edge TPU / Intel Neural Compute Stick: Hardware accelerators that allow faster inference on edge devices, reducing power consumption and latency.

3. **Cloud and Backend Infrastructure**

Google Cloud / AWS / Microsoft Azure: These platforms offer scalable computing power, model coordination, and storage for federated learning experiments.
Firebase: For real-time database and lightweight user authentication in mobile health applications.
Docker + Kubernetes: Used to containerize model training jobs and orchestrate them across edge and cloud servers.

4. **Healthcare-Specific Data Repositories**

MIMIC-III / MIMIC-IV: Real-world ICU datasets containing electronic health records.
UCI Human Activity Recognition (HAR) Dataset: Contains sensor data useful for recognizing physical movements of individuals.
PhysioNet: A repository for complex physiological signals such as ECG, EEG, and other biosignals used in cardiology and neurology.

5. **Data Privacy & Security Tools**

Differential Privacy Libraries (Google, IBM, OpenMined): These inject statistical noise to ensure individual data points can't be reverse-engineered.
Homomorphic Encryption (TenSEAL / Microsoft SEAL): Allow computation on encrypted data, preserving privacy during training.
Secure Aggregation Protocols: Used to combine model updates from devices without revealing individual contributions.

6. **Blockchain and Smart Contract Frameworks**

Ethereum / Solidity: Used to create decentralized applications (dApps) that manage health data permissions and incentive distribution.
Truffle Suite + Ganache: Tools for developing and testing smart contracts locally before deploying them to the main blockchain.
IPFS (InterPlanetary File System): For decentralized and secure storage of health data and audit logs.

7. **Visualization & Monitoring Tools**

Grafana + Prometheus: For monitoring edge device performance, system latency, and federated learning progress in real-time.
Power BI / Tableau: To visualize trends and outcomes from model predictions and personalized insights for patients and clinicians.

8. **Mobile Health (mHealth) Integration**

React Native / Flutter: For building cross-platform mobile apps that allow users and doctors to view personalized health analytics.
Bluetooth Low Energy (BLE) SDKs: Used for connecting wearable devices to mobile health platforms.

Mobile Health (mHealth) refers to the use of mobile devices, such as smartphones, tablets, and wearable gadgets, to support medical and public health practices. In the context of in-home health monitoring, mHealth plays a crucial role by enabling realtime interaction between patients and healthcare systems — all from the convenience of a mobile device.

Moreover, mHealth empowers patients to take a more active role in managing their health, fostering better adherence to treatment plans and lifestyle changes. Bridges the gap between advanced AI systems and user-centric healthcare delivery.

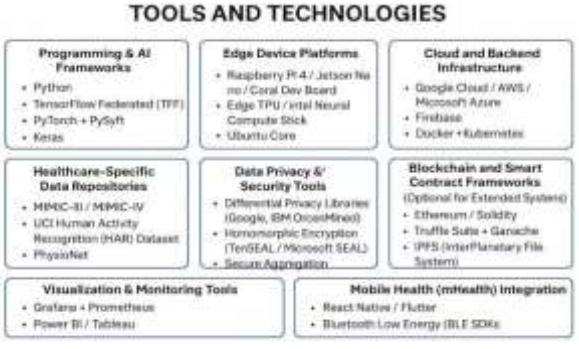

5 SMART CONTRACTS

As healthcare systems increasingly move toward decentralized models for data management and analytics, smart contracts—automated, self-executing agreements stored on a blockchain—

are becoming essential in maintaining trust, security, and transparency. In the context of Personalized Federated Learning (PFL) for inhome health monitoring, smart contracts can play a critical role in governing how data and model updates are exchanged and validated across multiple devices and stakeholders.

1. Role of Smart Contracts in PFL:

Smart contracts can automate and enforce policies related to:Data access control: Only authorized devices or institutions can participate in training or view aggregated models.
Contribution tracking: Devices (like wearables or smartphones) that contribute more to model accuracy can be fairly rewarded, encouraging participation.
Auditability: All actions—such as model training, data exchange, or access attempts—are logged immutably on the blockchain.
Compliance: Smart contracts help enforce rules tied to data protection laws like HIPAA or GDPR by restricting how and when data can be accessed.

2. Use Case Example:

Imagine a scenario where multiple patients wear healthmonitoring devices at home. Each device locally trains a model using patient-specific data. When an update is ready to be submitted to the cloud or edge node, a smart contract is triggered:
It verifies the authenticity of the submitting device. It logs the model update event on the blockchain for transparency.
It executes rewards (like tokens) if the update contributes significantly to the global model.
It ensures privacy policies are respected before accepting the update.

3. Technologies Used:

Blockchain Platform: Ethereum or Hyperledger Fabric
Smart Contract Languages: Solidity (for Ethereum), Chaincode (for Hyperledger)
Integration Tools: Web3.js (for DApp integration), Truffle Suite for testing and deployment.

How Do Smart Contracts Work?

Predefined Logic is Written
Rules are coded into the smart contract — for example: "If a patient device sends a valid model update, reward it with credits."
Deployment on Blockchain
Once written, the smart contract is deployed on a blockchain network like Ethereum. This means: It's immutable (can't be altered after deployment)
It's transparent (anyone can audit the code and results)
Triggering the Contract
When conditions are met (e.g., a wearable device sends a local model update), the contract automatically triggers.
Action is Executed
Depending on the logic, the contract might:
Log the update securely on the blockchain
Grant access to the cloud-edge system
Reward the device/user with a token or incentive
Update model participation records
Trustless Verification
There's no need to trust a third party — the contract handles enforcement automatically and fairly.

Challenges and Limitations

- **Data Heterogeneity:**
  Patient data collected from various devices (smartwatches, ECG monitors, motion sensors, etc.) is often non-uniform and inconsistent. This variability—known as non-IID data—can hinder model convergence and reduce the accuracy of generalized models.
- **Unstable Network Connectivity:**
  In many residential settings, internet reliability may not be ideal. Since federated learning requires frequent communication between clients, edge nodes, and the cloud, unstable connections can delay or disrupt training rounds.
- **Limited Edge Computing Resources:**
  Edge devices such as Raspberry Pi or Jetson Nano are resource-constrained in terms of memory, CPU, and power. Running sophisticated models or encryption mechanisms locally can cause overheating, slow response times, or even failure.
- **Privacy Leakage from Model Updates:** Even though raw data is not shared in FL, model gradients and updates may still leak private information through inference or reconstruction attacks.
- **Communication Overhead:**
  Transmitting model updates, especially large ones, across multiple training rounds introduces significant communication costs. In bandwidth-constrained or remote areas, this can slow down training drastically.

Future Scope of Smart Contracts

1. Automated Privacy Enforcement:

Smart contracts can be designed to automatically enforce data access policies. For example, they can ensure that only authorized edge devices or cloud servers participate in training, strictly based on healthcare compliance standards like HIPAA or GDPR.

2. Incentivization and Fair Contribution:

One of the major challenges in federated learning is the lack of motivation for edge nodes (especially personal devices) to contribute computing power or data. Smart contracts can handle automated incentive distribution, rewarding clients fairly based on how much they contributed to improving the model.

3. Transparent Participation Tracking:

In a federated system, especially one dealing with lifecritical data, transparency and trust are essential. Smart contracts can log each participant's contribution and behavior on a blockchain, creating a tamper-proof, auditable history of learning activity.

4. Self-Regulating Learning Cycles:

Smart contracts can be programmed to automatically trigger model training cycles, distribute tasks, validate model performance, and even halt participation if rules are violated (e.g., poor model accuracy or suspicious behavior).

5. Cross-Institutional Collaboration:

Hospitals, clinics, and remote healthcare providers can collaborate on a global scale without sharing raw data. Smart contracts can act as digital legal agreements that govern how models are trained across multiple entities while preserving data sovereignty.

6. Flowchart Representation:

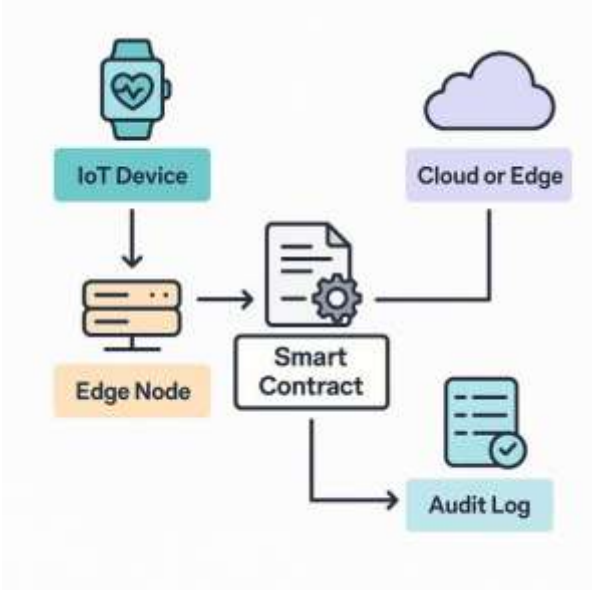

6 CASE STUDY

1.Personalized Federated Learning for Cardiac Monitoring in Elderly Patients:

With the increasing prevalence of heart disease among elderly populations, early detection of abnormal heart rhythms is critical. Traditional health systems often rely on hospital-based monitoring, which can delay interventions and increase costs. To address this, a pilot study was conducted using a Cloud-Edge Personalized Federated Learning (PFL) system for in-home cardiac health monitoring.

Implementation
Data remained on local devices.
Each edge node trained a local model using patientspecific data.
Weekly updates were sent (encrypted) to the cloud server for model aggregation.
The personalized models were then pushed back to each edge node.
Smart contracts logged the update cycles, model accuracy, and access records for transparency.

Key Technologies Used
TensorFlow Federated for distributed learning
Secure Aggregation Protocol to maintain data confidentiality Blockchain Ledger (Ethereum Testnet) for smart contract validation.
MIMIC-III dataset for initial model pre-training.

2.Remote Cardiovascular Health Monitoring Using CloudEdge Federated Learning:

Cardiovascular diseases remain the leading cause of death worldwide. Early detection and continuous monitoring are critical, especially for elderly patients or those living in rural areas. A hospital collaborated with a health-tech company to test an intelligent, privacy-preserving, in-home monitoring system for heart patients using a Cloud-Edge based Personalized Federated Learning (PFL) architecture.

Implementation
Data Collection:Patients wore sensors 24/7. Data such as heart rate, body temperature, and motion were continuously collected and processed at the edge node.
Model Training:A base FL model for arrhythmia detection was sent to edge devices. Each device trained the model on local data, without uploading sensitive data to the cloud.
Personalization: Personalized Federated Learning techniques such as FedPer and FedAvgM were applied to fine-tune each model based on the individual's heart data patterns.
Smart Contracts:Blockchain smart contracts were triggered approve training rounds.

3.Post-Surgery Recovery Monitoring using Cloud-Edge Based Personalized Federated Learning:

Implementation:
Smart wearables and a mobile app collect patient data like movement, pain level, and medication use. An edge device (e.g., Raspberry Pi) processes this data locally and trains a personalized federated learning model. The model is updated without sharing raw data, maintaining privacy. A cloud server aggregates encrypted model updates from many patients to improve the global model. Smart contracts on a blockchain manage access, send alerts, and ensure data integrity. This system enables realtime recovery monitoring, accurate predictions, and secure health data handling.

Technologies Used:
Sensors: Motion and pain-reporting apps
Edge Devices: Localized gateways in homes
Personalized FL Models: Adjusted to individual recovery curves
Smart Contracts: Used for scheduling follow-ups and tracking progress securely

7. CHALLENGES AND LIMITATIONS

Despite the promise of Cloud-Edge based Personalized Federated Learning (PFL) in transforming in-home health monitoring, several technical and practical challenges remain. These limitations must be addressed for widespread adoption and longterm success.

- **Data Heterogeneity:**
  One of the biggest hurdles in personalized FL is the non-identically distributed (non-IID) nature of health data. Patients differ in age, medical history, lifestyle, and device usage. This variation creates inconsistencies in training data across devices, making it difficult to build a universal model that performs equally well for all users.

- **Resource Constraints at the Edge:** Edge devices like smartphones, Raspberry Pi units, or IoT health trackers have limited computational power, storage, and energy capacity. Running FL algorithms, even lightweight

ones, on these devices can be inefficient or unsustainable, especially over long durations or in resource-poor regions. This raises concerns around model size, battery drain, and device wear and tear
.

- **Communication Overhead:**
  FL requires frequent communication between edge devices and the central/cloud server for model synchronization. In environments with poor network connectivity or limited bandwidth (such as rural areas or mobile scenarios), this can lead to delays, dropped updates, and unstable model convergence. Optimizing the number and size of updates remains an active research area.

- **Security and Privacy Risks:** Although FL is more privacy-preserving than centralized learning, it is not immune to attacks: Inference attacks can exploit model updates to reconstruct sensitive user data.
  Byzantine attacks involve malicious participants injecting poisoned updates to compromise the model.
  Secure aggregation and differential privacy can help, but they often reduce model accuracy or increase computation overhead.

- **Scalability Concerns:**
  As the number of participating clients grows, so does the complexity of managing updates, coordinating learning rounds, and maintaining synchronization. This limits how easily the system can scale to handle millions of devices across different health systems, especially when using personalized models.

- **Data Labeling and Quality Issues:**
  In healthcare, acquiring accurate and well-labeled data is often difficult. Many devices generate raw signals (e.g., heart rate, oxygen levels), but without proper clinical context or labeling, their utility is limited. Additionally, noise in sensor data can degrade model performance unless rigorous filtering or preprocessing is applied.

- **Lack of Standardization:**
  There's currently no unified standard for implementing FL in healthcare, especially when combining edge computing and blockchain. Variations in device types, software platforms, and regulatory policies make it challenging to develop interoperable, reusable frameworks.

- **Legal and Ethical Concerns:**
  Healthcare data is highly regulated under frameworks like HIPAA, GDPR, and HL7. Ensuring compliance while using decentralized learning is difficult, particularly when deploying across international borders. Ethical concerns also arise regarding data ownership, consent, and bias in personalized models.

- **Blockchain Complexity:**
  While integrating blockchain adds transparency and trust, it also introduces additional complexity. Managing smart contracts, ensuring consensus, and handling blockchain overhead may not always be feasible in realtime, resource-constrained healthcare applications.

- **Model Personalization Trade-offs:**
  While personalized federated learning tailors the model to individual users, it creates a trade-off between generalizability and specificity. Over-personalization may lead to overfitting, where the model performs well for a specific patient but poorly for new or unseen data, making it less clinically useful across broader populations.

- **Intermittent Device Participation:**
  In FL settings, device availability is unpredictable. Devices may go offline due to power issues, connectivity problems, or user inactivity. This results in client dropout, which affects the learning cycle and can reduce the effectiveness of the global model. Ensuring robustness against client unavailability is still a developing area.

- **Latency in Decision-Making:**
  Although edge computing reduces response times, certain health conditions require ultra-low-latency responses, like seizures or cardiac arrests. In such critical scenarios, even minor delays in model processing or communication can be life-threatening. Current FL systems are not yet optimized for such time-sensitive applications.

- **Usability and Human Factors:**
  Deploying these systems in real homes means the technology must be user-friendly. Many patients—especially the elderly or non-tech-savvy users—may face difficulties operating or maintaining smart devices. Additionally, inconsistent usage patterns (e.g., removing wearables, skipping app updates) can interrupt data flow and model learning.

- **Energy Consumption:**
  Running federated training on mobile or edge devices repeatedly can drain battery life, leading to reduced usage or resistance from users. This is especially problematic for wearables and portable medical devices where energy efficiency is critical.

- **Difficulty in Real-World Validation:** While simulations or lab testing may show promising results, real-world deployments are far more complex. Variability in environment, user behavior, device calibration, and health conditions makes it hard to validate models at scale or across diverse demographics.

- **Maintenance and Upgradability:**
  Keeping all devices up-to-date with the latest models, patches, and smart contract versions is a logistical challenge. Any failure to update devices may result in inconsistent model performance or security vulnerabilities.

- **Lack of Medical Certification:** Most FL and blockchain-based health monitoring systems are still experimental and lack formal approval from bodies like

the FDA or CE. Without regulatory endorsement, clinical adoption remains limited, despite technological readiness.

- **Data Synchronization and Time Drift:** Health data is time-sensitive. In federated settings where devices operate asynchronously, time drift between data sources can corrupt the temporal integrity of the dataset, affecting the quality of real-time predictions or event correlations.

- **Edge-Cloud Coordination Failures:** If the coordination between edge devices and the central cloud server is weak—due to network instability or software errors—the model convergence process can be interrupted, causing learning stalls, inconsistencies, or data loss.

- **Economic Cost:** Deploying and maintaining a cloud-edge federated learning infrastructure with blockchain integration involves significant financial costs—hardware, connectivity, cloud computation, and energy usage. This may limit its use in resourceconstrained settings or developing countries unless subsidized or made open-source.

8. FUTURE SCOPE

As healthcare becomes increasingly digitized, the integration of Personalized Federated Learning (PFL) within cloud-edge frameworks is poised to revolutionize how we monitor and manage patient health at home. Looking ahead, several promising directions can enhance the scalability, intelligence, and reliability of this technology.

- **Integration of AI-Driven Predictive Analytics:** In the near future, federated learning models could be enhanced with advanced AI prediction engines that detect subtle health anomalies before symptoms appear. By leveraging time-series data from wearables and IoT sensors, AI can foresee critical events like heart attacks, seizures, or postsurgical infections, enabling real-time, preventive healthcare.

- **Cross-Device and Cross-Platform Compatibility:** Future systems will support seamless interoperability across various wearable brands, smart home devices, and hospital databases. This standardization will allow healthcare providers to access richer, unified patient profiles and personalize care more effectively, regardless of device vendor or platform.

- **Blockchain Evolution for Smart Contracts:** The smart contracts used in this system today can evolve into self-updating, adaptive contracts that modify permissions and data flows based on real-time patient conditions. Future contracts could also integrate with decentralized identity (DID) systems, allowing patients to retain full control of their health records across providers and platforms.

- **Federated Transfer Learning:** By integrating federated transfer learning, models trained on one patient population or disease condition could be adapted to others with minimal retraining. This would greatly reduce training time and resource consumption, allowing rapid deployment in emergency situations or novel disease outbreaks.

- **Voice and Emotion Recognition Integration:** Future in-home monitoring systems may incorporate emotion detection, speech analysis, and mental health indicators into the federated framework. For example, identifying signs of depression or cognitive decline through voice patterns or facial recognition—while maintaining privacy—will expand the scope of in-home care beyond just physical health.

- **Energy-Efficient Edge Computing:** Next-gen edge devices will likely be equipped with lowpower AI chips (e.g., Google Coral, NVIDIA Jetson Nano) that allow efficient, high-performance model training directly on wearable or home-based health devices. This will further reduce cloud dependency and preserve energy, making the system more sustainable.

- **Clinical Trial and Medical Certification:** Future iterations of this framework could be tested in large-scale clinical trials and eventually gain regulatory approval from health authorities like the FDA, WHO, or national healthcare systems. This would pave the way for mainstream adoption in hospitals, home-care services, and rehabilitation centers.

- **Global Health Application:** This technology has the potential to extend its benefits to remote or underserved regions, where access to hospitals and specialists is limited. Lightweight edge devices combined with FL and blockchain can support low-resource, privacy-aware deployments in rural and developing areas, improving global health equity.

- **Adaptive Learning Models:** Future federated systems will include adaptive models that learn from new patient behaviors and automatically re-personalize themselves without full retraining. This "lifelong learning" capability will ensure that models remain accurate and up-to-date over time.

- **Emergency Response Automation:** In critical situations, like a sudden drop in oxygen level or irregular heartbeat, future frameworks could trigger automated emergency responses—such as calling a doctor, alerting a family member, or dispatching an ambulance—through integrated smart contracts and IoT connectivity.

9. CONCLUSION

This paper has presented a comprehensive and forward-thinking framework for Cloud-Edge Based Personalized Federated Learning (PFL) tailored for in-home health monitoring systems.

With the growing demand for real-time, secure, and personalized healthcare solutions, especially in aging populations and chronic disease management, the proposed model offers a significant leap toward decentralized, privacy-preserving, and intelligent healthcare.
By leveraging edge computing, federated learning, and blockchain-integrated smart contracts, this architecture allows patient data to remain at the source while still contributing to powerful AI models that learn from global patterns.
The integration of Personalized Federated Learning helps overcome data heterogeneity, ensuring that insights generated are both accurate and tailored to individual needs.

Meanwhile, smart contracts bring trust, transparency, and automated access control to the entire data-sharing and model training pipeline.
Through real-world case studies, such as post-surgery recovery and chronic disease monitoring, this paper has demonstrated the practical viability and adaptability of the framework. Despite challenges in energy efficiency, device compatibility, and regulatory compliance, the approach offers a solid foundation for next-generation digital healthcare.

Ultimately, this research opens new pathways for delivering proactive, patient-centric, and privacy-first healthcare at scale, laying the groundwork for future innovations in AI-driven health systems globally. The integration of Cloud-Edge computing, Personalized Federated Learning (PFL), and blockchain-powered smart contracts marks a transformative shift in how we approach in-home healthcare. This academic paper has proposed a scalable and privacy-preserving framework that enables intelligent health monitoring without compromising sensitive patient data.

Unlike traditional centralized systems, our approach ensures that personal health data remains within the patient's control—processed at the edge and refined with local personalization. This decentralization not only mitigates the risks of data breaches and latency but also empowers real-time decision-making at the point of care. The federated learning model collaboratively trains AI systems across edge devices, while personalization adapts these models to reflect the unique health patterns of each individual.

Furthermore, the incorporation of blockchain and smart contracts provides an immutable and transparent system for logging participation, enforcing data access permissions, and incentivizing data contributions. This adds an essential layer of trust and accountability—crucial in sensitive environments like healthcare. Real-world use cases such as post-surgery recovery tracking, chronic disease management, fall detection for the elderly, and mHealth app integration reinforce the practical applicability of this architecture. These implementations showcase how edge-enhanced PFL can be deployed in diverse scenarios to improve health outcomes, reduce hospital readmissions, and empower individuals in managing their own health.

Despite its advantages, this approach comes with a few limitations—such as limited processing power at the edge, communication bottlenecks, and regulatory complexity. However, with ongoing advancements in hardware miniaturization, efficient AI models, and 5G connectivity, these challenges are being progressively mitigated. Looking ahead, this architecture opens exciting opportunities for fully decentralized, intelligent health ecosystems—where wearable devices, smart homes, and healthcare providers interact seamlessly to deliver preventive, personalized, and patientcontrolled healthcare.
In essence, this research not only addresses current pain points in digital healthcare but also lays a strong foundation for the next generation of smart, secure, and personalized health monitoring systems.

The convergence of edge computing, federated learning, and blockchain represents not just a technical evolution, but a paradigm shift in how health data is managed—bringing the focus back to user-centric, privacy-first design in digital health.
With wearable and IoMT devices becoming more sophisticated, the accuracy and richness of real-time health data will only improve, allowing the federated models to become even more predictive and adaptive over time.

The personalization layer in federated learning ensures that healthcare recommendations aren't one-size-fits-all but are tailored to the individual's medical history, environment, and behavior—making interventions smarter and more effective.
As regulatory bodies like HIPAA and GDPR continue to raise the bar for data protection, architectures like the one proposed in this study will become not just beneficial, but necessary for legal and ethical compliance in healthcare AI.

This research also lays the groundwork for future developments in cross-hospital federated collaborations, where anonymized model updates from various medical institutions can collectively improve diagnosis and treatment quality—without compromising data sovereignty.Finally, by aligning with sustainable computing goals, this decentralized architecture has the potential to minimize energy consumption and reduce network load, making it suitable for widespread, long-term deployment even in rural or resourceconstrained regions.

10.REFERENCES

[1] Wu et al. (2020) – FedHome: Cloud-Edge based Personalized Federated Learning for In-Home Health Monitoring, introducing
generative auto-encoder techniques to tackle non-IID data

[2] Nguyen et al. (2022) – Federated Learning in Smart Healthcare: Applications, Challenges, Future Directions (survey)

[3] Al-Muzaffar et al. (2023) – Federated learning-based AI approaches in smart healthcare (taxonomy and open issues)

[4] Sheller et al. (2020) – Federated learning in medicine: multi-institution collaboration without sharing patient data

[5] Kuo & Ohno-Machado (2018) – ModelChain: Decentralized privacy-preserving healthcare predictive modeling using
blockchain

[6] Liang et al. (2023) – Blockchain-Enabled Federated Learning
Platform for Algorithmic Fairness in Predictive Healthcare

[7] Biken Singh & Pratap (2022) – BPFISH: Blockchain and
Privacy-Preserving Federated Learning in Smart Healthcare

[8] Behera et al. (2021/2022) – Federated Learning Using Smart
Contracts on Blockchains, based on Reward Driven Approach

[9] Preuveneers et al. (2021/2022) – Blockchain-based Federated Machine Learning with Proof-of-Training
Consensus & scalable AUC performance

[10] Heiss et al. (2022) – Verifiable Off-chain Computations for Blockchain-Based Federated Learning
(using zero-knowledge proofs)

[11] Karargyris et al. (2023) – Federated benchmarking of
medical AI with MedPerf platform

[12] Chen et al. (2019) – FedHealth: A Federated Transfer Learning Framework for Wearable Healthcare

[13] Rahman et al. (2023) – Federated learning-based AI approaches in smart healthcare: concepts and challenges

[14] Elayan et al. (2022) – Sustainability of Healthcare Data Analysis in IoT using Deep Federated Learning

[15] Liu et al. (2024) – Real-time classification of artifacts
during biosignal acquisition (ECG) in IoMT

[16] Gupta et al. (2021) – Hierarchical Federated Learning-based Anomaly Detection Using Digital Twins in IoMT

[17] Alzubi et al. (2023) – Cloud IIoT-based healthcare record privacy-preserving by CNN and blockchainenabled Federated Learning

[18] Pace et al. (2019) – Edge-based architecture to support efficient healthcare Industry 4.0 applications

[19] Diao et al. (2020/2021) – HeteroFL: Efficient Federated Learning for Heterogeneous Edge Clients

[20] Bonawitz et al. (2019) – Towards Federated Learning
at Scale: system design for decentralized learning

[21] Konečný et al. (2017) – Strategies for Improving Communication Efficiency in Federated Learning

[22] Preuveneers et al. (2020) – DeepChain: Blockchainbased collaborative deep training with incentives

[23] Lei et al. (2023) – BlocFL: Federated Learning with Consortium Blockchain for decentralized model aggregation

[24] Ying et al. (2023) – GCFL: Graph-coordinator FL with DAG blockchain to reduce redundancy and improve scalability

[25] Myrzashova et al. (2023) – Systematic review of blockchain-enabled FL in healthcare

[26] Zai et al. (2023) – Isomerism Learning: Distributed deep learning with dynamic blockchain-enabled FL

[27] Rieke et al. (2020) – The future of digital health with federated learning (npj Digital Medicine)

[28] Dayan et al. (2021) – FL for predicting clinical outcomes in COVID-19 patients (Nature Medicine)

[29] Rodríguez-Barroso et al. (2020) – Federated Learning and Differential Privacy: software tools and guidelines [30] Jatain et al. (Year) – Secure aggregation of private healthcare data via blockchain-based FL