Differential Privacy for Secure Machine Learning in Healthcare IoT-Cloud Systems

N Mangala, Murtaza Rangwala, S Aishwarya, B Eswara Reddy, Rajkumar Buyya, KR Venugopal, SS Iyengar, LM Patnaik

Abstract—Healthcare has become exceptionally sophisticated, as wearables and connected medical devices revolutionize remote patient monitoring, emergency response, medication management, diagnosis, and predictive and prescriptive analytics. Internet of Things and Cloud computing integrated systems (IoT-Cloud) facilitate sensing, automation, and processing for these healthcare applications. While real-time response is crucial for alleviating patient emergencies, protecting patient privacy is paramount in data-driven healthcare. In this paper, we propose a multi-layer IoT, Edge, and Cloud architecture to enhance emergency healthcare response times by distributing tasks based on response criticality and data permanence requirements. We ensure patient privacy through a Differential Privacy framework applied across several machine learning models: K-means, Logistic Regression, Random Forest, and Naive Bayes. We establish a comprehensive threat model identifying three adversary classes and evaluate Laplace, Gaussian, and hybrid noise mechanisms across varying privacy budgets, with supervised algorithms achieving up to 83.6% accuracy. The proposed hybrid Laplace-Gaussian noise mechanism with adaptive budget allocation provides a balanced approach, offering moderate tails and better privacy-utility trade-offs for both low- and high-dimension datasets. At the practical threshold of ε=5.0, supervised algorithms achieve 80–81% accuracy while reducing attribute inference attacks by up to 18% and data reconstruction correlation by 70%. We further enhance security through Blockchain integration, which ensures trusted communication through time-stamping, traceability, and immutability for analytics applications. Edge computing demonstrates 8$\times$ latency reduction for emergency scenarios, validating the hierarchical architecture for time-critical operations.

Index Terms—Differential Privacy, Healthcare IoT, Privacy Protection, Secure Machine Learning.

- *N Mangala is Research Scholar of the Department of Computer Science and Engineering, JNTU Anantapur, AP and Senior Director at C-DAC, India. E-mail: mangala.natampalli@gmail.com*
- *Murtaza Rangwala is a Researcher at the Quantum Cloud Computing and Distributed Systems (qCLOUDS) Laboratory at the University of Melbourne, Australia. E-mail: mrangwala@student.unimelb.edu.au*
- *S Aishwarya is a student of M.E. Computer Science and Engineering, UVCE, Bangalore, India. E-mail: aishwaryas0520@gmail.com*
- *B Eswara Reddy is Professor of the Department of Computer Science and Engineering, JNTU-A, and Director of Research and Development, JNTU Anantapur, AP, India. E-mail: eswar.cse@jntua.ac.in*
- *Rajkumar Buyya is a Redmond Barry Distinguished Professor and Director of the Quantum Cloud Computing and Distributed Systems (qCLOUDS) Laboratory at the University of Melbourne, Australia. E-mail: rbuyya@unimelb.edu.au*
- *KR Venugopal is Former Vice-Chancellor of Bangalore University and Hon. Professor of the Department of Computer Science and Engineering, UVCE, Bangalore, India. E-mail: venugopalkr@gmail.com*
- *SS Iyengar is a Distinguished Professor and Director of Computer Science, Florida International University, Miami, USA. E-mail: iyengar@cs.fiu.edu*
- *LM Patnaik is Adjunct Professor and NASI Senior Scientist in the National Institute of Advanced Studies, Bangalore, India. E-mail: lalitblr@gmail.com*

1 Introduction

The Internet of Things (IoT) has become an integral part of the modern application ecosystem, transforming not only organizational and industrial operations but also our daily lives. In the healthcare sector, digitalization has revolutionised patient care through enhanced data collection, personalised medicine, and preventative treatment approaches. Wearable devices such as Fitbit trackers, smart glasses, and insulin pumps now facilitate continuous monitoring of vital parameters including heart rate, blood sugar levels, and sleep patterns. Users must configure these Healthcare IoT devices with personal information such as age, gender, and location. Healthcare Analytics leverages this patient data to generate descriptive, diagnostic, predictive, and prescriptive insights, thereby supporting evidence-based medical decision-making. Moreover, through AI/ML applications, governmental health agencies and research institutions are developing socially beneficial applications such as epidemiological forecasting, disease progression modeling, healthcare resource allocation, and actuarial risk assessment.

Medical data from IoT devices, imaging technologies (MRI, CT scans), test results, Electronic Health Records (EHR), medication monitoring, and environmental sensors are stored in Cloud systems, creating comprehensive repositories for medical history, diagnosis, research, and analytics. This cloud infrastructure provides secure access to various stakeholders such as physicians, clinicians, paramedics, pharmaceutical companies, insurance firms, researchers, and analysts, enabling them to examine and analyse critical healthcare information [1]. However, this centralised data storage presents significant security vulnerabilities. Adversaries may exploit patient information for blackmail, ransom demands, manipulation, creating artificial market demand, and targeted marketing campaigns [2]. Therefore, protecting patient data is critically important, as misuse can lead to numerous harms including psychological distress, compromised treatment through data corruption, and various forms of extortion.

Healthcare expenditure in the USA reached nearly 20% of GDP in 2021, yet system security remains problematic, with 692 large healthcare data breaches reported between July 2021 and June 2022. To address these vulnerabilities, healthcare organizations must comply with regulatory frameworks such as the Health Insurance Portability and Accountability Act (HIPAA), established by the U.S. Congress in 1996 [3]. HIPAA encompasses three primary

components: the Security Rule, the Privacy Rule, and the Breach Notification Rule. These regulations mandate essential protective measures for Electronic Health Record systems, including access control, data encryption, and comprehensive audit trails [4]. The HIPAA Privacy Rule, in particular, establishes national standards governing the security of health information during transfer, reception, processing, and sharing [5]. Beyond the United States, the European Union's General Data Protection Regulation (GDPR) establishes a comprehensive legal framework for the protection of personal data, including health records, applicable to any organisation processing data of EU residents [6]. GDPR classifies health data as a special category requiring explicit consent and heightened safeguards, mandates data protection by design and by default, and imposes strict accountability obligations including Data Protection Impact Assessments for high-risk processing activities. Together, HIPAA and GDPR represent the two principal regulatory regimes governing healthcare data privacy internationally, and any framework intended for cross-institutional or multi-jurisdictional deployment must satisfy the compliance requirements of both. Despite these regulatory safeguards, persistent data breaches demonstrate the need for more robust privacy-preserving techniques. Differential Privacy (DP) has emerged as a particularly effective approach to securing personally identifiable information, offering advantages over conventional methods through its mathematical framework that anonymises data by introducing calibrated noise to datasets [7].

Motivation

Medical records contain various personal details of patients and are stored in cloud systems for use by multiple stakeholders. Suppressing explicit parameters such as name and address does not guarantee patient privacy because combinations of other fields such as postcode and date of birth can uniquely re-identify individuals.

A significant case illustrating privacy vulnerabilities in healthcare data occurred when Sweeney [8] demonstrated how medical records could be re-identified despite anonymization. She successfully re-identified the medical records of the Governor of Massachusetts by executing merely three queries on an anonymized hospital dataset released by the Massachusetts Group Insurance Commission (GIC). Cross-referencing public information—news reports of the governor's collapse on 18 May 1996 and subsequent hospitalization—with voter registration records containing demographic identifiers (names, addresses, postcodes, birth dates, and gender) facilitated the re-identification. This case exemplifies how quasi-identifiers can be matched across datasets to compromise patient privacy. In contemporary healthcare systems, Machine Learning (ML) techniques are routinely employed to analyse medical data and address research questions [9], yet these analytical capabilities simultaneously present heightened privacy risks if not properly safeguarded. By scripting fine-tuned sequences of ML queries, adversaries can steal personal information of patients from healthcare databases. DP can help safeguard against such privacy issues.

The *Differential Privacy guarantee* states that adding or removing a single row from a database does not significantly alter the output of the analysis, preventing an adversary from determining whether a particular person is present in the database or not. Consider, for instance, a health department investigating the causal relationship between smoking and cancer incidence. Potential participants may exhibit reluctance to engage in such research due to concerns regarding social stigma, legal implications, and possible insurance premium discrimination if their identities become traceable within the dataset. In this scenario, robust data privacy protections are essential to facilitate accurate analysis of the smoking-cancer relationship while ensuring participant anonymity. DP offers a methodological framework that effectively balances individual privacy preservation with data utility for research purposes. Although alternative approaches such as homomorphic encryption [10] and multi-party secure computing [11] enable operations on protected data while maintaining confidentiality, these methods frequently present substantial computational complexity challenges [12]. The principal advantage of DP lies in its capacity to maintain data utility for statistical analyses while specifically obfuscating individual-level identifying information.

Key Contributions

This research enhances response time and security in healthcare data systems, with particular emphasis on protecting patient privacy. We employ DP techniques to safeguard patient data while maintaining its utility for analysis. While prior works have explored DP-ML integration [13], [14], blockchain-based distributed ML [15], and blockchain-based healthcare systems [3], [16]–[18] independently, our work advances the state-of-the-art through three key innovations that distinguish it from existing approaches.

First, unlike existing hybrid DP mechanisms that typically combine noise distributions in fixed proportions or apply them sequentially to the same data, our adaptive hybrid noise mechanism presents guidance and insights on dynamic selection of noise distributions based on the sensitivity characteristics of individual data features and the specific ML algorithm employed. This feature-aware approach exploits the complementary strengths of Laplace noise (superior performance for low-dimensional, sparse features) and Gaussian noise (optimal for high-dimensional, dense features with composition properties), resulting in improved privacy-utility trade-offs across heterogeneous healthcare datasets compared to uniform noise application strategies.

Second, we address a critical gap in the literature by providing the first systematic empirical evaluation of DP mechanisms across multiple ML algorithms (K-Means, Logistic Regression, Random Forest, Naive Bayes) specifically within a multi-layer IoT-Edge-Cloud architecture. Existing DP-ML frameworks focus primarily on deep learning models in centralized cloud environments. In contrast, our work characterizes how the computational constraints and latency requirements of edge layers influence the practical achievability of privacy guarantees, providing actionable guidance for deploying differentially private ML at various architectural tiers.

Third, our integrated architecture uniquely combines three orthogonal security mechanisms: DP for input privacy, blockchain for data integrity and provenance, and hierarchical access control across computing layers, into a unified framework optimized for healthcare emergency response scenarios. While prior work has explored pairwise combinations

of these security mechanisms, our approach formally analyses the security guarantees of the complete integrated system under a comprehensive threat model encompassing inference attacks, data tampering, and consensus attacks. This holistic security analysis demonstrates that the three mechanisms provide complementary protections without introducing conflicting requirements or degrading system performance.

The specific contributions of this work include:

- Design and implementation of a multi-layered IoT-Edge-Cloud architecture to improve response time for healthcare applications.
- Development of a hybrid privacy protection mechanism for medical data in ML applications through the implementation of DP using a novel combined Laplace-Gaussian noise approach that provides insights on dynamic selection of noise distributions based on data characteristics.
- Comprehensive evaluation of the privacy-utility trade-off for different noise distribution types across four distinct ML algorithms, providing algorithm-specific recommendations for privacy budget allocation.
- Enhancement of data integrity, tamper resistance, trust, and security for healthcare applications through integration with Blockchain technology.

Organization

The rest of the paper is organized as follows. An overview of ML, DP and the state-of-the-art Computing Platforms is provided in Section 2. A comparison of the latest literature is presented in Section 3 followed by the problem statement in Section 4. Section 5 establishes a comprehensive threat model defining adversarial capabilities and attack surfaces. The details of the proposed solution are explained in Section 6. Implementation nuances are presented in Section 7. The experimentation and analysis of results are presented in Section 8, threats to validity are discussed in Section 9, followed by conclusions and future directions in Section 10.

2 Preliminaries

2.1 Machine Learning

ML is a branch of artificial intelligence that focuses on creating systems that learn patterns from data and improve automatically with experience. In healthcare, ML supports diverse applications including disease identification, predictive diagnostics, image analysis, epidemic control, treatment optimization, genomic analysis, and surgical automation [9]. These techniques enable researchers to analyse population-level data such as demographic patterns, symptom clusters, and disease correlations, classifying data based on features to derive insights for diagnosis, prediction, and analysis of diseases.

ML approaches can be categorised into two main types: supervised and unsupervised learning. Supervised learning utilises labelled data, where input features are paired with corresponding outputs. The algorithm learns to predict outputs for new inputs by recognising patterns from training data. This approach is commonly employed for classification, regression, and object detection tasks. Key supervised algorithms include Logistic Regression [19], Random Forests [20], and Naive Bayes [21]. In regression problems, the output variable represents a continuous value (e.g., 'pounds' or 'kilogrammes'), while classification problems involve categorical outputs (e.g., 'disease' or 'no disease'). Unsupervised learning, conversely, operates on unlabelled data without predefined outputs. These algorithms identify inherent patterns and relationships within datasets, making them suitable for clustering, dimensionality reduction, and anomaly detection. K-Means [22] represents a prominent unsupervised learning technique used for grouping similar data points into clusters based on their characteristics.

This work evaluates one unsupervised algorithm (K-Means) and three supervised methods (Random Forest, Logistic Regression, Naive Bayes).

(i) K-means is an unsupervised clustering algorithm that partitions data into K distinct clusters by iteratively minimising the within-cluster sum-of-squares. The algorithm assigns data points to the nearest centroid and subsequently updates these centroids until convergence is achieved. In healthcare applications, K-means facilitates patient stratification by identifying cohorts with similar clinical characteristics, thereby enabling more targeted therapeutic interventions. This approach has demonstrated efficacy in medical image segmentation, particularly for tumour delineation, and in identifying latent patterns within complex clinical datasets that may elude conventional analytical methods.

(ii) Logistic Regression is a supervised learning algorithm that models the probability of a binary outcome through the logistic function, which maps the linear combination of predictors to a probability range $[0, 1]$. This parametric approach is particularly valuable for medical diagnostics and prognostics, where quantifying the probability of disease occurrence is essential. The algorithm has been extensively applied in predicting the presence of chronic conditions such as diabetes, cardiovascular risk assessment, and clinical decision support. Its interpretability offers significant advantages in healthcare contexts where transparency in predictive modelling is paramount.

(iii) Random Forest represents an ensemble learning technique that constructs multiple decision trees during training and aggregates their predictions, typically through majority voting. This approach mitigates overfitting and enhances generalization, rendering it particularly suitable for both binary and multi-class classification challenges in medical contexts. Random Forest algorithms have demonstrated considerable utility in analysing complex medical datasets, including radiological images and electrocardiographic data, predicting disease progression, and identifying optimal therapeutic compounds based on molecular characteristics.

(iv) Naive Bayes classifiers apply Bayes' theorem with the naive assumption of conditional independence among features given the class label. This probabilistic approach is computationally efficient and performs remarkably well in various text classification tasks despite its simplifying assumptions. In medical contexts, Naive Bayes algorithms effectively process symptom-based diagnostic classification and medical document categorization. They have shown particular efficacy in diabetes detection when applied to clinical datasets, offering a balance between computational simplicity and predictive performance.

2.2 Differential Privacy

DP is a data sequestration technique that aims to safeguard individual privacy while enabling analysis of large datasets. This mathematical approach introduces controlled noise to data, making the outputs of queries differing by at most one record indistinguishable, thereby ensuring individual privacy [23]. The DP guarantee for two datasets is given by:

$$\Pr[\mathrm{M(D)} \in \mathrm{S}] \leq e^{\varepsilon} \Pr[\mathrm{M(D')} \in \mathrm{S}]$$

where:

- M: Randomised algorithm (i.e., $query(db) + noise$ or $query(db + noise)$)
- S: All potential outputs of M that could be predicted
- D: Entries in the database
- D': Entries in adjacent database
- ε : Privacy parameter that bounds the ratio of probabilities of obtaining the same output when the mechanism is run on adjacent databases

The databases D and D' differ by at most a single record, and the mechanism is differentially private if the results of $M(D)$ and $M(D')$ are almost indistinguishable for every choice of D and D'. The ε parameter quantifies the privacy guarantee. Rényi divergence [24] is used in variants of DP to measure the difference between the probability distributions of the mechanism's outputs when applied to adjacent datasets.

2.2.1 Types of DP

DP can be classified along two primary dimensions. First, based on implementation architecture, into Centralised DP and Local DP. Second, based on methodology, into Pure DP, Query-Level DP, and Moment DP [25].

(i) Centralised DP applies noise to a centralised database maintained by a trusted curator, protecting the privacy of the entire dataset while enabling useful analysis. The trusted curator manages both the raw data and the privacy-preserving mechanisms.

(ii) Local DP introduces noise directly to individual data points before their collection or aggregation, ensuring that sensitive information is protected at its source. This approach eliminates the need for a trusted curator as privacy guarantees are applied locally.

(iii) Pure DP provides mathematical guarantees that analysis results cannot reveal information about any individual record, regardless of the type of analysis performed or background knowledge available to potential adversaries.

(iv) Query-level DP applies privacy protections to specific queries rather than the entire dataset, allowing for fine-grained privacy budgeting based on query sensitivity and importance.

(v) Moment DP focuses on protecting statistical moments of the data distribution (e.g., mean, variance) by adding calibrated noise to these summary statistics rather than to raw data points, often resulting in improved utility for statistical analyses.

2.2.2 Noise Types in DP

DP achieves data protection by introducing calibrated noise to dataset elements or query results, preserving individual privacy while maintaining analytical utility. Various noise mechanisms can be employed, each with distinct statistical properties and application domains:

(i) Laplace Noise Mechanism applies noise drawn from the Laplace distribution, characterised by a symmetric probability density function centered at zero with exponentially decaying tails. This mechanism is widely implemented due to its mathematical tractability and strong theoretical guarantees for bounded-sensitivity functions, particularly in contexts requiring ε-DP with no additional parameters.

(ii) Gaussian Noise Mechanism introduces noise sampled from the Gaussian (normal) distribution, yielding the bell-shaped probability curve familiar in statistical applications. This approach provides (ε, δ)-DP and typically demonstrates greater resilience to outliers than Laplace noise. The Gaussian mechanism is particularly effective for high-dimensional data and functions with $L2$-sensitivity constraints.

(iii) Exponential Noise Mechanism extends DP beyond numerical queries to selection problems by sampling outputs with probability exponentially proportional to their utility scores. This mechanism facilitates privacy-preserving selection from discrete sets and is widely employed for non-numerical applications such as identifying maxima or selecting optimal elements.

(iv) Poisson Noise Mechanism incorporates noise from the Poisson distribution, which models discrete count events. This approach is particularly well-suited for count queries, contingency tables, and histogram analyses where integral outputs are required. Its discrete nature makes it appropriate for applications involving event frequencies or integer-valued datasets.

(v) Discrete Noise Mechanism applies specially designed noise distributions to categorical data or discretised numerical values. This approach preserves the discrete structure of the underlying data while providing DP guarantees, making it particularly valuable for privacy-preserving analysis of categorical attributes and binned numerical data.

2.2.3 Underlying Concepts of DP

(i) Privacy Parameters: DP establishes quantifiable boundaries regarding how much information about an individual's presence in a database may be disclosed to external parties. The parameters ε and δ define these boundaries, characterising the privacy guarantees offered by a randomised privacy-preserving algorithm (M) applied to a specific database (D).

- *Privacy Budget/Privacy Loss (ε)*: The ε parameter quantifies the strength of privacy protection. A smaller ε value corresponds to stronger privacy guarantees, whilst larger values indicate greater privacy loss and potentially more useful but less protected data.
- *Probability to Fail/Probability of Error (δ)*: The delta parameter accounts for the probability of privacy protection failure, specifically, the likelihood that a query might reveal an individual's presence in the dataset. Such failures may occur $\delta \times n$ times, where n represents the number of records [25].

(ii) Centralised Versus Local DP: As discussed in Section 2.2.1, two principal implementation strategies exist for DP. Centralized DP employs a trusted curator who applies precisely

calibrated noise to query results. This approach typically utilises Laplace or Gaussian noise mechanisms and is commonly referred to as the trusted curator model. Conversely, local DP operates without requiring a trusted intermediary, hence its designation as the untrusted curator model. In local DP implementations, data undergoes randomization before curator access. A trusted entity may also employ local DP to simultaneously randomise all database records. Local DP algorithms frequently produce more heavily perturbed data, as noise is applied at the individual record level. This approach provides a particularly rigorous privacy standard with plausible deniability guarantees, establishing local DP as a leading methodology for privacy-preserving data collection and distribution, but at the cost of model utility.

(iii) Correlated Sensitivity: Whilst Global Sensitivity effectively measures the maximum number of correlated records, it fails to account for the degree of data correlation. The concept of Correlated Sensitivity addresses this limitation by quantifying the cascading impact on related records when a single record undergoes modification [26]. This refinement proves particularly valuable when analysing datasets containing interdependent entries, a common characteristic of complex relational databases.

2.2.4 Laplace Noise in Differential Privacy

The Laplace distribution's probability density function (PDF) is defined by the location parameter (μ) and the scale parameter (b):

$$f(x|\mu, b) = \frac{1}{2b} \exp\left(-\frac{|x-\mu|}{b}\right)$$

where:

- x is a random variable.
- μ represents the center of the distribution.
- b controls the width of the distribution.

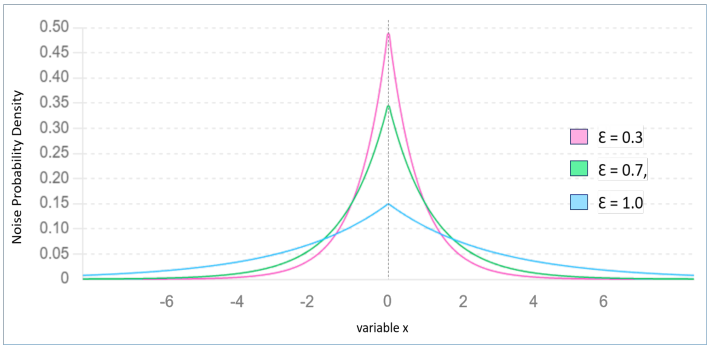

Fig. 1: Laplace Noise Distribution for Different ε Values

In DP applications, this mechanism injects calibrated noise drawn from this distribution into query results. The noise calibration depends on the query sensitivity Δf and the privacy budget parameter ε, as illustrated in Figure 1.

- *Sensitivity (Δf)*: The sensitivity parameter quantifies the maximum possible change in a query's output when a single record is either added to or removed from the dataset. For count queries with binary attributes, the sensitivity typically equals one, as an individual record can influence the count by at most one unit.
- *Privacy Parameter (ε)*: The privacy budget parameter establishes the privacy-utility trade-off, with smaller values providing stronger privacy guarantees at the expense of analytical precision.

The Laplace mechanism determines the scale parameter (b) according to:

$$b = \frac{\Delta f}{\varepsilon}$$

This formulation ensures that the noise magnitude is proportional to the query sensitivity and inversely proportional to the privacy budget, thus maintaining mathematical guarantees of ε-DP.

2.2.5 Gaussian Noise in Differential Privacy

The Gaussian distribution's PDF is defined by the mean (μ) and standard deviation (σ):

$$f(x|\mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{|x-\mu|^2}{2\sigma^2}\right)$$

where:

- x is a random variable for which the probability density is calculated.
- μ represents the center of the distribution.
- σ is the standard deviation of the distribution, which controls the spread or width of the distribution.

In DP implementations, the Gaussian mechanism provides (ε, δ)-DP, where δ represents a small probability of privacy failure. The mechanism is (ε, δ)-differentially private when the standard deviation σ of the Gaussian noise satisfies:

$$\sigma = \sqrt{2log(1.25/\delta)}\frac{\Delta_2 f}{\varepsilon}$$

where ε is the privacy parameter and δ is the probability of failure.

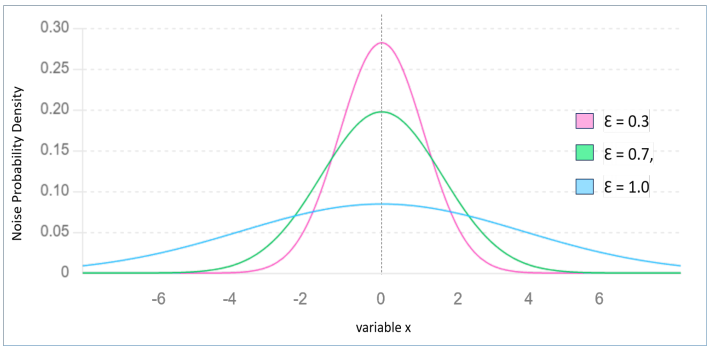

Fig. 2: Gaussian Noise Distribution for Different ε Values

As illustrated in Figure 2, the Gaussian distribution produces the characteristic bell-shaped curve, with its noise distribution varying based on the privacy parameter ε.

- *Sensitivity (Δf)*: The Gaussian mechanism utilises $L2$ sensitivity (squared sensitivity), measuring the impact of changing a single data point through Euclidean distance. This approach is particularly effective for high-dimensional data analysis, mean calculations, and ML model training. For instance, in a dataset of n entries ranging from 0 to 1, the $L2$ sensitivity equals $1/n$, as

adding or removing one value alters the mean by at most $1/n$.

- *Privacy Parameter (ε)*: As in the previous section, the privacy budget parameter establishes the trade-off between privacy protection and data utility.

2.2.6 Applications of DP

Table 1 presents various domains that benefit from DP guarantees for statistical analysis and applications while preserving individual privacy. Key application areas include:

(i) Healthcare: DP enables secure sharing of electronic health records and clinical trial data, allowing researchers to extract insights whilst maintaining patient confidentiality [27], [28].

(ii) Financial Services: Financial institutions employ DP for fraud detection, risk assessment, and customer behaviour modelling. This protects individuals' financial information from exposure that could lead to identity theft or fraudulent activities [29].

(iii) Social Media and Advertising: Platforms implement DP to analyse user behaviour for personalised experiences. Google has applied this approach in its advertising platforms to create user profiles without exposing individual behaviours, addressing growing privacy concerns in digital advertising [30], [31].

(iv) Transportation: DP extracts insights from journey patterns and traffic flow data without revealing individual travel habits. This protection is essential as location data can reveal sensitive information about individuals' routines and behaviours [32].

(v) Research and Academia: DP supports research by enabling examination of sensitive datasets while preserving participant anonymity. This balance is crucial for maintaining ethical standards in research whilst facilitating valuable data-driven discoveries [33].

(vi) Government and Policy-making: DP facilitates evidence-based policy development using large-scale datasets. This protection is necessary as government data often contains comprehensive citizen information that, if compromised, could lead to significant privacy violations [34].

(vii) Web Browsing Analysis: Google implemented DP techniques through their RAPPOR framework for Chrome browsing data to identify performance issues such as slow-loading web pages. This approach enables system improvements whilst protecting individual browsing patterns from surveillance or profiling [35].

(viii) Location-Based Services: Google Maps utilises DP to provide real-time traffic updates and location-based services from millions of users' data. This implementation preserves user anonymity whilst delivering accurate collective insights about traffic conditions and business popularity [36].

(ix) Search Query Analysis: Google's search engine employs DP to enhance recommendations by analysing queries and user interactions. This enables improved search results without compromising individual search histories, which may contain sensitive personal information [31].

(x) Public Health Monitoring: During the COVID-19 pandemic, Google released aggregated mobility reports using DP to protect individual identities. These reports assisted health authorities in understanding social distancing patterns whilst maintaining location privacy of individual users [37].

2.3 Blockchain

Blockchain technology represents a distributed, decentralised ledger architecture designed to record transactions across multiple computing nodes. The architecture comprises multiple interconnected blocks, with each block incorporating cryptographic references to its predecessor. This structure ensures that records cannot be altered or manipulated without modifying all subsequent blocks in the chain. Once data is entered onto the blockchain, it becomes immutable. This immutability confers strong resistance to tampering attempts by malicious actors [28].

Several key features make blockchain particularly valuable for healthcare applications:

- *Immutability:* Patient data, once stored on the blockchain, cannot be altered or tampered with, ensuring the integrity of medical records.
- *Time-stamping:* All transactions are chronologically recorded with precise timestamps, providing an auditable history of healthcare events.
- *Traceability:* Patient-related data, including hospitalizations and treatments, can be traced and categorized by geographical location and temporal parameters.
- *Transparency:* All transactions are transparently viewable by authorised stakeholders within the healthcare ecosystem, facilitating improved coordination of care.
- *Automation:* Smart contracts, which are self-executing protocols stored on the blockchain, automatically implement predefined actions when specific conditions are met without requiring intermediary involvement. These programmable agreements enable automated healthcare workflows such as insurance claim processing, clinical trial management, and medication supply chain verification.

These features collectively enhance the reliability and security of medical data systems [38]. Blockchain's immutable architecture provides an effective framework for regulatory compliance with health authority mandates while enabling secure information exchange among healthcare professionals. This technology is particularly valuable in healthcare contexts characterised by distributed trust requirements, where multiple stakeholders must verify data integrity and provenance without centralised authority. The resultant interoperability across disparate systems addresses a significant challenge in contemporary healthcare information management.

2.4 Different layers of Computing Architecture

The evolution of computing paradigms has facilitated a hierarchical approach to processing healthcare data, addressing various requirements including latency, resource constraints and data volume. This section examines the distinctive characteristics of each architectural layer and their specific healthcare applications.

(i) *Edge Computing*: Edge nodes provide computational offloading, storage and caching capabilities for IoT management, facilitating reduced computation requirements and rapid response times for time-sensitive applications [39]. In healthcare contexts,

TABLE 1: Applications of Differential Privacy

Application	Description	Factors to Consider
Data Analysis	Aggregating and analyzing sensitive data	- Privacy budget (ε) for the desired level of privacy - Noise mechanism (e.g., Laplace, Gaussian) based on data distribution - Data sensitivity and scale - Query types and frequency
Recommendation Systems	Personalized recommendations	- Protecting user preferences while providing personalized suggestions - Trade-off between utility and privacy - Ensuring diversity in recommendations
Healthcare Research	Medical data analysis and research	- Compliance with healthcare regulations (e.g., HIPAA) - Preserving patient privacy while conducting studies - Anonymization techniques for sharing medical records
Location Privacy	Protecting user location data	- Adding noise to location data for anonymity - Balancing location accuracy with privacy guarantees - Considering adversarial attacks on location data
Social Media Analysis	Analyzing social media posts and trends	- Anonymizing user data to protect identity - Maintaining the ability to detect trends and sentiment - Privacy implications of social network analysis
Census Data	Collecting and sharing demographic information	- Protecting individuals' privacy in census data - Balancing accuracy of demographic data with privacy guarantees - Adherence to legal and ethical guidelines
Smart Cities	Analyzing data from IoT sensors for urban planning	- Anonymizing sensor data to protect privacy - Ensuring data integrity and security - Handling diverse data sources and formats - Public awareness and consent
Finance and Banking	Protecting financial transaction data	- Compliance with financial regulations (e.g., GDPR) - Ensuring transaction privacy - Detecting fraud while preserving customer privacy - Secure data sharing mechanisms
Educational Research	Analyzing student performance and learning outcomes	- Balancing the need for research with student privacy - Consent and data sharing agreements - Anonymizing student data to prevent identification - Ethical considerations
Transportation	Analyzing traffic patterns and congestion	- Protecting user privacy while analyzing traffic data - Noise addition to GPS data for anonymity - Data aggregation techniques for traffic analysis - Privacy-aware routing algorithms

edge computing serves as an intermediary processing layer, managing data that requires swift analysis but can tolerate minimal processing delays. A remote healthcare facility might deploy edge servers to process and store patient vital signs such as blood pressure, oxygen saturation and cardiac rhythms. Medical practitioners can examine this locally processed data to make timely clinical decisions, while selectively determining which information warrants transmission to cloud infrastructure for comprehensive analysis and long-term storage.

(ii) *Cloud Computing*: Cloud computing provides on-demand network access to a shared pool of configurable computing resources, including servers, storage, applications and databases [40]. This model centralises data processing and storage on remote server infrastructure, offering significant computational capacity and storage volumes at the expense of increased latency compared to edge-oriented paradigms. While not optimised for time-critical healthcare functions, cloud computing excels in scenarios requiring substantial resources for complex analysis or extensive data retention. Healthcare organizations leverage cloud platforms to maintain comprehensive historical patient records, conduct longitudinal research and perform sophisticated analytical operations. Cloud-based solutions securely store EHRs from multiple healthcare institutions, enabling researchers and clinicians to identify long-term health trends, population-level insights and evidence-based treatment protocols through analysis of aggregated datasets.

3 Literature Review

Recent research has explored various approaches to implementing DP in IoT-Cloud systems. Sun et al. [41] examined the relationship between inference and data privacy, comparing DP implementations using Bayes probability and Gauss-Seidel methods for enhanced data privacy, scalability and reduced network latency. Their work highlighted the need for further investigation regarding the impact of communication delays and packet losses on accuracy and privacy in IoT environments.

While addressing communication challenges in IoT networks, researchers have explored blockchain integration with DP mechanisms, despite blockchain's own consensus overhead contributing to communication delays. Jia et al. [13] implemented a blockchain-enabled federated data protection aggregation scheme with DP, utilising K-means clustering with Adaptive Boosting and Laplacian noise to improve security across multiple datasets. Their research indicated that homomorphic encryption techniques could further enhance secure data exchange in Industrial IoT (IIoT) contexts. Similarly, Gai et al. [14] proposed a blockchain-based Internet of Edge model incorporating Q-Learning algorithms with Laplace noise distribution to implement tamper-resistant DP, though they noted high energy consumption as a limitation requiring future work.

Several studies have investigated noise distribution mechanisms for DP. Hu et al. [42] implemented Federated Learning DP using Gaussian methods, demonstrating that heterogeneous noise perturbation achieves robust accuracy suitable for real-world scenarios. Complementing this work, Cai et al. [43] developed a multimodal DP framework for local DP that demonstrated the flexibility of Gaussian noise in balancing privacy protection with high data utility. Chowdhury et al. [44] described a DP framework utilising a noisy max algo-

rithm with Laplace mechanism that provided accuracy guarantees, suggesting potential for practical application with extensive empirical evaluation on real-world datasets.

Foundational work by Dwork [45] presented various privacy-preserving data analysis techniques established since the early 2000s, providing systematic approaches to achieving privacy protection. Taking a different architectural approach, Bi et al. [46] designed a privacy-isolation zone where sensitive user information is removed from data collected at the IoT endpoints before transmission to cloud systems for analytics.

Table 2 summarises notable recent efforts in guaranteeing privacy protection for sensitive datasets, highlighting the diversity of approaches and their respective strengths in addressing the privacy-utility trade-off in IoT-Cloud systems.

4 Problem Statement

Medical records containing patient histories, diagnoses, treatments, and health-related information represent sensitive data requiring careful management. The central challenge in healthcare IoT-Cloud systems is balancing immediate emergency response capabilities with secure data handling that enables valuable research without compromising patient privacy. This research addresses the tension between processing IoT data with minimal latency for emergency response while simultaneously ensuring secure transmission of this data to cloud storage for long-term preservation, diagnosis, and analytics.

Objectives:
(i) Enhance the response speed of IoT-Cloud healthcare systems for emergency scenarios
(ii) Provide robust privacy guarantees for patient data while maintaining research utility
(iii) Ensure reliable and secure storage of healthcare transaction data in cloud environments

Constraints:
The fundamental constraint governing this research is the necessity to maintain a balance between robust patient privacy protection and preserving sufficient data utility for scientific inquiry. This represents the classical privacy-utility trade-off that requires careful optimization within the proposed framework to ensure both ethical data handling and meaningful analytical capabilities.

5 Threat Model and Security Analysis

This section establishes a comprehensive threat model for healthcare IoT-Cloud systems, defining adversary capabilities, attack surfaces, and security objectives that our proposed solution must address.

5.1 System Model and Trust Assumptions

Healthcare IoT-Cloud systems comprise three hierarchical computational tiers: IoT devices, Edge computing infrastructure and Cloud data centers as shown in Figure 3. Each tier presents distinct trust characteristics and security vulnerabilities. IoT devices including patient wearables and medical sensors operate under direct patient or healthcare provider control, typically with physical security measures limiting unauthorized access. However, these devices often possess limited computational resources for implementing sophisticated security protocols.

Edge computing nodes occupy intermediate positions in the architectural hierarchy, processing data from multiple IoT sources while forwarding aggregated information to cloud infrastructure. These nodes may be operated by third-party service providers or healthcare organizations, introducing varied trust assumptions. The computational resources available at these layers enable more sophisticated processing but also present larger attack surfaces for potential adversaries. Cloud infrastructure, while offering substantial computational and storage capabilities, operates under the control of external service providers and may be accessible to multiple stakeholders including researchers, pharmaceutical companies, and insurance providers, each with potentially conflicting interests regarding data access and privacy.

In distributed ledger implementations for healthcare data integrity, we must consider the trust model of consensus participants. Permissioned blockchain networks restrict participation to authorized entities, whereas public blockchains allow arbitrary participation. The security guarantees of these systems depend fundamentally on assumptions about the proportion of honest versus malicious consensus participants. Byzantine fault tolerance models typically assume that fewer than one-third of participants exhibit arbitrary malicious behavior, while other consensus mechanisms may require different trust assumptions.

5.2 Adversary Model

We classify potential adversaries into three categories with progressively increasing capabilities, reflecting the diverse threat landscape confronting healthcare information systems.

5.2.1 Type I: Passive Adversary (Honest-but-Curious)

The Type I adversary represents entities that faithfully execute system protocols but attempt to extract sensitive information through observation and analysis. Such adversaries possess the capability to observe all data passing through compromised nodes within their control, including encrypted traffic patterns, query sequences, and aggregate statistical outputs. They maintain access to auxiliary information sources such as public demographic databases, voter registration records, medical literature, and previously published health statistics. These adversaries possess unlimited computational resources for offline cryptanalysis and statistical inference attacks, enabling them to perform sophisticated correlation analyses and apply advanced ML techniques to infer sensitive attributes.

The Type I adversary can execute multiple queries on databases or trained ML models, potentially crafting query sequences designed to maximize information extraction while remaining within nominal system usage patterns. However, this adversary class is constrained by its adherence to system protocols: it cannot modify data in transit or at rest, cannot corrupt other system components beyond those under its direct control, and cannot forge cryptographic signatures or break established cryptographic primitives.

The primary objectives of Type I adversaries include re-identification attacks, where anonymized medical records are linked to specific individuals through correlation with auxiliary information sources and exploitation of quasi-identifiers such as birthdate, postal code, and demographic attributes. Attribute inference represents another critical threat, wherein adversaries deduce sensitive health attributes including disease status, genetic predispositions, medication regimens, or

TABLE 2: DP Literature Review

Author Concept/Model	Algorithm Implementation	Performance Advantages	Research Gaps Future Work
Huang et al., [47] 2020, DP-ADMM:ADMM-Based Distributed Learning with DP	- DP-Alternating direction method of multipliers (DP-ADMM) - Real-world dataset: Adult data set from UCI ML Repository - Gaussian noise	- Distributed Learning - Noise-resilient, convergent, High privacy guarantee	Increased computation time and sensitive to hyper parameters
Saeidian et al., [34] 2021, Quantifying Membership Privacy via Information Leakage	- Private Aggregation of Teacher Ensembles (PATE) framework - Public data set - Laplace distributed noise	- Novel framework for measuring membership privacy - Accurate measure of membership privacy	Difficult to implement in practice it make challenges for organization to apply framework in real world situation
Gai et al., [14] 2020, DP-Based Blockchain for Industrial Internet-of-Things	- Blockchain Internet of Edge model - Q-learning algorithm - Laplace's distribution noise	- BIoE model enhance the Privacy-Preserving capacity - Tamper resistant	Study impact of data and noise on energy cost in future works
Li et al., [48] 2023, Optimal Trading Mechanism Based on DP and Stackelberg Game in Big Data	- Optimized Unary Encoding (OUE) - LDP based gradient iteration (LGI) algorithm	Guarantee both Privacy and utility for data provider and the users	Optimizing our model with more measured data from data platform and optimizing game theory using Reinforcement learning theory
Zhang et al., [26] 2019, Correlated DP: Feature Selection in ML	- CR-FS Scheme - Mean absolute error (MAE) - ML algorithm SVM and LR - Data sets: Adult, Breast Cancer, Titanic, Porto Seguro	- Better prediction results with ML tasks - Better trade-off between data utility and privacy leaks - Reduce Data correlation	Data correlation may bring new errors with different queries
Bi Jia et al., [13], 2022, Federated Learning Data Aggregation with DP, Homomorphic Encryption and Blockchain in IIoT	- K-means clustering with DP, Homomorphic encryption - Random Forest, AdaBoost	Improved F1-score and accuracy compared normal K-Means	Test for secure exchange and sharing for Enterprise IIoT
Zhang et al., [49] 2023, APDP: Attack-Proof Personalized DP Model for Smart Homes	- APDP model - Fog computing - Real-world Smart home data set	- Improved performance with enhanced bandwidth and reduced service latency - Defeat collusion attack under multiple circumstances - Achieved optimize trade-off between Privacy protection and data utility	Explore cross discipline techniques to further optimize the trade-offs
Zhang et al., [50] 2022, DP-Based double Auction for data market in Blockchain IoT	- Double-Auction Normal Transaction Method (DANTM) - Double-Auction Transaction Method Based on DP (DADPM) - Gaussian mechanism	- Protects participants bid information - Good truthfulness, Privacy, performance	Difficult to determine the size of the Noise
Ali et al., [51] 2022, A privacy enhancing model for IoT using three-way decisions and DP	- Attribute Division Algorithm for DP (3WADD) - Laplace noise - Data sets: Titanic, Adult, Bank, Marketing, Heart Disease, Student Performance	- Automatic division of attributes for DP - Information content and stability of data set	More sophisticated method three way decision attributes of IoT to motivate more organizations to use IoT
Miao Du et al., [52], 2018, DP of Training Model in Wireless Big Data with Edge Computing	Laplace mechanisms with Output Perturbation (OPP) and Objective Perturbation (OJP)	Effectively privacy protection of training data ensuring $> 95\%$ accuracy and data utility	Apply for practical usecases

behavioral risk factors for individuals of interest. Membership inference attacks attempt to determine whether a specific individual's data was included in a training dataset or statistical database, potentially revealing participation in sensitive medical studies or presence of stigmatized conditions.

5.2.2 Type II: Active Adversary (Malicious)

The Type II adversary extends beyond passive observation to actively manipulating system components and data flows. This adversary class possesses all capabilities of Type I adversaries while additionally being able to modify, inject, or delete data packets during network transmission. Active adversaries can compromise and assume control of Edge computing nodes, potentially affecting data processing for multiple IoT sources. They can submit carefully crafted queries specifically designed to amplify privacy leakage beyond what would be revealed through legitimate usage patterns, exploiting potential vulnerabilities in privacy protection mechanisms.

Collusion represents a significant threat multiplier for Type II adversaries, as multiple compromised entities may pool their observations and capabilities to breach privacy or integrity protections that would withstand individual attackers. However, several constraints limit Type II adversary capabilities. Physical security measures protect IoT devices and Edge nodes from direct compromise in most deployment scenarios. Established cryptographic primitives including digital signature schemes, hash functions, and encryption algorithms are assumed to remain computationally infeasible to break. Distributed consensus mechanisms in blockchain implementations impose constraints on the ability of adversaries to unilaterally modify ledger contents.

The attack objectives for Type II adversaries encompass data tampering, wherein patient medical records are modified to influence clinical decisions, insurance claim determinations, or legal proceedings. Privacy budget exhaustion attacks involve carefully sequenced queries designed to deplete privacy protection mechanisms, enabling subsequent queries to extract sensitive information with reduced protection. In distributed ML scenarios, gradient leakage attacks attempt to reconstruct training data from model parameter updates or gradient information exchanged during collaborative learning. Model poisoning attacks inject malicious training examples designed to corrupt the behavior of ML models, potentially causing misdiagnosis or inappropriate treatment recommendations.

5.2.3 Type III: Insider Adversary

The Type III adversary represents perhaps the most challenging threat class: authorized system users who abuse their legitimate access privileges for malicious purposes. Insider adversaries possess all capabilities of Type II adversaries

while additionally holding valid system credentials and authorized access to various system components. These adversaries can access raw, unprotected data before privacy-preserving transformations are applied, possess detailed knowledge of system architecture, implementation details, and security mechanisms, and may abuse privileged access to bypass certain security controls intended to constrain external adversaries.

Despite these extensive capabilities, insider adversaries face several constraints. System activity logging mechanisms record access patterns and operations performed by authenticated users. Role-based access control systems restrict the scope of data and operations accessible even to privileged users, implementing principle of least privilege and separation of duties. In systems employing cryptographic audit trails such as blockchain-based logging, insiders cannot forge or repudiate transactions without detection, as each operation is cryptographically signed and timestamped.

Insider adversaries pursue several attack objectives. Privilege escalation attacks attempt to access data or perform operations beyond the adversary's authorized scope, potentially by exploiting software vulnerabilities or social engineering. Unauthorized disclosure involves exfiltration of protected health information for purposes including financial gain through sale to data brokers, competitive intelligence, blackmail, or personal curiosity. Cover-up attacks attempt to delete or modify audit logs to conceal prior malicious activities, though cryptographic logging mechanisms may render such attempts detectable or impossible.

5.3 Attack Surfaces and Threat Vectors

The multi-layer architecture of healthcare IoT-Cloud systems presents distinct attack surfaces at each computational tier, with threat vectors exploiting various system vulnerabilities.

At the IoT and Edge computing layers, physical device compromise represents a primary threat vector. Adversaries may gain unauthorized physical access to wearable medical devices or home healthcare equipment, potentially extracting cryptographic keys, implanting malware, or tampering with sensor readings. Devices with limited computational resources may lack sophisticated security features such as secure enclaves or hardware-based attestation, making them vulnerable to firmware modification. The resource constraints of IoT devices also render them susceptible to denial-of-service attacks that exhaust battery power or overwhelm processing capacity.

The Edge computing layer faces threats from network traffic analysis. Adversaries monitoring network communications may perform traffic analysis attacks, correlating encrypted packet sizes, timing patterns, and communication frequencies to infer sensitive information about patient conditions or healthcare activities even without decrypting packet contents. These intermediate layers process data from multiple IoT sources, presenting opportunities for cross-patient correlation attacks if data from different individuals is not properly isolated.

A particularly significant threat across both Edge and Cloud layers involves inference attacks on ML models and aggregate statistics. Membership inference attacks analyse the behavior of trained models to determine whether specific individuals' data was included in training datasets, potentially revealing participation in sensitive medical studies. Attribute inference exploits correlations in released statistics or model predictions to deduce sensitive attributes that were not directly disclosed. Model inversion attacks attempt to reconstruct training data from model parameters or predictions, potentially recovering sensitive medical records. Property inference attacks deduce aggregate properties of training data that were not intended for release, such as prevalence of specific conditions or demographic correlations.

Data integrity threats emerge at the Cloud storage layer. Adversaries with write access to storage systems may modify patient records to corrupt medical decision-making, fraudulently alter insurance claims, or manipulate research datasets. More subtle integrity violations involve selective deletion of records to bias statistical analyses or hide evidence of medical errors. The temporal dimension presents additional vulnerabilities, as adversaries may attempt to manipulate timestamps to obscure the chronology of medical events or treatment decisions.

For systems employing distributed ledger technology to ensure data integrity, consensus mechanisms themselves present attack surfaces. Majority attacks in proof-of-work or proof-of-stake systems involve adversaries controlling sufficient computational power or stake to override consensus and modify ledger contents. Selfish mining strategies allow adversaries to gain disproportionate influence over consensus by strategically withholding and releasing blocks. Eclipse attacks isolate specific nodes from the honest network, feeding them false information about the state of the distributed ledger.

Cross-layer threats span multiple architectural tiers. Linkage attacks combine information from multiple queries, potentially issued to different system components or at different times, to circumvent privacy protections that would be effective against individual queries. Collusion attacks involve multiple adversaries, potentially with compromised nodes at different architectural layers, pooling their observations to amplify information extraction. Composition attacks exploit the accumulation of privacy loss across multiple privacy-preserving mechanisms or repeated queries to the same underlying dataset.

5.4 Security Objectives

The threat landscape outlined above necessitates five fundamental security objectives that healthcare IoT-Cloud systems must satisfy.

Confidentiality requires that individual patient records remain confidential even when aggregate statistics, trained ML models, or research findings derived from the data are publicly released. This objective extends beyond simple access control to require that released information provably limits what can be inferred about individuals. Quantifiable privacy guarantees must bound the maximum information leakage that can occur through any sequence of queries or analyses, even when adversaries possess arbitrary auxiliary information and unlimited computational resources for inference attacks.

Integrity demands protection of patient data, ML models, audit logs, and system configurations against unauthorized modification. Any tampering attempts must be detectable through cryptographic verification mechanisms, and the system must maintain evidence of data provenance enabling verification of authenticity. Integrity protections must be tamper-evident, meaning that modifications leave detectable traces

even if the adversary controls storage infrastructure. For critical healthcare data, integrity requirements may extend to non-repudiation, ensuring that entities cannot deny having performed specific operations.

Availability ensures that healthcare services remain operational despite adversarial disruption attempts. This encompasses resilience against denial-of-service attacks, infrastructure failures, and resource exhaustion. For emergency healthcare scenarios, availability requirements include strict latency bounds on critical operations, as delays in responding to physiological emergencies can result in patient harm or mortality. Availability must be maintained even under partial system compromise, requiring redundancy and graceful degradation of service quality rather than complete failure.

Accountability requires that all data access and modification operations be attributable to specific entities through immutable audit trails. Authenticated users must not be able to perform operations anonymously or repudiate actions they have taken. Audit mechanisms must themselves resist tampering, as adversaries may attempt to cover their tracks by modifying logs. Accountability enables forensic investigation of security incidents, deters insider attacks through the threat of detection, and supports regulatory compliance with healthcare privacy regulations requiring access logging.

Privacy Preservation mandates formal, quantifiable guarantees limiting information leakage across multiple queries and analyses. Unlike confidentiality, which focuses on access control, privacy preservation addresses the fundamental tension between data utility and individual privacy in statistical databases and ML systems. Privacy guarantees must hold under worst-case adversarial behavior, accounting for arbitrary auxiliary information and composition of multiple analyses. The quantification of privacy loss must enable healthcare organizations to make informed decisions about acceptable risk levels while complying with regulatory requirements such as HIPAA in the United States or GDPR in the European Union.

6 Proposed Solution

The threat model established in Section 5 identifies three adversary classes (honest-but-curious, malicious, and insider threats) along with attack surfaces spanning inference attacks, data tampering, and consensus attacks. To address this comprehensive threat landscape, our research proposes an integrated framework comprising three complementary components that provide defense-in-depth security:

(i) A multi-layered IoT-Edge-Cloud computing architecture to improve response time for healthcare applications through strategic workload distribution

(ii) Privacy-preserving DP techniques to maintain data utility for medical research whilst protecting patient identity, specifically addressing the inference attacks (re-identification, membership inference, attribute inference) posed by Type I and Type II adversaries

(iii) Secure cloud storage utilising blockchain technology for healthcare transaction data, providing tamper-evident audit trails and integrity verification to counter data tampering threats from Type II and Type III adversaries

To facilitate real-time processing for delay-intolerant emergency healthcare responses, the proposed architecture incorporates multiple computational layers as depicted in Fig. 3. This IoT, Edge and Cloud (IEC) framework is characterized by increasing compute capacity and storage permanence in the upper layers, counterbalanced by corresponding increases in communication latency. This design enables task distribution based on response urgency, with critical operations deployed to lower layers and persistent data directed to cloud infrastructure. The hierarchical nature of this architecture also implements defense-in-depth against the cross-layer threats identified in Section 5, limiting the scope of compromise at each tier through role-based access control and computational isolation.

We achieve privacy protection by applying DP techniques during the analysis phase, where systems are most vulnerable to compromise. As established in our threat model, Type I adversaries with unlimited computational resources and access to auxiliary information can perform sophisticated inference attacks through multiple queries on aggregated datasets to extract personal patient information. By applying calibrated noise to ML training datasets, we provide mathematical guarantees bounding information leakage even under worst-case adversarial conditions. For any two datasets D and D' differing by a single record, the probability ratio of observing any output is bounded by ε, ensuring that healthcare institutions can share data with researchers whilst preventing individual patient re-identification, thus preserving population-level insights for medical advancement.

6.1 Multi-Layered Compute-Storage Architecture

The proposed solution incorporates a hierarchical edge and cloud architecture to ensure rapid response times and enhanced security for healthcare IoT applications, as illustrated in Fig. 3.

The IEC architecture strategically distributes computational workloads and data based on proximity to data sources, processing speed requirements, and storage permanence needs. This distribution optimizes system performance for various healthcare scenarios with different response criticality profiles.

The effectiveness of this approach is demonstrated through several healthcare applications. In emergency response scenarios, patient wearables detecting physiological abnormalities trigger immediate data processing at proximal Edge Computing nodes, while simultaneously transmitting data to Cloud storage via other Edge nodes for comprehensive analytics. This dual-path approach ensures both immediate intervention and long-term data utilization for population health insights by public health authorities. Child Health Information systems represent intermediate computational requirements, utilizing Edge nodes to compare current sensor readings against historical baselines for generating caregiver recommendations. Telemedicine applications employ Edge Computing resources for processing patient vitals and Electronic Medical Records while supporting video consultations.

This layered approach optimizes resource allocation based on response time criticality and computational complexity, ensuring appropriate performance characteristics across diverse healthcare applications without unnecessary resource expenditure.

This layered architecture directly addresses the availability and resilience requirements identified in Section 5. By distributing processing across multiple tiers, the system maintains emergency response capabilities even if higher-layer Edge or Cloud nodes are compromised or unavailable. The

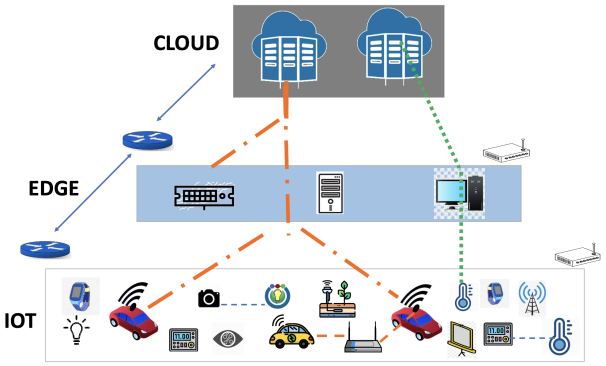

Fig. 3: Multi-Layered Compute-Storage Architecture

hierarchical access control implemented at each layer constrains unauthorized access attempts by Type II and Type III adversaries, as compromise of a single layer does not grant access to all system resources.

6.2 Secure IoT-Cloud Healthcare Application Architecture

The proposed secure architecture for healthcare analytics integrates DP with distributed ledger technology to address the confidentiality, integrity, and accountability requirements established in Section 5. Fig. 4 illustrates the data flow and control mechanisms of this architecture. The security framework comprises six integrated components:

(i) Data Collection: Patient data is aggregated from distributed IoT devices, wearables, and clinical systems. This heterogeneous data contains sensitive personally identifiable information (PII) requiring robust protection against the re-identification attacks described in Section 5.

(ii) Preprocessing: Prior to privacy transformations, the data undergoes cleansing, normalization, and initial anonymization to remove direct identifiers and minimize re-identification risks through quasi-identifiers, addressing the linkage attack vectors identified in our threat model.

(iii) Privacy Budget: The framework implements a quantifiable privacy budget using the ε parameter to constrain information leakage across multiple queries. This quantification establishes the maximum permissible privacy loss when adding or removing individual records from the dataset, directly addressing the composition attacks where adversaries accumulate information through repeated queries.

(iv) Noise Injection: Calibrated statistical noise is introduced to dataset calculations according to mathematically rigorous DP mechanisms. This process obscures individual data points while preserving statistical validity for aggregate analysis, providing formal guarantees against membership inference and attribute inference attacks by Type I adversaries.

(v) Blockchain Security: Blockchain technology is employed to address the integrity and accountability requirements from Section 5, particularly the data tampering threats posed by Type II adversaries and audit trail tampering by Type III insider adversaries. The system utilizes off-chain storage mechanisms such as Ethereum off-chain solutions to mitigate transaction costs while maintaining data integrity across distributed stakeholders. Each transaction is cryptographically hashed and timestamped, creating tamper-evident records that detect any modification attempts. The current implementation employs the Raft consensus protocol, which provides crash fault tolerance (CFT) with deterministic finality, tolerating up to $f < n/2$ crashed nodes in a network of n ordering service nodes. This CFT model operates under the assumption that ordering service operators within the permissioned consortium are trusted and do not exhibit arbitrary malicious behaviour. For deployment scenarios requiring resilience against Byzantine adversaries (Type II and Type III), the architecture is designed to accommodate a Byzantine fault-tolerant (BFT) ordering service, such as those supported by Hyperledger Fabric's pluggable consensus framework, which would provide integrity guarantees with up to $f < n/3$ arbitrarily compromised consensus nodes. Cryptographic transaction signing and hash chaining provide tamper-evidence at the ledger level independent of the consensus mechanism, ensuring that any post-consensus manipulation of committed blocks remains detectable.

(vi) Privacy-Preserving Analytics: The resulting differentially-private dataset enables sophisticated statistical and computational analyses without compromising individual privacy. This approach balances clinical utility with robust privacy guarantees, ensuring that even Type III insiders with elevated privileges cannot bypass DP protections, as noise injection occurs before data leaves the trusted computational environment.

This dual-protection framework provides complementary defenses against the distinct threat classes in our adversary model: DP mechanisms protect against inference attacks on aggregate data and ML models (Type I and Type II threats), while distributed ledger technology ensures integrity verification and accountability through immutable audit trails (Type II and Type III threats). Hierarchical access control across computational layers constrains unauthorized access and limits the scope of compromise (all adversary types).

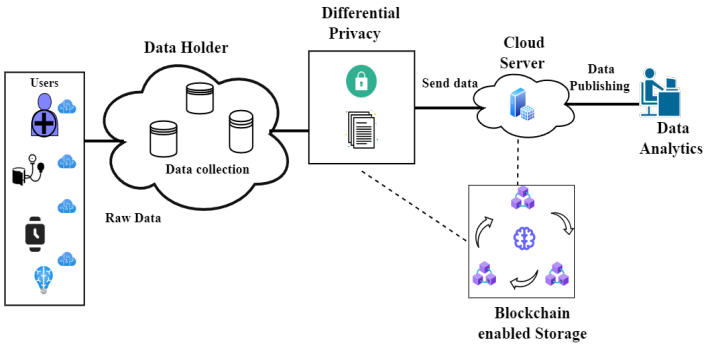

Fig. 4: Securing Healthcare Data by DP and Blockchain

6.3 Procedure for Differential Privacy with Machine Learning

The proposed methodology for implementing DP within ML workflows is illustrated in Fig. 5. This process comprises a structured sequence of operations designed to maintain analytical utility while providing mathematically rigorous privacy guarantees.

(i) Data Acquisition: The process begins with a dataset containing sensitive healthcare information that requires privacy

protection while retaining analytical value.

(ii) Data Preprocessing: The dataset undergoes normalization, cleansing, and standardization procedures to ensure consistency and quality prior to analysis.

(iii) ML Implementation: Supervised or unsupervised learning algorithms are applied to the preprocessed data. For ensemble methods, this includes the generation of multiple decision trees for classification or regression tasks.

(iv) Statistical Perturbation: Controlled stochastic noise is introduced to the model outputs. This perturbation prevents adversarial reconstruction of individual data points while preserving aggregate statistical properties.

(v) DP Mechanism: A formal DP framework calibrates noise introduction according to sensitivity analysis and privacy budget constraints, ensuring mathematical guarantees of individual privacy.

(vi) Aggregation Procedures: Statistical measures such as cluster centroids, distribution parameters, and confidence intervals are calculated from the privacy-protected outputs, maintaining population-level insights while obscuring individual contributions.

(vii) Analytical Interpretation: The differentially private aggregated results are analyzed to extract clinically relevant insights, with careful attention to potential implications for healthcare decision-making.

(viii) Privacy-Utility Assessment: Quantitative evaluation of both privacy preservation (through DP guarantees) and analytical utility (through accuracy metrics) is performed to validate the methodology.

(ix) Knowledge Synthesis: The final privacy-protected analytical outcomes are synthesized into actionable healthcare intelligence that satisfies both privacy and utility requirements.

(x) Process Completion: The workflow concludes with documentation of privacy parameters and methodological constraints to ensure reproducibility and transparency.

The algorithms detailed in the following subsections implement DP at various stages of ML workflows, providing formal guarantees against the gradient leakage and model inversion attacks identified in Section 5. By introducing calibrated noise during training (input perturbation) or to model outputs (output perturbation), these mechanisms bound the influence of any single training example on model parameters, inherently providing robustness against model poisoning attacks by Type II adversaries.

6.4 Algorithms

This section delineates the integration of DP mechanisms with four ML algorithms: Random Forest, K-Means, Logistic Regression, and Naive Bayes. Each method incorporates Laplace, Gaussian, or hybrid noise calibrated to the sensitivity (Δ) and privacy budget (ε), ensuring (ε, δ)-DP guarantees. Comprehensive pseudocode and procedural descriptions are provided for reproducibility and practical implementation.

6.4.1 Differentially Private Random Forest

The Random Forest ensemble is modified by injecting Gaussian noise into leaf node predictions post-training. Let f denote the sensitivity of the prediction function, computed as the maximum L_2-norm difference in outputs between adjacent datasets. For each leaf, Gaussian noise $\mathcal{N}(0, \sigma^2)$ is sampled, where $\sigma = \frac{\Delta f \sqrt{2\ln(1.25/\delta)}}{\varepsilon}$. The introduction of noise at the leaf level, rather than during the training process itself, represents a post-processing approach to DP that maintains the fundamental splitting criteria of the individual trees whilst privatising their outputs. Algorithm 1 formalises this process in detail.

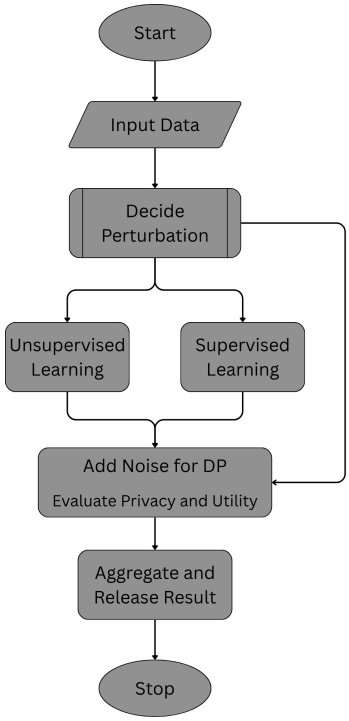

Fig. 5: ML with DP Mechanism

Algorithm 1: Differentially Private Random Forest

Data: Training data X, labels y, privacy budget ε, failure probability δ, sensitivity Δ
Result: Differentially private random forest DP_{RF}

foreach *tree* $t \in$ *RandomForest* **do**
  Train base decision tree: $t \leftarrow \textsc{TrainTree}(X, y)$;
  foreach *leaf* $l \in t$ **do**
    $\sigma \leftarrow \Delta\sqrt{2\ln(1.25/\delta)}/\varepsilon$;
    $noise \leftarrow \texttt{SampleGaussian}(0, \sigma)$;
    l.prediction $\leftarrow l$.prediction $+ noise$;
  end
end
return DP_{RF}

This algorithm ensures that each individual tree's predictions are perturbed with calibrated noise, whilst the overall ensemble maintains its predictive power through aggregation, as the independent noise additions tend to average out across multiple trees.

6.4.2 Differentially Private K-Means Clustering

The K-Means clustering algorithm, an unsupervised learning technique, requires modification to ensure DP during the iterative centroid refinement process. Cluster centroids are perturbed using Laplace noise during each iteration. Let Δf represent the L_1-sensitivity of centroid updates. At each

iteration, Laplace noise $\text{Lap}(\Delta f/\varepsilon)$ is added component-wise to each centroid coordinate.

The iterative nature of K-Means presents a particular challenge for privacy preservation, as each update potentially leaks information. By calibrating the noise to the sensitivity and privacy budget at each step, the algorithm maintains its clustering efficacy whilst providing formal DP guarantees. Algorithm 2 presents the complete procedure.

Algorithm 2: Differentially Private K-Means

Data: Dataset D, clusters k, max iterations T, sensitivity Δ, ε
Result: Private centroids C, clusters S

$C \leftarrow \textsc{InitialiseCentroids}(D, k)$;
for $t = 1$ **to** T **do**
  $S \leftarrow \texttt{AssignClusters}(D, C)$;
  foreach *cluster* $c_i \in C$ **do**
    $\tilde{\mu}_i \leftarrow \text{mean}(S_i) + \text{Lap}(\Delta/\varepsilon)$;
    $C \leftarrow C \cup \tilde{\mu}_i$;
  end
end
return C, S

This differentially private adaptation of K-Means preserves the algorithm's ability to identify natural groupings in the data whilst ensuring that the resulting clusters and centroids do not compromise the privacy of individual observations in the dataset.

6.4.3 Differentially Private Logistic Regression

Logistic Regression, a cornerstone supervised learning algorithm for binary classification, can be rendered differentially private through gradient perturbation during the optimization process. This approach injects calibrated noise into the gradient computations used for parameter updates, ensuring that the trained model does not reveal sensitive information about individual training examples. Let $\nabla L(w)$ denote the loss gradient with respect to model weights w. The sensitivity of this gradient, Δ, is defined as the maximum L_1-norm difference in gradients that could result from adding or removing a single training example, calculated as $\Delta = \max \|\nabla L_i(w)\|_1$. Laplace noise $\eta \sim \text{Lap}(\Delta/\varepsilon)$ is injected into each gradient component prior to weight updates.

This gradient perturbation approach represents an input perturbation method for DP, in contrast to the output perturbation used in Random Forest. By introducing noise during the training process itself, the algorithm ensures that the entire model fitting procedure maintains privacy guarantees. Algorithm 3 details this procedure.

The function `PerturbGradient` adds Laplace noise to each component of the gradient vector, with the noise magnitude calibrated according to the sensitivity and privacy budget. This approach ensures that the resulting logistic regression model maintains its predictive capabilities whilst providing formal privacy guarantees.

Algorithm 3: Differentially Private Logistic Regression

Data: Training set D, ε, Δ, learning rate η, iterations T
Result: Private weights w

Initialise $w \leftarrow \mathbf{0}$;
for $t = 1$ **to** T **do**
  Compute $\nabla L(w) = \sum_{(x_i, y_i) \in D} (p_i - y_i) x_i$;
  $\nabla L(w) \leftarrow \texttt{PerturbGradient}(\nabla L(w), \Delta, \varepsilon)$;
  $w \leftarrow w - \eta \nabla L(w)$;
end
return w

6.4.4 Differentially Private Naive Bayes

The Naïve Bayes classifier can be rendered differentially private through two complementary approaches depending on the feature representation, both presented in Algorithm 4.

Variant A (Discrete Features): For datasets with discrete or binary features, the classical approach perturbs the sufficient statistics (counts) used to estimate conditional probabilities. For each feature x_j and class c_k, Laplace noise drawn from $\text{Lap}(1/\varepsilon)$ is added to both the feature-class count and the class count, where the sensitivity equals 1 since adding or removing a single record changes any individual count by at most one unit. Because independent noise on numerator and denominator does not guarantee that the resulting ratio lies in $[0, 1]$ or that probabilities sum to unity across features, the perturbed estimates are clamped to $[0, 1]$ and renormalized per class to form valid probability distributions.

Variant B (Continuous Features): For datasets with continuous features, input perturbation adds calibrated noise (Laplace, Gaussian, or hybrid) directly to the training feature values, after which a standard Gaussian Naive Bayes classifier estimates class-conditional mean and variance parameters from the perturbed data. This approach leverages the post-processing immunity property of DP, as the Gaussian likelihood estimation on already-privatised data does not consume additional privacy budget. Variant B is appropriate for healthcare datasets containing continuous clinical measurements (e.g., CD4 and CD8 counts) that are better suited to Gaussian likelihood estimation than discrete count models.

6.4.5 Theoretical Guarantees

Each of the described algorithms satisfies (ε, δ)-DP, with mathematical proofs derived from the fundamental properties of DP, particularly the post-processing immunity theorem. This theorem establishes that any function of a differentially private output remains differentially private, without requiring additional privacy budget expenditure.

For the Random Forest algorithm, the addition of noise to leaf node predictions preserves privacy as predictions depend solely on these perturbed aggregates, with no further access to the original training data. The K-Means and Logistic Regression algorithms adhere to the sequential composition theorem, which quantifies the cumulative privacy loss across multiple operations on the same data. The Naïve Bayes discrete variant (Variant A) satisfies DP through the perturbation of sufficient statistics, with post-hoc clamping and renormalization preserving privacy guarantees via the post-processing immunity theorem. The continuous (Gaussian) variant (Variant B) inherits its privacy guarantee directly from the input perturbation step, as fitting class-conditional Gaussian distributions to already-privatised data constitutes post-processing of a differentially private output.

These theoretical guarantees ensure that the privacy properties of the algorithms hold regardless of the adversary's computational power or background knowledge, providing a robust foundation for privacy-preserving ML in sensitive

Algorithm 4: Differentially Private Naive Bayes

Data: Training set $D = \{(x_i, y_i)\}$, ε, sensitivity Δ, classes C, features F
Result: Differentially private Naive Bayes model NB_{priv}

```
// Variant A: Discrete count perturbation
```
foreach *class* $c_k \in C$ **do**
  $n_k \leftarrow |\{x \in D : y = c_k\}|$;
  foreach *feature* $x_j \in F$ **do**
    $count_{j,k} \leftarrow \sum_{x \in D} \mathbf{1}(x_j = 1 \wedge y = c_k)$;
    $\tilde{P}(x_j|c_k) \leftarrow \frac{count_{j,k} + \text{Lap}(1/\varepsilon)}{n_k + \text{Lap}(1/\varepsilon)}$;
    $\tilde{P}(x_j|c_k) \leftarrow \text{clamp}(\tilde{P}(x_j|c_k),\ 0,\ 1)$;
  end
  Renormalise: $\tilde{P}(\cdot|c_k) \leftarrow \tilde{P}(\cdot|c_k) / \sum_j \tilde{P}(x_j|c_k)$;
end
return $\tilde{P}(x_j|c_k)$

```
// Variant B: Continuous input perturbation
```
foreach $x_i \in D$ **do**
  $\tilde{x}_i \leftarrow x_i + \text{Noise}(\Delta, \varepsilon)$ `// Laplace, Gaussian, or hybrid`
end
Fit Gaussian NB: $\hat{\mu}_{j,k}, \hat{\sigma}^2_{j,k} \leftarrow \text{MLE}(\{\tilde{x}_i : y_i = c_k\})$;
return $\hat{\mu}_{j,k}, \hat{\sigma}^2_{j,k}\ \forall\ j, k$

healthcare applications. The formal nature of these guarantees distinguishes DP from heuristic anonymization approaches that lack such mathematical rigor.

6.5 Blockchain-Based Trust Framework

To establish robust trust in our healthcare data ecosystem, we propose integrating blockchain technology with our multi-layered architecture and DP framework. This blockchain component creates a trusted foundation for healthcare transactions where centralized trust cannot be assumed, particularly in multi-stakeholder environments involving patients, providers, insurers, and researchers.

6.5.1 Trust Architecture using Blockchain

The proposed blockchain implementation addresses the fundamental trust challenges in healthcare data sharing by creating a decentralized trust framework where no single entity controls the entire system. This approach is particularly valuable in healthcare contexts, where patients must trust multiple parties with their sensitive information [53].

Our design employs a permissioned blockchain network where trusted healthcare entities (hospitals, clinics, research institutions, and regulatory agencies) serve as validator nodes. Unlike public blockchains, this consortium model balances efficiency with trusted verification processes [38]. The architecture establishes trust through:

(i) Distributed Consensus: Critical healthcare transactions achieve validity only through agreement among multiple independent validators, eliminating single points of trust failure.

(ii) Cryptographic Verification: Digital signatures and hash functions create mathematical proof of data integrity, replacing institutional trust with cryptographic certainty.

(iii) Immutable Record-Keeping: The append-only structure of blockchain creates tamper-evident records, allowing stakeholders to trust the permanence and integrity of healthcare data histories.

(iv) Trust Transparency: The blockchain provides visibility into who accessed what data and when, creating accountability without requiring blind trust in any single record-keeper.

6.5.2 Establishing Trust Between Privacy and Utility

A significant innovation in our proposed framework is resolving the traditional trust tension between data privacy and research utility. By integrating blockchain with DP mechanisms, we create a system where:

(i) Patients can trust that their privacy is mathematically guaranteed through DP, with provable limits on information disclosure.

(ii) Researchers can trust the authenticity and integrity of aggregated healthcare data without needing access to individual records.

(iii) Regulators can trust the compliance status of the system through cryptographically verified audit trails.

This balanced approach addresses the trust paradox in healthcare informatics, maintaining trust in both privacy protection and data utility simultaneously [28]. The blockchain would store cryptographic commitments to privacy budgets, creating verifiable records of compliance with privacy guarantees.

6.5.3 Smart Contracts as Trust Automation

The proposed framework would implement specialized smart contracts that codify trust relationships into executable agreements. These self-enforcing protocols would automate trust in consent by converting patient preferences into cryptographically enforced access rules, allowing patients to trust that their sharing preferences are honored without ongoing monitoring. They would create trustworthy audit trails by generating immutable records of every data access event that all stakeholders can independently verify, establishing trust through transparency. Additionally, they would enable trusted multi-party research by facilitating complex data sharing arrangements between competing institutions that might otherwise lack sufficient trust for collaboration.

6.5.4 Blockchain Architecture Specifications

The proposed blockchain integration employs a permissioned consortium architecture to address security and auditability requirements of privacy-preserving healthcare analytics. This subsection delineates the technical specifications for blockchain implementation within the multi-layer IoT-Cloud framework.

(i) Platform Selection: The architecture specifies Hyperledger Fabric [54] as the blockchain substrate, selected for its permissioned network access compatible with regulatory frameworks (HIPAA, GDPR), modular architecture enabling healthcare-specific customization, private data collections for confidential information sharing among authorized participants, and deterministic transaction finality through crash fault-tolerant ordering services. This selection was informed by established healthcare deployment precedents [55] and

enterprise-grade tooling.

(ii) Consensus Mechanism: The framework employs Raft consensus protocol [56], providing crash fault tolerance with deterministic finality suitable for permissioned healthcare networks. The network topology comprises peer nodes deployed at the Edge layer maintaining ledger replicas and executing smart contract logic, ordering service nodes at Cloud infrastructure establishing transaction sequence and block generation, certificate authority infrastructure managing cryptographic identities and access credentials across the multi-layer architecture, and client applications at IoT and Edge layers submitting transactions through SDK interfaces.

(iii) Smart Contract Architecture: The blockchain layer implements chaincode (Hyperledger terminology for smart contracts) encoded in Go language. The Privacy Budget Ledger maintains immutable records of epsilon allocation, tracks cumulative privacy expenditure across analytical queries, enforces budget constraints through transaction validation logic, and records privacy parameters (ε, δ, sensitivity Δf, noise distribution) for audit verification. The Data Provenance Chain establishes cryptographic linkage between data transformations across architectural layers, recording SHA-256 hashes at each processing stage (IoT, Edge and Cloud), maintaining temporal metadata such as ISO 8601 timestamps and processing duration, and enabling end-to-end traceability for regulatory compliance. The Access Control Enforcement implements attribute-based access control policies, logs access attempts with requestor identity verification, and records patient consent states for data processing operations.

(iv) Transaction Structure: Blockchain transactions encapsulate transaction identifier (UUID v4), timestamp with nanosecond precision, operation type enumeration, namely, DATAINGESTION, DPQUERY, BUDGETUPDATE, and ACCESSREQUEST, cryptographic hash of data or query result (SHA-256 digest), privacy parameters encoded as JSON object containing $\{\varepsilon, \delta, \text{noiseType}, \Delta f\}$, digital signature (ECDSA with NIST P-256 curve) authenticating authorized entity, and metadata fields including device identifier, processing layer designation, and analytical purpose classification. This structure provides cryptographic proof of privacy guarantee compliance without exposing sensitive healthcare information.

(v) Integration with Differential Privacy: The blockchain interfaces with the DP module through REST API endpoints, enabling atomic operations that precede each analytical query. Prior to executing any DP-protected query, the system invokes the Privacy Budget Ledger smart contract to verify remaining budget sufficiency and record the impending privacy expenditure. This integration creates cryptographically verifiable proof that privacy budgets have not been exceeded, addressing the composition attack vulnerability where adversaries accumulate information through repeated queries [7]. Post-query execution, the system commits query results as cryptographic hashes and privacy parameters to the ledger, establishing an immutable audit trail.

(vi) Cryptographic Primitives: The blockchain implementation employs standardized cryptographic algorithms. ECDSA with NIST P-256 curve (secp256r1) provides digital signatures with 128-bit security strength [57]. SHA-256 hash function ensures data integrity verification and block linking, conforming to NIST FIPS 180-4 specifications [58]. X.509 v3 certificates manage identity within the permissioned network, issued by the Hyperledger Fabric Certificate Authority. These selections align with NIST recommendations for cryptographic algorithm standards in healthcare information systems [59].

This blockchain specification provides comprehensive technical grounding for implementing tamper-proof privacy guarantees within the proposed healthcare system. While this work focuses on architectural design and DP implementation validation through experimental evaluation (Section 8), the detailed blockchain specifications establish a rigorous foundation for future system integration.

6.5.5 Trust Recovery Mechanisms

A critical aspect of any trust framework is the ability to recover from failures. Our blockchain design incorporates mechanisms for fault detection, identifying crashed or unresponsive nodes through heartbeat monitoring and consensus timeouts. Under the current Raft-based CFT ordering service, detection is limited to crash faults; extending to Byzantine fault detection, capable of identifying arbitrarily malicious behaviour through consensus discrepancies, would require migration to a BFT ordering service. Trust revocation processes allow for removing compromised entities from trusted operations, while transparent remediation documents all trust violations and recovery actions on the blockchain itself [53]. These mechanisms acknowledge that trust violations will occasionally occur while providing structured, transparent processes for maintaining system-wide trust even when individual components fail. The resilience of the trust framework is particularly important in healthcare contexts where continuity of operations directly impacts patient outcomes.

7 IMPLEMENTATION

The implementation of DP in healthcare systems necessitates meticulous handling of medical records, encompassing patient histories, diagnoses, and treatments, to balance privacy preservation with analytical utility. This section delineates the methodologies for noise injection, key considerations for medical data, and the framework for evaluating ML algorithms under DP constraints.

7.1 Noise addition in Differential Privacy

Noise introduction in DP ensures individual privacy while maintaining statistical validity. In healthcare contexts, noise is applied during aggregation (e.g., calculating disease prevalence) or statistical computations (e.g., average treatment duration). The calibration of noise magnitude depends on the sensitivity of the query and the privacy budget (ε), which governs the trade-off between privacy guarantees and data accuracy. Table 3 presents a comparison of different noise mechanisms.

(i) *Laplace Noise Mechanism:* Laplace noise, drawn from a symmetric exponential distribution, is scaled by the sensitivity (Δ) of the query and ε. This mechanism is optimal for low-dimensional datasets (e.g., patient counts per diagnosis) due to its heavy-tailed distribution, which provides robust privacy guarantees. However, excessive noise at low ε values

may degrade utility.

(ii) *Gaussian Noise Mechanism:* Gaussian noise, characterized by a bell-shaped normal distribution, is governed by (ε, δ)-DP, where δ represents the probability of privacy leakage. It is suitable for high-dimensional data (e.g., electronic health records) due to its lighter tails, which preserve utility in complex analyses.

(iii) *Hybrid Laplace-Gaussian Noise Mechanism:* The proposed hybrid mechanism combines Laplace and Gaussian noise to leverage the complementary strengths of both distributions, Laplace's robust privacy guarantees through heavy tails and Gaussian's utility preservation through concentrated mass near zero. This approach addresses the limitation that Laplace noise may be excessive for high-dimensional data while Gaussian noise may provide insufficient protection for outlier-sensitive queries. Given a query function $f : \mathcal{D} \to \mathbb{R}^d$ with sensitivity Δf and total privacy budget $\varepsilon_{\text{total}}$, the hybrid mechanism generates noisy output as:

$$M(D) = f(D) + \text{Lap}\left(\frac{\Delta f}{\varepsilon_L}\right) + \mathcal{N}\left(0, \frac{2\ln(1.25/\delta) \cdot \Delta f^2}{\varepsilon_G^2}\right) \quad (1)$$

where ε_L and ε_G denote the privacy budget allocated to Laplace and Gaussian components respectively, satisfying $\varepsilon_L + \varepsilon_G = \varepsilon_{\text{total}}$ under sequential composition theorem [7], and δ is the probability of privacy failure for the Gaussian component. The privacy budget is partitioned between noise mechanisms based on data characteristics. For datasets with mixed dimensionality or varying sensitivity profiles, we employ an equal allocation strategy ($\varepsilon_L = \varepsilon_G = \varepsilon_{\text{total}}/2$) as the baseline. Alternative allocation strategies include Laplace-dominant ($\varepsilon_L = 0.7\varepsilon_{\text{total}}, \varepsilon_G = 0.3\varepsilon_{\text{total}}$) for low-dimensional data requiring stronger tail guarantees, and Gaussian-dominant ($\varepsilon_L = 0.3\varepsilon_{\text{total}}, \varepsilon_G = 0.7\varepsilon_{\text{total}}$) for high-dimensional data prioritizing utility preservation. By the sequential composition property of DP, applying two independent randomized mechanisms with privacy guarantees ε_L-DP and ε_G-DP respectively yields an overall privacy guarantee of $(\varepsilon_L + \varepsilon_G)$-DP = $\varepsilon_{\text{total}}$-DP [7]. The resulting hybrid noise distribution exhibits intermediate tail behavior; heavier than pure Gaussian but lighter than pure Laplace, creating a more nuanced privacy-utility trade-off. This distribution decays gradually over a larger range than either individual distribution, making it particularly suitable for protecting datasets with mixed dimensionality or varying sensitivity profiles. The moderate tail behavior provides robust privacy guarantees similar to Laplace while preserving analytical utility comparable to Gaussian noise, especially beneficial for healthcare datasets containing both low-dimensional aggregates (patient counts) and high-dimensional features (electronic health records). The optimal allocation of the privacy budget between components can be determined through empirical evaluation via grid search over the allocation parameter $\alpha \in [0, 1]$ where $\varepsilon_L = \alpha\varepsilon_{\text{total}}$ and $\varepsilon_G = (1 - \alpha)\varepsilon_{\text{total}}$.

7.2 Key Considerations in Differential Privacy for Medical Health Data

(i) *Privacy Budget*: Cumulative ε across multiple queries must be constrained to prevent excessive privacy loss. Dynamic budgeting strategies, such as zero-concentrated DP, can optimize allocations for longitudinal studies.

(ii) *Sensitivity of Data*: Attributes like genetic markers or rare diseases necessitate higher noise due to their identifiability risks. Sensitivity analysis should precede noise calibration.

(iii) *Noise Mechanism Selection*: Laplace suits bounded, low-dimensional queries (e.g., disease counts), while Gaussian is preferable for unbounded, high-dimensional analyses (e.g., ML model training).

(iv) *Granularity vs Utility*: Aggregating data (e.g., age groups instead of exact ages) reduces sensitivity but may obscure critical patterns. Context-aware aggregation preserves utility for clinical decision-making.

(v) *Analysis Type*: The type of analysis being performed matters. Simple counts or averages may require less noise compared to complex statistical analyses or ML tasks.

(vi) *Data Size*: Larger data sets can tolerate more noise, while smaller data sets might need careful handling to avoid excessive distortion.

(vii) *Privacy vs Utility*: A critical consideration is the trade-off between preserving individual privacy and maintaining the usefulness of the data. Striking the right balance is essential to ensure meaningful results without compromising privacy.

(viii) *Regulations and Standards*: HIPAA and GDPR mandate strict anonymization. DP parameters must align with legal thresholds for de-identification.

(ix) *Expert Consultation*: Input from clinicians ensures noise levels do not invalidate medical insights, while privacy experts validate compliance with ethical standards.

7.3 Evaluation of ML Algorithms with DP techniques

Healthcare research and analytics require accurate data for delivering correct and error free decisions beneficial to humanity. Healthcare researchers and analysts utilise medical datasets to improve community health indices, hospital systems and public health policies. Various analytical approaches are required for different situations, including: (i) descriptive (e.g., number of hospitalised patients in the previous week), (ii) diagnostic (e.g., hospitalization causes), (iii) predictive (e.g., likely hospitalizations in the coming week) and (iv) prescriptive (e.g., preventative medicine recommendations). Consequently, different ML algorithms are employed in healthcare analytics. For instance, patients visiting hospitals may be grouped according to symptom intensity using the K Means ML algorithm to cluster patients into k=3 groups (no symptoms/mild symptoms/strong symptoms). K means is an unsupervised ML algorithm that clusters data points into groups based on similarity. To predict infection likelihood, Logistic Regression is valuable, identifying relationships between patient features such as comorbidity, age and present symptoms to generate binary predictions about future infection status. Such predictions

TABLE 3: Comparison of Noise

Factor	Laplace Noise	Gaussian Noise	Combined L-G
Probability Distribution	Laplace distribution	Gaussian distribution	Combined distribution
Symmetry	Symmetric around 0	Symmetric around 0	Symmetric around 0
Scale	Controlled by sensitivity	Controlled by sensitivity	Controlled by sensitivity
Privacy Budget	ε parameter	ε parameter	ε parameter
Tail Behavior	Heavier tails	Lighter tails	Moderate tails
Data Type	Suitable for bounded data, lower dimension	Suitable for any data type, higher dimension	Suitable for any data type, higher dimension
Adding Noise Mechanism	Add noise to each data point	Add noise to aggregate	Add to individual or aggregate
Trade-off	More noise, high privacy	Lesser noise, high privacy	Moderate noise, high privacy

facilitate hospital/bed/medicine capacity planning and preventative care provision. Naive Bayes, a supervised ML algorithm using labelled data, classifies instances into predefined classes based on independent features, aiding diagnosis. For example, Naive Bayes can identify jaundice when symptoms include yellow eye colouration, turbid urine and elevated body temperature. Random Forest ML algorithms are employed for predicting drug sensitivity. DP is applied to datasets to enable ML analysis whilst protecting individual privacy. Generally, K means, Logistic Regression, Random Forest and Naive Bayes cannot be directly compared as they address different tasks. A primary objective of this work is to assess ML query/analytics accuracy on differentially private datasets compared to accuracy on original datasets. Therefore, experiments evaluate accuracy against privacy budget ε. The efficacy of DP lies in its ability to maintain data usability whilst protecting individual privacy, measured by accuracy:

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + TN + FN}$$

Where,
TP (True Positive) i.e. the count of positive outcomes correctly classified under positive class
TN (True Negative) i.e. the count of negative outcomes correctly classified under negative class
FP (False Positive) i.e. the count of negative outcomes incorrectly classified under positive class
FN (False Negative) i.e. the count of positive outcomes that are incorrectly classified under negative class

8 Performance Evaluation

This section presents the experimental methodology, dataset characteristics, and systematic assessment of DP techniques. The efficacy of the proposed differentially private data aggregation methods, each employing different noise characteristics for various ML models, is evaluated and analyzed[1].

8.1 Experimental Setting

Experiments were conducted in Python 3.9.6 using scikit-learn for ML algorithms, NumPy for numerical computation, and pandas for data processing. Differential privacy mechanisms were implemented via custom functions to control noise calibration and privacy budget allocation. All experiments used fixed random seeds for reproducibility. Security validation employed a publicly available medical dataset from the UCI repository, retrieved programmatically using the `ucimlrepo` package, a standard benchmark for privacy-preserving healthcare ML.

Multi-layer architecture performance was evaluated through discrete-event simulations incorporating stochastic models of network and computational behavior. Network delays were modeled as uniformly distributed variables calibrated against large-scale RTT measurements reported by Charyyev et al. [60], who measured ping latency from 8,456 end-users to 6,341 edge servers and 69 cloud locations. Their results show that 58% of users reach a nearby edge server in under 10 ms, with median edge RTT of approximately 5–10 ms for the majority of users, while cloud RTT typically ranges from 30–100 ms. Our simulation parameters—IoT-to-Edge (2–8 ms) and Edge-to-Cloud (40–80 ms)—fall within these empirically observed ranges, representing network-layer round-trip time for hospital-proximate edge deployments and well-provisioned WAN connections respectively. Data transfer latency was computed deterministically from payload size and bandwidth assumptions. Computational processing times, including ML inference were modeled as uniform distributions reflecting hardware capabilities at each tier. Blockchain consensus performance was simulated using Hyperledger Fabric's Raft protocol across varying topologies (4–13 nodes) to assess integrity verification overhead. Monte Carlo evaluation with 100 independent trials per configuration provided statistical confidence intervals for all performance metrics.

8.2 Data Set

The evaluation employed the UCI AIDS Clinical Trials Group Study 175 dataset, which comprises healthcare statistics from AIDS patients [61]. This dataset encompasses both medical parameters (CD4 counts at baseline and months 20, CD8 counts at baseline and month 20, prescribed medications) and demographic attributes (date, time, age, ethnicity, and gender) that constitute sensitive information whose disclosure could potentially result in privacy violations and social implications. The UCI AIDS dataset was selected for its appropriate dimensionality, containing 2,139 observations across 27 variables, of which 23 serve as input features for model training ($d = 23$). The dataset exhibits significant class imbalance with 75.7% of samples in class 0 (censored) and 24.3% in class 1 (failure), reflecting realistic clinical trial outcomes. To address this imbalance, we employed undersampling to create balanced training sets with equal representation (417 samples per class, 834 total), improving model convergence and preventing majority-class bias under DP noise. Table 4 reports the exact instance counts per class at each stage of the data pipeline to facilitate independent replication. The dataset was partitioned into 80% training and 20% test sets using stratified random splitting (random

1. Code and experiment artifacts are available at: https://github.com/Cloudslab/DP-Healthcare

seed 42) to preserve class proportions. No separate validation set was employed; hyperparameter selection was based on established defaults from the literature rather than data-driven tuning, ensuring that the full training set was available for DP noise calibration.

TABLE 4: Dataset Composition: Instance Counts per Class at Each Pipeline Stage

Stage	Class 0	Class 1	Total
Full dataset	1618 (75.6%)	521 (24.4%)	2139
Train (80%, stratified)	1294 (75.6%)	417 (24.4%)	1711
Test (20%, stratified)	324 (75.7%)	104 (24.3%)	428
Train after undersampling	417 (50.0%)	417 (50.0%)	834
Discarded (majority class)	877	—	877

8.3 Experiment

The dataset was processed using K-Means clustering, Logistic Regression, Random Forest, and Naive Bayes algorithms with three DP techniques applied: Laplace, Gaussian, and hybrid Laplace-Gaussian noise mechanisms. All combinations of ML algorithms and DP techniques were evaluated across varying epsilon values ($\varepsilon \in \{0.5, 1.0, 2.0, 3.0, 5.0, 10.0\}$) to quantify accuracy degradation following DP obfuscation and identify practical epsilon thresholds for healthcare analytics. This methodical approach determined the optimal noise distribution and epsilon configuration for protecting ML training datasets while preserving analytical utility.

The privacy budget parameters were extended to include higher epsilon values to investigate the practical epsilon threshold where input perturbation becomes viable for healthcare applications. The range spans from $\varepsilon = 0.5$ (strong privacy for highly sensitive HIV/AIDS patient data, aligning with recommendations for stigmatizing medical conditions [62]) to $\varepsilon = 10.0$ (moderate privacy, consistent with real-world deployments such as the U.S. Census Bureau [63]). This extended spectrum enables evaluation across different healthcare use cases: individual patient records (lower ε) to population-level analytics (higher ε), balancing re-identification risk with model utility as mandated by healthcare privacy regulations.

The DP implementation employed input perturbation, where calibrated noise is added to training data before model learning. This approach ensures privacy guarantees independent of the learning algorithm. Sensitivity was established through per-record norm clipping, a standard technique for bounding query sensitivity in DP [7]. Specifically, each record x_i in the standardised AIDS dataset was projected onto the L_2 ball of radius C via the clipping operation $\bar{x}_i = x_i \cdot \min(1, C/\|x_i\|_2)$, ensuring that $\|\bar{x}_i\|_2 \leq C$ for all records. The clipping threshold was set to the 95th percentile of the empirical per-record L_2 norm distribution ($C \approx 6.47$), a heuristic that balances bounded sensitivity against data distortion: lower thresholds increase clipping-induced bias while higher thresholds necessitate larger noise magnitudes. We note that the threshold selection itself is data-dependent; in a deployment setting, C should be determined from public domain knowledge or a held-out sample with separate privacy accounting.

L_2 clipping directly establishes the L_2 sensitivity $\Delta_2 = C$ required by the Gaussian mechanism. For the Laplace mechanism, which requires L_1 sensitivity, the Cauchy-Schwarz inequality yields the bound $\Delta_1 = \max \|\bar{x}_i\|_1 \leq C\sqrt{d}$, where d is the number of features used in model training ($d = 23$ for the standardised AIDS dataset). In our experimental implementation, Laplace noise was calibrated to the L_2 bound ($\lambda = C/\varepsilon$) rather than the L_1 bound ($\lambda = C\sqrt{d}/\varepsilon$). This calibration provides a relaxed privacy guarantee for the Laplace mechanism: the effective privacy parameter is $\varepsilon_{\text{eff}} = \varepsilon\sqrt{d} \approx 4.8\varepsilon$ rather than the nominal ε. Consequently, the Laplace results in Tables 6 and 7 should be interpreted at this effective privacy level (e.g., nominal $\varepsilon = 1.0$ corresponds to $\varepsilon_{\text{eff}} \approx 4.8$ for Laplace). The Gaussian mechanism results carry the stated (ε, δ)-DP guarantees exactly, and the hybrid mechanism inherits its guarantee from its component budgets under sequential composition. Relative accuracy comparisons between noise mechanisms at matched nominal ε remain valid for algorithm selection guidance, as all mechanisms received noise calibrated to the same L_2 bound. In our dataset, approximately 5% of records underwent clipping, introducing bounded distortion dominated by the subsequent DP noise at all tested privacy budgets. Gaussian noise standard deviation was set to $\sigma = C\sqrt{2\ln(1.25/\delta)}/\varepsilon$ with $\delta = 10^{-5}$, and the hybrid mechanism employed a weighted-noise combination: noise $= \alpha \cdot \text{Laplace}(C/\varepsilon) + (1-\alpha) \cdot \text{Gaussian}(\sigma)$ with $\alpha = 0.5$, ensuring distributional consistency across all data points and preserving algorithm-specific statistical assumptions [7].

The DP algorithm was executed with five independent runs to ensure statistical robustness, with mean performance reported. K-Means utilized two clusters (matching the binary classification task) determined through optimal cluster-to-class mapping using the Hungarian algorithm [64]. Random Forest employed a five-level decision tree architecture with 100 trees, minimum 20 samples per split, and 10 samples per leaf to prevent overfitting. Logistic Regression used L2 regularization ($C = 1.0$) with maximum 1000 iterations for convergence. Naive Bayes employed Variant B of Algorithm 4 (continuous input perturbation with Gaussian likelihood estimation) with default priors. Evaluation metrics include accuracy, precision, recall, F1-score, AUC, computed on a held-out test set (20% of data, 428 samples) using stratified splitting to preserve class distribution, as detailed in Tables 5, 6, 7, and 8.

8.4 Result Analysis

Table 5 presents baseline performance without DP, establishing reference metrics for privacy-utility trade-off analysis. Table 6 presents the accuracy of various supervised and unsupervised ML techniques across an extended epsilon range ($\varepsilon \in \{0.5, 1.0, 2.0, 3.0, 5.0, 10.0\}$) using Laplace, Gaussian, and weighted-combination hybrid noise distributions. Table 7 provides detailed precision, recall, F1-score, and AUC metrics for $\varepsilon = 10.0$, demonstrating near-baseline performance recovery at moderate privacy levels. These values are plotted graphically in Fig. 6, illustrating the accuracy of DP-enabled K-Means, Logistic Regression, Random Forest and Naive Bayes respectively. Additionally, Figs. 7 and 8 quantify privacy protection through adversarial evaluation, measuring attribute inference and data reconstruction attack success rates across privacy budgets.

(i) ML Algorithms: The selection of an ML algorithm depends on the specific use case and problem type. Whilst K-Means performs clustering tasks, the supervised algorithms (Naive Bayes, Random Forest, and Logistic Regression) handle classification and prediction problems with varying strengths. For instance, Random Forest excels at detailed multi-class

TABLE 5: Baseline Model Performance (No Differential Privacy)

Algorithm	Acc.	Prec.	Recall	F1	AUC
Logistic Regression	84.8	0.66	0.79	0.72	0.887
Random Forest	85.0	0.64	0.88	0.74	0.922
Naive Bayes	80.8	0.58	0.77	0.66	0.844
K-Means	54.9	0.31	0.70	0.43	0.650

TABLE 6: ε vs. Accuracy (%): Mean±Std over 5 Runs (Hybrid: $\alpha = 0.5$)

Noise Type/ ε Value	Laplace	Gaussian	Hybrid
K-Means			
0.5	53.8±10.5	42.1±23.0	37.6±18.0
1.0	51.2±5.7	38.6±14.8	45.1±21.1
2.0	56.4±3.8	40.0±18.6	53.6±14.8
3.0	58.4±6.6	42.8±20.1	59.1±9.2
5.0	58.6±2.8	51.4±14.0	53.7±14.5
10.0	55.1±0.1	60.6±8.3	57.9±1.2
Logistic Regression			
0.5	63.2±8.3	51.6±26.1	64.9±14.6
1.0	68.6±2.5	53.0±22.7	66.9±10.9
2.0	74.3±2.2	58.2±15.8	71.8±4.7
3.0	77.2±2.2	64.4±10.1	73.9±2.7
5.0	80.4±1.3	72.0±5.0	77.3±1.1
10.0	83.6±0.7	77.8±2.7	80.4±0.6
Random Forest			
0.5	57.6±20.3	52.4±26.3	50.3±23.9
1.0	58.4±20.0	47.1±22.0	54.7±25.4
2.0	79.2±1.9	54.1±25.5	54.7±17.5
3.0	81.7±1.9	48.4±18.7	66.9±13.4
5.0	81.2±2.5	56.4±26.5	72.8±8.6
10.0	83.2±1.9	63.0±20.2	79.9±2.8
Naive Bayes			
0.5	55.1±28.2	55.1±28.2	51.2±25.8
1.0	55.5±28.5	55.1±28.2	54.0±27.2
2.0	57.3±28.1	57.6±25.1	55.6±28.5
3.0	62.7±24.5	62.7±22.5	60.7±24.0
5.0	75.7±10.3	66.7±23.8	75.4±8.1
10.0	80.0±2.6	70.5±23.0	80.3±1.3

TABLE 7: Comprehensive Performance Metrics at $\varepsilon = 10.0$: Mean±Std over 5 Runs (Hybrid: $\alpha = 0.5$)

Algorithm/Noise	Acc.	Prec.	Recall	F1
Logistic Regression				
Laplace	83.6±0.7	0.63±0.02	0.79±0.02	0.70±0.01
Gaussian	77.8±2.7	0.53±0.04	0.74±0.03	0.62±0.03
Hybrid	80.4±0.6	0.57±0.01	0.77±0.04	0.65±0.01
Random Forest				
Laplace	83.2±1.9	0.63±0.04	0.78±0.02	0.69±0.03
Gaussian	63.0±20.2	0.43±0.15	0.71±0.17	0.50±0.08
Hybrid	79.9±2.8	0.58±0.05	0.71±0.10	0.63±0.05
Naive Bayes				
Laplace	80.0±2.6	0.57±0.04	0.74±0.05	0.64±0.02
Gaussian	70.5±23.0	0.56±0.18	0.63±0.25	0.53±0.11
Hybrid	80.3±1.3	0.58±0.03	0.69±0.03	0.63±0.01
K-Means				
Laplace	55.1±0.1	0.31±0.00	0.70±0.00	0.43±0.00
Gaussian	60.6±8.3	0.35±0.08	0.65±0.03	0.45±0.06
Hybrid	57.9±1.2	0.33±0.01	0.69±0.01	0.44±0.01

TABLE 8: Hybrid Mechanism: Impact of Budget Allocation Parameter α across $\varepsilon \in \{2.0, 5.0, 10.0\}$ (5-run mean accuracy %)

Algorithm	$\alpha = 0.3$	$\alpha = 0.5$	$\alpha = 0.7$	Best α
$\varepsilon = 2.0$				
Logistic Regression	69.1	71.8	74.4	0.7
Random Forest	48.8	54.7	48.8	0.5
Naive Bayes	55.0	55.6	61.6	0.7
K-Means	43.4	53.6	50.2	0.5
$\varepsilon = 5.0$				
Logistic Regression	74.9	77.3	78.7	0.7
Random Forest	70.4	72.8	78.3	0.7
Naive Bayes	66.1	75.4	78.2	0.7
K-Means	60.1	53.7	61.1	0.7
$\varepsilon = 10.0$				
Logistic Regression	79.1	80.4	82.0	0.7
Random Forest	75.3	79.9	82.6	0.7
Naive Bayes	79.4	80.3	80.8	0.7
K-Means	60.1	57.9	55.9	0.3

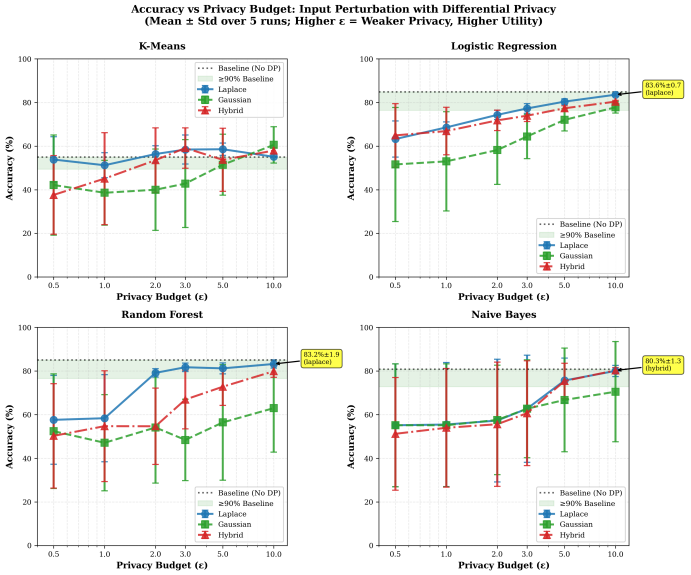

Fig. 6: Accuracy of ML algorithms across different noise mechanisms and varying privacy budgets. Horizontal dotted lines indicate baseline performance without DP. Green shaded regions show ≥90% baseline retention zone. Logarithmic x-axis emphasizes the wide privacy budget range tested.

classification, whilst Logistic Regression is particularly effective for binary classification and probability estimation.

As shown in Table 5, supervised methods significantly outperformed unsupervised clustering, with Logistic Regression achieving the highest baseline accuracy (84.8%, AUC: 0.887), followed closely by Random Forest (85.0%, AUC: 0.922) and Naive Bayes (80.8%, AUC: 0.844). K-Means clustering achieved 54.9% accuracy, reflecting the fundamental limitations of unsupervised learning for classification tasks.

Under DP with $\varepsilon = 10.0$ (Table 7), Laplace noise enabled near-baseline performance recovery across supervised algorithms. Logistic Regression achieved 83.6±0.7% accuracy (98.6% of baseline), followed by Random Forest at 83.2±1.9% (97.9% of baseline), both with low variance confirming reliable performance. Naive Bayes recovered to 80.3±1.3% with hybrid noise (99.4% of baseline). K-Means accuracy reached 60.6±8.3% with Gaussian noise (110.4% of baseline), though the high standard deviation reflects fundamental instability of centroid-based clustering under data perturbation. The elevated variance observed for Gaussian noise across K-Means, Naive Bayes, and Random Forest (std: 8.3–23.0%) contrasts with the stable Laplace results (std: 0.1–2.6%), indicating that Laplace noise produces more consistent outcomes across independent noise realisations.

(ii) Effect of Privacy Budget (ε): The epsilon parameter (ε) governs the fundamental trade-off between privacy protection and data utility. Smaller epsilon values ($\varepsilon < 1$) provide stronger privacy guarantees but result in larger accuracy losses, while larger values ($\varepsilon > 1$) provide better accuracy with reduced privacy guarantees. This relationship arises because Laplace and Gaussian mechanisms add noise with scale parameters inversely proportional to epsilon and directly

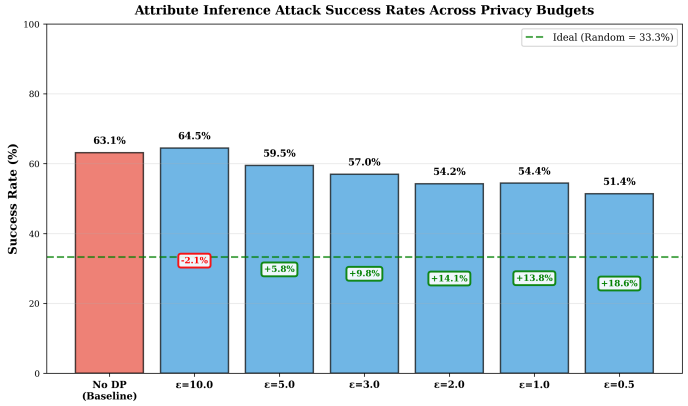

Fig. 7: Attribute inference attack success rates across privacy budgets. Results averaged across all algorithms. Lower values indicate better privacy protection, with 33.3% representing random guessing for the three-class sensitive attribute (ideal privacy).

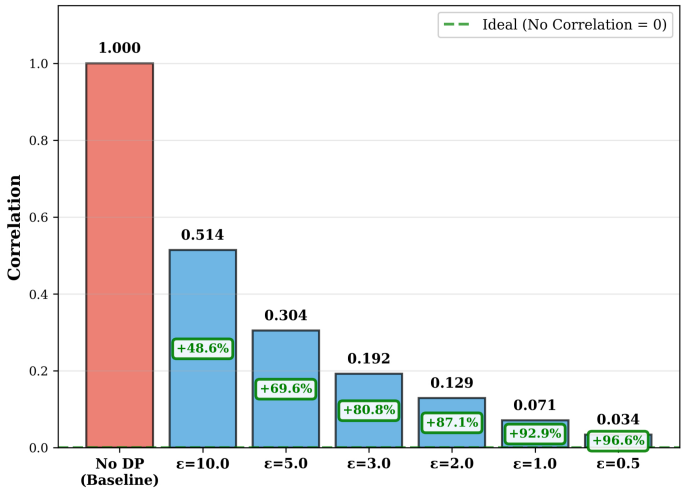

Fig. 8: Data reconstruction attack correlation across privacy budgets. Results averaged across all algorithms. Lower correlation values indicate better privacy protection, with 0 representing no reconstruction capability (ideal privacy). The baseline bar is clipped with hatching to improve readability.

proportional to query sensitivity.

As illustrated in Table 6 and Fig. 6, the extended epsilon range reveals clear practical thresholds for input perturbation viability. For Random Forest with Laplace noise, mean accuracy increases from 57.6±20.3% at $\varepsilon = 0.5$ to 83.2±1.9% at $\varepsilon = 10.0$, with variance decreasing substantially as the privacy budget increases. The critical transition occurs at $\varepsilon \geq 2.0$, where accuracy stabilises above 79% with std <2%. Logistic Regression exhibits the most consistent Laplace performance: 63.2±8.3% at $\varepsilon = 0.5$ to 83.6±0.7% at $\varepsilon = 10.0$, achieving practical utility (>80%) at $\varepsilon \geq 5.0$. Naive Bayes displays high variance at low ε (std: 25–28%), reflecting sensitivity to noise in distributional parameter estimation, before stabilising at 80.0±2.6% by $\varepsilon = 10.0$. K-Means shows persistent instability across all epsilon values, with standard deviations of 1–23% reflecting fundamental limitations of centroid-based clustering under input perturbation.

These results establish $\varepsilon \geq 5.0$ as the practical epsilon threshold for input perturbation on healthcare datasets of this scale (∼800 training samples, 23 features). At this threshold, supervised algorithms achieve 80–81% accuracy (94–96% baseline retention), balancing moderate privacy protection with high analytical utility suitable for population-level clinical research and regulatory reporting.

(iii) Noise Mechanism Comparison: Table 6 reveals that algorithm-noise interactions differ substantially across mechanisms. Laplace noise demonstrates superior and more consistent performance for supervised methods: Logistic Regression achieves 83.6±0.7% at $\varepsilon = 10.0$ (best overall), and Random Forest reaches 83.2±1.9%, both with low variance. This stems from Laplace's heavier tails better preserving decision boundary information. Gaussian noise exhibits substantially higher variance across runs (std: 2.7–23.0% for supervised algorithms at $\varepsilon = 10.0$), suggesting sensitivity to specific noise realisations. For K-Means, Gaussian achieves 60.6±8.3% at $\varepsilon = 10.0$ compared to Laplace's 55.1±0.1%, though this advantage comes with considerably higher variance.

The weighted-noise hybrid mechanism achieves performance intermediate between pure Laplace and Gaussian with moderate variance. At $\varepsilon = 10.0$, hybrid yields 80.4±0.6% (Logistic Regression), 79.9±2.8% (Random Forest), and 80.3±1.3% (Naive Bayes), consistently falling between the pure mechanisms. The low standard deviations confirm that the weighted-noise approach provides reliable, reproducible results across independent noise realisations. Table 8 presents an ablation study of the budget allocation parameter α across $\varepsilon \in \{2.0, 5.0, 10.0\}$, evaluated over 5 independent runs. The results reveal a consistent pattern: Laplace-dominant allocation ($\alpha = 0.7$) is optimal for all three supervised algorithms at $\varepsilon \geq 5.0$, with the advantage strengthening at higher privacy budgets. At $\varepsilon = 10.0$, $\alpha = 0.7$ yields 82.0% (Logistic Regression), 82.6% (Random Forest), and 80.8% (Naive Bayes). At $\varepsilon = 2.0$, the pattern is less stable due to high noise magnitude, with Random Forest and K-Means preferring $\alpha = 0.5$ and the supervised algorithms showing narrower margins between allocations. K-Means exhibits an opposite preference toward Gaussian-dominant allocation ($\alpha = 0.3$) at $\varepsilon = 10.0$, consistent with Gaussian noise's compatibility with centroid-based distance calculations. These results demonstrate that optimal α is indeed ε-dependent, supporting the adaptive allocation rationale: at low ε where noise dominates, the allocation choice has limited impact, while at moderate-to-high ε where the signal-to-noise ratio permits meaningful model learning, Laplace-dominant allocation consistently benefits supervised algorithms. For practical deployment, $\alpha = 0.7$ is recommended for supervised learning at $\varepsilon \geq 5.0$, with $\alpha = 0.5$ as a conservative default when the privacy budget is uncertain.

(iv) ε Threshold Analysis for Healthcare Applications: The extended epsilon range enables identification of use-case-specific deployment recommendations. For **individual patient analytics** requiring strong privacy ($\varepsilon \in \{0.5, 1.0\}$): all algorithms exhibit high variance (std: 2.5–28.5%), with Logistic Regression Laplace achieving the most reliable performance (68.6±2.5% at $\varepsilon = 1.0$). The elevated variance at low ε indicates that input perturbation at this privacy level produces unreliable results on datasets of this scale, suggesting the need for either larger datasets ($n > 5000$) or alternative approaches such as output perturbation. For **cohort-level research** with moderate privacy ($\varepsilon \in \{2.0, 3.0\}$): Random Forest Laplace achieves 81.7±1.9% and Logistic Regression 77.2±2.2% at $\varepsilon = 3.0$ with low variance, providing acceptable

utility for multi-site clinical trials with informed consent. For **population-level analytics** with regulatory-compliant privacy ($\varepsilon \in \{5.0, 10.0\}$): supervised algorithms with Laplace noise achieve 80–84% accuracy with std <2.5%, enabling epidemiological studies, public health surveillance, and quality improvement initiatives while maintaining census-level privacy standards [63].

(v) Privacy Protection Validation: To empirically validate the privacy guarantees of the DP mechanisms, we evaluate two complementary adversarial attacks. The *attribute inference attack* follows the model-assisted paradigm: an adversary who observes the target model's predictions attempts to infer a sensitive feature value for records in the test set. Specifically, the adversary trains a Random Forest classifier (50 trees, balanced class weights) on the first half of the training data, using as input features all non-sensitive attributes concatenated with the target model's predicted class label and class probability vector. The sensitive attribute (feature index 0, representing the first clinical variable) is discretised into three ordinal classes using the 33rd and 67th percentile thresholds, and the adversary's goal is to predict the correct bin. Attack success is measured as classification accuracy on the held-out test set, with 33.3% representing random guessing for three classes. The *data reconstruction attack* quantifies how closely an adversary can recover the original training data from the DP-perturbed version by computing the mean Pearson correlation coefficient across all features between the original and perturbed datasets, where 1.0 indicates perfect reconstruction and 0.0 indicates no recoverable signal.

Figures 7 and 8 quantify empirical privacy protection across the extended epsilon range. Attribute inference attack success (Fig. 7) decreases from 63.1% (no DP baseline) to 51.4-64.5% under DP ($\varepsilon \in \{0.5, 10.0\}$), with strongest protection at $\varepsilon = 0.5$ (51.4%, approaching the random guessing baseline of 33.3% for the three-class sensitive attribute). Importantly, even at the highest tested epsilon ($\varepsilon = 10.0$), attack success remains suppressed (64.5%), representing 2.1% improvement over baseline and confirming that moderate privacy budgets still provide meaningful protection against sensitive attribute disclosure. Attack success shows weak epsilon-dependence across the range, varying only 13.1 percentage points, suggesting that inference vulnerability saturates quickly and that privacy gains beyond $\varepsilon = 5.0$ provide diminishing returns for this attack type.

Data reconstruction correlation (Fig. 8) demonstrates dramatic privacy gains across all epsilon values: baseline correlation of 1.000 (perfect reconstruction) decreases to 0.034-0.514 under DP, representing 48.6-96.6% privacy improvement. Reconstruction protection strengthens monotonically with decreasing epsilon (0.514 at $\varepsilon = 10.0$ to 0.034 at $\varepsilon = 0.5$), confirming theoretical guarantees. At the recommended threshold $\varepsilon = 5.0$, reconstruction correlation of 0.304 represents 69.6% improvement, effectively preventing individual-level data recovery while enabling the near-baseline model utility (80–81% accuracy) required for practical healthcare analytics. These results validate that input perturbation provides robust empirical privacy protection against realistic adversaries across the full tested epsilon spectrum, with optimal privacy-utility balance achieved at $\varepsilon \in \{5.0, 10.0\}$ for population-level healthcare applications.

8.5 Architectural Performance Evaluation

To validate the proposed multi-layer architecture's effectiveness for time-critical healthcare operations, we conducted simulations measuring latency and throughput across IoT-Edge-Cloud layers, alongside blockchain consensus performance for data integrity verification.

8.5.1 Edge-Cloud Latency Analysis

Table 9 presents end-to-end latency measurements for representative healthcare scenarios across Edge and Cloud processing layers. Edge computing demonstrates substantial latency advantages for time-critical operations, achieving 8.0× speedup for emergency response scenarios (15.4ms vs 123.2ms) and 7.2× for vital signs monitoring (26.8ms vs 193.5ms). This performance differential stems from reduced network propagation delay and localized processing, validating the architectural decision to deploy latency-sensitive operations at Edge nodes.

For data-intensive operations such as radiological imaging, the speedup factor decreases to 3.0× (158.7ms vs 476.3ms for high-resolution images), as Cloud infrastructure's superior computational resources partially offset network latency penalties. These results confirm that the proposed hierarchical task distribution—emergency response and vital monitoring at Edge, batch analytics at Cloud—optimally balances response time requirements with computational capacity constraints.

TABLE 9: Edge vs Cloud Processing Latency Comparison

Scenario	Data Size (KB)	Edge (ms)	Cloud (ms)	Speedup Factor
Emergency Response	1	15.4	123.2	8.0×
Vital Signs Monitoring	10	26.8	193.5	7.2×
ECG Analysis	50	42.9	248.1	5.8×
Diagnostic Imaging	256	85.6	356.4	4.2×
Radiological Imaging	1024	158.7	476.3	3.0×
Batch Analytics	5120	421.3	897.6	2.1×

System throughput evaluation under varying load conditions (Table 10) demonstrates that Edge infrastructure sustains 186.5 requests/second for light workloads (5 concurrent users, 10KB data) but experiences degradation to 46.6 req/s under heavy load (20 users, 50KB data). Cloud infrastructure exhibits superior scalability, maintaining 186.9 req/s under light load and 105.0 req/s under heavy load (50 concurrent users), confirming its suitability for high-throughput batch analytics. The 2.25× throughput advantage under heavy load validates the architectural partitioning of sporadic high-priority requests to Edge nodes while directing sustained analytical workloads to Cloud infrastructure. The 75% Edge throughput degradation under heavy load warrants consideration for real-world emergency healthcare deployments where concurrent monitoring of 20 or more patients per Edge node is realistic. However, emergency response payloads are lightweight (1 KB, Table 9) compared to the 50 KB payloads used in the heavy-load simulation, and critical alerts are inherently sporadic rather than sustained. Charyyev et al. [60] similarly observe that edge server capacity limitations under high workloads necessitate either non-uniform server provisioning or cloud offloading to avoid over-subscription. For deployments requiring higher concurrent capacity, horizontal scaling through multiple Edge nodes or priority queuing that pre-empts non-critical requests during emergency events would preserve sub-20 ms response times.

TABLE 10: System Throughput Under Different Load Conditions

Scenario	Concurrent Users	Avg Latency (ms)	Throughput (req/s)
Edge - Light Load	5	26.8	186.5
Edge - Heavy Load	20	42.9	46.6
Cloud - Light Load	10	53.5	186.9
Cloud - Heavy Load	50	47.6	105.0

8.5.2 Blockchain Consensus Performance

The blockchain component employing Hyperledger Fabric with Raft consensus protocol was evaluated across varying network configurations (Table 11). Transaction finality latency for a 4-node network averaged 144.8ms (std: 45.2ms), supporting 2068 transactions per second (TPS), sufficient for non-emergency audit trail and privacy budget tracking requirements. As network size increased to 13 nodes, average latency rose to 284.9ms with throughput decreasing to 1140 TPS, reflecting the inherent scalability-latency trade-off in consensus protocols.

The measured latencies align with healthcare data integrity requirements where sub-second finality suffices for access logging, data provenance tracking, and privacy budget ledger updates. For emergency response scenarios requiring ¡20ms latency (Table 9), the architecture appropriately bypasses blockchain verification, deferring cryptographic commitment to post-emergency audit phases. This design ensures that data integrity mechanisms do not compromise critical response times while maintaining comprehensive auditability for regulatory compliance.

TABLE 11: Blockchain Consensus Performance (Raft Protocol)

Nodes	Avg Latency (ms)	Std Dev (ms)	Min (ms)	Max (ms)	Throughput (TPS)
4	144.8	45.2	72.4	282.1	2068
7	209.7	66.9	94.5	414.3	1503
10	250.3	78.4	113.8	512.7	1278
13	284.9	88.1	128.6	598.4	1140

Component-level analysis reveals that consensus protocol overhead (log replication and acknowledgment) constitutes 35-42% of total transaction latency across all configurations, with block propagation and validation contributing 28-33% and 20-25% respectively. These measurements validate the architectural decision to employ Raft consensus for healthcare applications, as its deterministic finality and moderate throughput requirements align well with audit trail use cases while avoiding the computational expense of proof-of-work mechanisms unsuitable for resource-constrained healthcare environments.

9 Threats to Validity

While the proposed framework demonstrates promising results, several limitations warrant acknowledgment. First, the blockchain integration, though architecturally specified, introduces inherent scalability constraints: as the number of consensus nodes increases from 4 to 13, transaction throughput decreases from 2068 to 1140 TPS and average latency rises from 144.8ms to 284.9ms (Table 11), which may pose challenges for large-scale multi-institutional deployments with high transaction volumes. Second, the DP mechanisms impose a measurable impact on model accuracy, particularly at lower privacy budgets; at $\varepsilon \leq 1.0$, supervised algorithms experience significant accuracy degradation (e.g., Logistic Regression drops to 57.0–63.1%), limiting the applicability of strong privacy guarantees for individual-level patient analytics on datasets of this scale. Third, edge computing resources, while effective for latency reduction in emergency scenarios, face throughput degradation under heavy load (from 186.5 to 46.6 req/s), and possess limited computational capacity for executing sophisticated ML models or complex DP mechanisms, potentially constraining real-time privacy-preserving analytics at the edge layer. Fourth, the experimental evaluation relies on a single healthcare dataset (UCI AIDS Clinical Trials) with discrete-event simulations for architectural performance, and further validation on diverse clinical datasets and real-world deployments is needed to establish broader generalisability. Fifth, the Laplace noise calibration uses the L_2 clipping bound rather than the L_1 sensitivity bound required by the Laplace mechanism, resulting in an effective privacy parameter of $\varepsilon_{\text{eff}} = \varepsilon\sqrt{d} \approx 4.8\varepsilon$ ($d = 23$) for Laplace results (Section 8); the Gaussian mechanism results carry the stated (ε, δ)-DP guarantees exactly. Additionally, the clipping threshold ($C \approx 6.47$) was determined from the training data itself; a production deployment should derive this bound from domain knowledge or a held-out sample with separate privacy accounting to preserve formal DP guarantees. Sixth, the experimental evaluation employed balanced undersampling (417 samples per class) to address the dataset's inherent class imbalance (75.6% censored, 24.4% failure). While this strategy improves model convergence and prevents majority-class bias under DP noise, it may not reflect realistic clinical deployment conditions where censored outcomes vastly dominate. These constraints represent important considerations for practitioners seeking to deploy the framework in production healthcare environments.

10 Conclusions and Future Directions

Contemporary Internet of Things (IoT) and Cloud computing frameworks have facilitated advanced medical assistance and enabled novel insights in healthcare analytics. This investigation presents a comprehensive architectural framework that addresses critical response time optimization and security requirements in modern healthcare systems.

Our multilayered architectural approach integrating IoT, Edge, and Cloud components facilitates expeditious response capabilities for emergency medical scenarios through strategic workload distribution across computational layers, with data aggregation and analytical processes performed within Cloud nodes to maximise computational efficiency whilst maintaining system responsiveness.

To establish robust security provisions, the framework implements a dual mechanism approach incorporating DP and Blockchain technology. The DP methodology preserves individual patient confidentiality during both data storage operations and ML analytical processes. Experimental findings demonstrated that ML query accuracy maintained statistical integrity following the application of Laplace noise, Gaussian noise, and combined Laplace-Gaussian noise. Analytical utility and algorithmic accuracy remained statistically comparable across all evaluated DP implementations. The architectural design utilised Blockchain technology to enhance system reliability and establish cryptographic protection for transactional data and persistent storage mechanisms. This integration of DP mechanisms with ML methodologies, com-

bined with blockchain-based data integrity, provides a foundation for secure and ethically sound data-driven healthcare analytics.

The primary actionable finding of this work is the identification of $\varepsilon = 5.0$ as the practical deployment threshold for differentially private healthcare analytics. At this privacy budget, supervised algorithms with Laplace noise achieve 80–81% accuracy (94–96% of baseline) with low variance across independent runs, while attribute inference attack success is reduced and data reconstruction correlation drops by approximately 70%. For healthcare practitioners seeking to balance patient privacy with analytical utility, we recommend $\varepsilon \geq 5.0$ with Laplace noise for population-level analytics (epidemiological studies, public health surveillance, quality improvement), and Laplace-dominant hybrid allocation ($\alpha = 0.7$) when noise mechanism properties are uncertain. Below $\varepsilon = 2.0$, input perturbation on datasets of this scale produces unreliable results with high variance, and alternative approaches such as output perturbation or larger training datasets should be considered. Overall, this work has successfully designed and validated a healthcare system architecture that simultaneously addresses the critical requirements of enhanced responsiveness and secure, privacy-preserved analytics.

Future research directions include practical deployment of the blockchain infrastructure with healthcare-specific consensus mechanisms, integration of zero-knowledge proofs for privacy-preserving audit trails, and cross-chain interoperability protocols for multi-institutional data sharing. Extending the DP framework to deep neural networks and federated learning architectures, alongside integration with complementary privacy-enhancing technologies such as secure multi-party computation and homomorphic encryption, would strengthen privacy guarantees whilst enabling broader analytical capabilities. Developing adaptive privacy budget allocation mechanisms that dynamically adjust based on data sensitivity and query patterns represents another promising avenue. Real-world pilot studies in clinical settings, development of user-friendly interfaces for healthcare practitioners, and integration with existing Electronic Health Record systems are essential steps toward production deployment. These directions aim to advance the proposed framework toward a comprehensive, production-ready privacy-preserving healthcare analytics platform that addresses security, privacy, performance, and regulatory requirements whilst maintaining clinical utility for evidence-based medical decision-making.

REFERENCES

[1] L. M. Dang, M. J. Piran, D. Han, K. Min, and H. Moon, "A survey on internet of things and cloud computing for healthcare," *Electronics*, 2019.

[2] N. Mangala, K. R. Venugopal, and B. E. Reddy, "Short paper: Current challenges in iot cloud smart applications," *Proceedings - 2021 IEEE International Conference on Cloud Computing in Emerging Markets, CCEM 2021*, pp. 36–40, 2021.

[3] K. Azbeg, O. Ouchetto, and S. J. Andaloussi, "Access control and privacy-preserving blockchain-based system for diseases management," *IEEE Transactions on Computational Social Systems*, vol. 10, pp. 1515–1527, 8 2023.

[4] D. of Health and H. S. USA, "Privacy, security, and electronic health records."

[5] "Summary of the hipaa privacy rule — hhs.gov."

[6] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council (general data protection regulation)," 2016.

[7] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and trends in theoretical computer science*, vol. 9, pp. 211–407, 2014.

[8] L. Sweeney, "Foundations of privacy protection from a computer science perspective."

[9] A. Qayyum, J. Qadir, M. Bilal, and A. Al-Fuqaha, "Secure and robust machine learning for healthcare: A survey," *IEEE Reviews in Biomedical Engineering*, vol. 14, pp. 156–180, 2021.

[10] A. Wood, K. Najarian, and D. Kahrobaei, "Homomorphic encryption for machine learning in medicine and bioinformatics," *ACM Computing Surveys*, vol. 53, 7 2021.

[11] B. Knott, S. Venkataraman, A. Hannun, S. Sengupta, M. Ibrahim, and L. van der Maaten, "Crypten: Secure multi-party computation meets machine learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 4961–4973, 12 2021.

[12] N. Mangala, E. B. Reddy, and K. Venugopal, "Light weight circular error learning algorithm (cela) for secure data communication protocol in iot-cloud systems," *IJACSA) International Journal of Advanced Computer Science and Applications*, vol. 14, 2023.

[13] B. Jia, X. Zhang, J. Liu, Y. Zhang, K. Huang, and Y. Liang, "Blockchain-enabled federated learning data protection aggregation scheme with differential privacy and homomorphic encryption in iiot," *IEEE Transactions on Industrial Informatics*, vol. 18, pp. 4049–4058, 6 2022.

[14] K. Gai, Y. Wu, L. Zhu, Z. Zhang, and M. Qiu, "Differential privacy-based blockchain for industrial internet-of-things," *IEEE Transactions on Industrial Informatics*, vol. 16, pp. 4156–4165, 6 2020.

[15] M. Rangwala, K. Venugopal, and R. Buyya, "Blockchain-enabled federated learning," 2025.

[16] H. Wu, A. D. Dwivedi, and G. Srivastava, "Security and privacy of patient information in medical systems based on blockchain technology," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 17, 6 2021.

[17] J. Zhang, Y. Yang, X. Liu, and J. Ma, "An efficient blockchain-based hierarchical data sharing for healthcare internet of things," *IEEE Transactions on Industrial Informatics*, vol. 18, pp. 7139–7150, 10 2022.

[18] M. Wazid, A. K. Das, and S. Shetty, "Bsfr-sh: Blockchain-enabled security framework against ransomware attacks for smart healthcare," *IEEE Transactions on Consumer Electronics*, vol. 69, pp. 18–28, 2 2023.

[19] D. W. Hosmer, S. Lemeshow, and R. X. Sturdivant, *Applied Logistic Regression*. Wiley, 2013.

[20] L. Breiman, "Random forests," *Machine Learning*, vol. 45, pp. 5–32, 10 2001.

[21] D. J. Hand and K. Yu, "Idiot's bayes - not so stupid after all?," *International Statistical Review*, vol. 69, pp. 385–398, 12 2001.

[22] M. JB., "Some methods of classification and analysis of multivariate observations," *Proc. of 5th Berkeley Symposium on Math. Stat. and Prob.*, pp. 281–297, 1967.

[23] J. Soria-Comas, J. Domingo-Ferrer, D. Sanchez, and D. Megias, "Individual differential privacy: A utility-preserving formulation of differential privacy guarantees," *IEEE Transactions on Information Forensics and Security*, vol. 12, pp. 1418–1429, 6 2017.

[24] S. Asoodeh, J. Liao, F. P. Calmon, O. Kosut, and L. Sankar, "Three variants of differential privacy: Lossless conversion and applications," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, pp. 208–222, 3 2021.

[25] M. Abadi, H. B. McMahan, A. Chu, I. Mironov, L. Zhang, I. Goodfellow, and K. Talwar, "Deep learning with differential privacy," *Proceedings of the ACM Conference on Computer and Communications Security*, vol. 24-28-October-2016, pp. 308–318, 10 2016.

[26] T. Zhang, T. Zhu, P. Xiong, H. Huo, Z. Tari, and W. Zhou, "Correlated differential privacy: Feature selection in machine learning," *IEEE Transactions on Industrial Informatics*, vol. 16, pp. 2115–2124, 3 2020.

[27] Y. Sun, J. Liu, K. Yu, M. Alazab, and K. Lin, "Pmrss: Privacy-preserving medical record searching scheme for intelligent diagnosis in iot healthcare," *IEEE Transactions on Industrial Informatics*, vol. 18, pp. 1981–1990, 3 2022.

[28] Y. Jiang, X. Xu, and F. Xiao, "Attribute-based encryption with blockchain protection scheme for electronic health records," *IEEE Transactions on Network and Service Management*, vol. 19, pp. 3884–3895, 12 2022.

[29] J. Lin, J. Niu, X. Liu, and M. Guizani, "Protecting your shopping preference with differential privacy," *IEEE Transactions on Mobile Computing*, vol. 20, pp. 1965–1978, 5 2021.

[30] H. Jiang, J. Pei, D. Yu, J. Yu, B. Gong, and X. Cheng, "Applications of differential privacy in social network analysis: A survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, pp. 108–127, 1 2023.

[31] Y. Lindell and E. Omri, "A practical application of differential privacy to personalized online advertising," *Cryptology ePrint Archive*, 2011.

[32] Z. Cai, X. Zheng, and J. Yu, "A differential-private framework for urban traffic flows estimation via taxi companies," *IEEE Transactions on Industrial Informatics*, vol. 15, pp. 6492–6499, 12 2019.

[33] C. Zhan, S. Joksimovic, D. Ladjal, T. Rakotoarivelo, R. Marshall, and A. Pardo, "Preserving both privacy and utility in learning analytics," *IEEE Transactions on Learning Technologies*, vol. 17, pp. 1655–1667, 2024.

[34] S. Saeidian, G. Cervia, T. J. Oechtering, and M. Skoglund, "Quantifying membership privacy via information leakage," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 3096–3108, 2021.

[35] Úlfar Erlingsson, V. Pihur, and A. Korolova, "Rappor: Randomized aggregatable privacy-preserving ordinal response," *Proceedings of the ACM Conference on Computer and Communications Security*, pp. 1054–1067, 11 2014.

[36] "Enabling developers and organizations to use differential privacy - google developers blog."

[37] A. Aktay, S. Bavadekar, G. Cossoul, J. Davis, *et al.*, "Google covid-19 community mobility reports: Anonymization process description (version 1.1)," 4 2020.

[38] L. Rachakonda, A. K. Bapatla, S. P. Mohanty, and E. Kougianos, "Sayopillow: Blockchain-integrated privacy-assured iomt framework for stress management considering sleeping habits," *IEEE Transactions on Consumer Electronics*, vol. 67, pp. 20–29, 2 2021.

[39] H. A. Kumar, J. Rakshith, R. Shetty, S. Roy, and D. Sitaram, "Comparison of iot architectures using a smart city benchmark," *Procedia Computer Science*, vol. 171, pp. 1507–1516, 1 2020.

[40] S. Yan, L. He, J. Seo, and M. Lin, "Concurrent healthcare data processing and storage framework using deep-learning in distributed cloud computing environment," *IEEE Transactions on Industrial Informatics*, vol. 17, pp. 2794–2801, 4 2021.

[41] M. Sun and W. P. Tay, "On the relationship between inference and data privacy in decentralized iot networks," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 852–866, 2020.

[42] R. Hu, Y. Guo, H. Li, Q. Pei, and Y. Gong, "Personalized federated learning with differential privacy," *IEEE Internet of Things Journal*, vol. 7, pp. 9530–9539, 10 2020.

[43] C. Cai, Y. Sang, and H. Tian, "A multimodal differential privacy framework based on fusion representation learning," *Connection Science*, vol. 34, pp. 2219–2239, 12 2022.

[44] A. R. Chowdhury, C. Wang, X. He, A. MacHanavajjhala, and S. Jha, "Crypte: Crypto-assisted differential privacy on untrusted servers," *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pp. 603–619, 6 2020.

[45] C. Dwork, "Differential privacy: A survey of results," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 4978 LNCS, pp. 1–19, 2008.

[46] H. Bi, J. Liu, and N. Kato, "Deep learning-based privacy preservation and data analytics for iot enabled healthcare," *IEEE Transactions on Industrial Informatics*, vol. 18, pp. 4798–4807, 7 2022.

[47] Z. Huang, R. Hu, Y. Guo, E. Chan-Tin, and Y. Gong, "Dp-admm: Admm-based distributed learning with differential privacy," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1002–1012, 2020.

[48] C. Li, A. He, Y. Wen, G. Liu, and A. T. Chronopoulos, "Optimal trading mechanism based on differential privacy protection and stackelberg game in big data market," *IEEE Transactions on Services Computing*, vol. 16, pp. 3550–3563, 9 2023.

[49] M. Zhang, J. Zhou, G. Zhang, L. Cui, T. Gao, and S. Yu, "Apdp: Attribute-based personalized differential privacy data publishing scheme for social networks," *IEEE Transactions on Network Science and Engineering*, vol. 10, pp. 922–933, 3 2023.

[50] J. Zhang and C. Zhong, "Differential privacy-based double auction for data market in blockchain-enhanced internet of things," *Wireless Communications and Mobile Computing*, vol. 2022, p. 8038846, 1 2022.

[51] W. Ali, M. Nauman, and N. Azam, "A privacy enhancing model for internet of things using three-way decisions and differential privacy," *Computers and Electrical Engineering*, vol. 100, p. 107894, 5 2022.

[52] M. Du, K. Wang, Z. Xia, and Y. Zhang, "Differential privacy preserving of training model in wireless big data with edge computing," *IEEE Transactions on Big Data*, vol. 6, pp. 283–295, 4 2018.

[53] M. Rangwala and R. Buyya, "Trustmesh: A blockchain-enabled trusted distributed computing framework for open heterogeneous iot environments," *Proceedings - 2025 IEEE 22nd International Conference on Software Architecture, ICSA 2025*, pp. 131–141, 2025.

[54] E. Androulaki, A. Barger, V. Bortnikov, S. Muralidharan, C. Cachin, *et al.*, "Hyperledger fabric: A distributed operating system for permissioned blockchains," *Proceedings of the 13th EuroSys Conference, EuroSys 2018*, vol. 2018-January, 4 2018.

[55] C. Stamatellis, P. Papadopoulos, N. Pitropakis, S. Katsikas, and W. J. Buchanan, "A privacy-preserving healthcare framework using hyperledger fabric," *Sensors (Basel, Switzerland)*, vol. 20, pp. 1–14, 11 2020.

[56] D. Ongaro and J. Ousterhout, *In Search of an Understandable Consensus Algorithm*. Usenix, 2014.

[57] W. Mehuron, "Public law (104-106), and the computer security act of 1996," *Technology Management Reform Act*, 1996.

[58] N. I. of Standards, T. (NIST), and E. Barker, "Secure hash standard (shs)," 1993.

[59] J. A. Marron, "Nist special publication 800 nist sp 800-66r2 implementing the health insurance portability and accountability act (hipaa) security rule a cybersecurity resource guide."

[60] B. Charyyev, E. Arslan, and M. H. Gunes, "Latency comparison of cloud datacenters and edge servers," in *GLOBECOM 2020-2020 IEEE Global Communications Conference*, pp. 1–6, IEEE, 2020.

[61] S. M. Hammer, D. A. Katzenstein, M. D. Hughes, H. Gundacker, R. T. Schooley, *et al.*, "A trial comparing nucleoside monotherapy with combination therapy in hiv-infected adults with cd4 cell counts from 200 to 500 per cubic millimeter. aids clinical trials group study 175 study team.," *New England Journal of Medicine*, vol. 335, pp. 1081–1090, 10 1996.

[62] F. K. Dankar and K. E. Emam, "Practicing differential privacy in health care: A review," *TRANSACTIONS ON DATA PRIVACY*, vol. 5, pp. 35–67, 2013.

[63] J. M. Abowd, "The u.s. census bureau adopts differential privacy," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, (New York, NY, USA), p. 2867, Association for Computing Machinery, 2018.

[64] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, pp. 83–97, 3 1955.